Technical Analysis, Central Bank Interventions and Exchange Rate Regimes

Dissertation

zur Erlangung des akademischen Grades eines Doktors der Wirtschafts- und Sozialwissenschaften (Dr. rer. pol.) des Fachbereichs Wirtschaftswissenschaften der Universität Hamburg

vorgelegt von
Dipl.-Volkswirt Peter Saacke MA
aus Köln

Hamburg, den 03. Mai 1999

Mitglieder der Promotionskommission

Vorsitzender: Prof. Dr. M. Holler
Erstgutachter: Prof. Dr. M. Funke
Zweitgutachter: Prof. Dr. B. Lucke

Das wissenschaftliche Gespräch fand am 06. August 1999 statt.

Technical Analysis, Central Bank Interventions and Exchange Rate Regimes

by

Peter Saacke

Diplom-Volkswirt, University of Freiburg (1995)
B.A., Mathematics and Philosophy, University of Oxford (1991)

Submitted to the Department of Economics
in partial fulfillment of the requirements for the degree of

Doctor Rerum Politicarum

at the

UNIVERSITY OF HAMBURG

May 1999

Signature of Author ....
Department of Economics
3.5.1999

Accepted by ....

Technical Analysis, Central Bank Interventions and Exchange Rate Regimes

by

Peter Saacke

Submitted to the Department of Economics
on 3.5.1999, in partial fulfillment of the
requirements for the degree of
Doctor Rerum Politicarum

Abstract

This investigation studies the usefulness of technical trading rules on foreign exchange markets and the way it depends on central bank interventions and exchange rate regimes. The first objective is to carry out a systematic analysis of the statistical and the economic significance of returns from following technical trading rules for flexible exchange rates. We find strong evidence for technical trading rule profitability and show that for several ways of adjusting trading rule returns for risk, technical trading rules tend to be less, rather than more risky than buying and holding currencies or stock market indices. The second objective of this investigation is to clarify the relationship between technical trading rule profitability and central bank interventions. Using daily data on foreign exchange interventions of both the Bundesbank and the Fed we provide further evidence that central banks earn profits with interventions and that technical trading rules are abnormally profitable on days on which interventions take place. We find that this seeming contradiction is due to the fact that whilst in the short run exchange rates move in a direction contrary to the central banks' intentions, in the long run this effect reverses. Moreover, we argue that the reason why technical trading rules are so profitable on days when central banks intervene is that since interventions are aimed at turning technically motivated trends (or, more metaphorically, are aimed at pricking speculative bubbles) they are a more reliable signal of the existence of a trend than a technical trading rule by itself. The third and final objective is to examine the influence of the participation in a fixed exchange rate regime on the usefulness of technical trading rules. We find that technical trading rules are unprofitable for exchange rates belonging to the European Monetary System (EMS) and that this was the case even during the early 1980s, when the fixed exchange rate parities in the EMS were non-credible and subject to frequent realignments. Moreover, examining the cases of the Italian lira and the British pound reveals that leaving/joining the EMS had a significant effect on the profitability of technical trading rules.

Contents

List of Tables

List of Figures

Acknowledgements

There are many people who have helped me in various ways to write this dissertation. In the first place, thanks are due to the professors who supervised my research: To Uwe Westphal, who unfortunately I only had the pleasure to work with for less than a year before his untimely death in December 1996; to Michael Funke, who was kind enough to take over the responsibility for my supervision and who greatly supported my dissertation project not only through discussions, but also by encouraging my research visits to the European University Institute; and finally to Mike Artis for his advice and support during my time at the EUI in Florence.

I am also grateful to a number of other researchers who have helped improve my thesis through comments and discussions. In particular I would like to thank Bernd Lucke, Chris Neely, Andrew Rose, Spyros Skouras, Rainer Kiefer, Klaus Adam and Bernd Wilfling. Special thanks are due to Nikolaus Siegfried, whose constant readiness to answer my questions about econometrics helped me to avoid some of the more serious blunders in empirical work. I also benefitted from the comments of seminar participants at the universities of Hamburg and Freiburg, as well at the 1998 meeting of the European Economic Association in Berlin.

On a more personal level I owe a great debt of gratitude to my parents both for the educational opportunities they offered me, as well as for encouraging me to follow my own ideas (even though they were usually rather naive). I am also grateful to Nina Wessels, without whose hospitality and enterprising spirit neither my time in Florence nor in Hamburg would have been the same. Moreover, I would like to thank the Pekrun and Maffucci families for the countless occasions on which I was their guest in Hamburg and in Florence, respectively. Last not least I would like thank Julie Bon for her patience and understanding of the psychological and the behavioural consequences of spending too much time in front of the computer.

Finally, I would like to thank the Bundesbank, the Federal Reserve Board as well as the Bank for International Settlements for providing the data used in this thesis.

Chapter 1

Introduction

One of the most puzzling findings in international economics is the almost complete inability of empirical exchange rate models to predict or even to explain ex-post the short and medium term movements of exchange rates. This first became apparent when Meese and Rogoff (1983) compared the out-of-sample forecast precision of a variety of different exchange rate models, including both structural and time series models. They found that in terms of the root-mean-square prediction error a simple random walk model outperformed all models for all exchange rates and at all forecasting horizons with the single exception of the USD/DEM exchange rate at the one month horizon.[1] Notably, for the structural models this finding holds in spite of the fact that they had the benefit of ex-post information concerning future values of explanatory variables. Whilst a number of authors have tried to improve exchange rate models and to demonstrate the predictive superiority of their models compared to the random walk, their attempts have been limited in their success.[2] Only at very long horizons (12 to 16 quarters) is there some evidence that it is possible to outforecast the random walk model consistently using econometric models.[3] At shorter horizons, in contrast, the consensus view in the literature is still well summed up in Dornbusch and Frankel's (1987) remark: 'Econometrically, most of the action is in the error term'.

A related and equally puzzling stylised fact in international economics is that whilst both

[1] Meese and Rogoff (1983), Table 1.

[2] See e.g. Frankel and Rose (1995) for a survey of the literature.

[3] See Mark (1995). There is also some evidence for the explanatory power of macroeconomic exchange rate models during periods of hyperinflations; e.g. Frenkel (1976).

nominal and real exchange rate variability is much greater in flexible than in fixed exchange rate regimes, there does not exist such a regime-specific difference in the volatility of macroeconomic variables. Mussa (1986) was one of the first to establish that real exchange rate variability varies with exchange rate regime. He did so by analysing exchange rate variability for bilateral USD exchange rates during and after the Bretton Woods regime. This finding in itself need not be very puzzling as long as exchange rate variability is a reflection of underlying economic disturbances, since in this case fixing exchange rates might lead to a reduction of exchange rate variability at the cost of increased volatility in other parts of the economy, for instance in the money market.[4] If this explanation was correct, one would expect that at least some other macroeconomic variables display regime-specific volatility. However, Baxter and Stockman (1989) examine the cyclic behaviour of output, consumption, trade flows, government consumption spending and the real exchange rate under alternative exchange rate regimes for 49 countries and find little evidence of systematic regime-specific differences apart from greater real exchange rate variability under flexible exchange rate regimes. Moreover, Flood and Rose (1995) find that the volatility of variables like money, output and prices differs only marginally across regimes, whilst exchange rate volatility is between 3 and 9 times higher during floating than during fixed exchange rate regimes. Flood and Rose (1995) argue that since models based on macroeconomic determinants cannot explain why exchange rate volatility is regime-specific, they are unlikely to be successful at explaining the short term movements of exchange rates.

These results suggest that macroeconomic variables are not only of little use for forecasting exchange rates (except for long horizons), but that if one wants to explain the short and medium term movements of exchange rates they are simply the wrong thing to look at.[5] As a result, the focus of attention of international economists has recently shifted towards an analysis of the microstructure of the foreign exchange market. One of the central issues in this so-called 'market microstructure' approach is the formation of market participants' exchange rate expectations. Here, particular attention is paid to potential heterogeneity in processes of expectation formation and to the effects of such heterogeneity on the dynamics of exchange rates. In this respect,

[4] As Jeanne and Rose (1999), p. 1, put it: "The economy can be thought of as a balloon; squeezing volatility out of one part (e.g. the exchange market) merely transfers the volatility elsewhere." For a more extensive argument to this effect see Frenkel and Mussa (1980), p. 379.

[5] As Flood and Rose (1999), p. 9, state: "To put it baldly, macroeconomics appears to be irrelevant in explaining high- and medium-frequency exchange rate dynamics for low inflation countries."

two groups of market participants are frequently distinguished: 'Fundamentalists', who form their expectations on the basis of (macro)economic fundamentals and 'Chartists' or 'Technical Analysts', who base their investment decisions solely on the past behaviour of exchange rates. Interest in the latter type of market participants increased considerably in the mid 1980's, which saw the unprecedented rise of the value of the US dollar against the currencies of all major industrial countries. A survey conducted by the Group of Thirty (1985) finds that the overwhelming majority of respondents from banks and securities houses believed that technical analysis had a significant impact on the foreign exchange market. In another survey, which was commissioned for the London foreign exchange market by the Bank of England in 1988, Taylor and Allen (1992) find that at least 90% of respondents place some weight on technical analysis when forming exchange rate expectations at short horizons of intra-day to one week. More recently, these results have been confirmed for foreign exchange professionals in Germany (Menkhoff 1997, 1998) and for foreign exchange dealers in Hong Kong (Lui and Mole 1998).

This evidence that technical analysis is in widespread use and that market participants consider it to have an influence on the behaviour of exchange rates has prompted a considerable body of research trying to explain the short term behaviour of exchange rates as a result of the interaction of fundamentalists and non-fundamentalists. There are two main ways in which the effects of non-fundamentalist trading on the dynamics of exchange rates has been modelled. The first was proposed in a paper by Frankel and Froot (1986). In their model, a market maker sets the price of the foreign currency as a weighted average of the expectations of both fundamentalists (who expect the exchange rate to return to its fundamental value) and chartists (who, for simplicity, take a random walk view of exchange rate changes). The weight given to each of the two groups by the market maker depends on their recent forecasting record. Frankel and Froot carry out simulations with this model in order to explain the last 20% of the appreciation of the US-Dollar in the early 80's, which took place in spite of the fact that macroeconomic fundamentals like real interest differentials had already started moving in the opposite direction.

Frankel and Froot's (1986) model of chartists and fundamentalists has been extended in a number of ways. De Grauwe and Dewachter (1992) and de Grauwe et al. (1993) introduce a sticky price monetary model for fundamentals (Dornbusch 1976) and consider commonly used technical trading rules to model chartists' expectations. They show that these assumptions can lead to chaotic paths in exchange rates. Moreover, they find that news do not necessarily have perceptible effects on exchange rates and that turbulence in the exchange rate market may come about in the absence of news. In a similar vein, Lux (1998) sets up a model of chartists and fundamentalists, in which chartist behaviour is characterised by trend following as well as by mimetic contagion. Again, chaotic attractors are found for a wide range of parameters, but in addition the model can also account for leptokurtosis in the distribution of exchange rate changes.[6] Finally, Vigfusson (1997) models the time-varying influence of technical and fundamental analysis in the framework of a Markov-Switching model. During the fundamentalist regime, exchange rates are assumed to be determined either by inflation differentials or a terms of trade proxy, whereas in the chartist state they are assumed to follow a moving average trading rule. Estimating the models Vigfusson "finds favourable though inconclusive evidence".

The main criticism that can be levied against these chartist-and-fundamentalist models is that the behaviour of market participants is introduced in an ad-hoc fashion rather than being derived within an optimising framework. This criticism is met in the other main approach to analysing the effects of non-fundamentalist trading in the foreign exchange market. This is the so called noise trader approach to finance. As the name suggests, it does not originate in international economics, but rather in financial economics and is a reaction to mounting evidence that asset prices do not always reflect fundamental values.[7] In these models, a distinction is made between rational, risk averse traders and noise traders, which Black (1986) defines as people who "trade on noise as if it were information". Whilst this field is large, two seminal papers are worth mentioning in particular.

De Long et al. (1990) propose an overlapping generations model in which noise traders have erroneous stochastic beliefs. In this model, a rational trader selling short an asset which is overpriced due to the fact that noise traders are (irrationally) optimistic runs the risk that noise traders will become even more bullish in the next period, which would push the price even further above its fundamental value. If rational traders are risk averse and have a limited

[6] Further models analysing contagion and epidemics of opinion include Lux (1995) as well as the simple, elegant and entertaining paper by Kirman (1993).

[7] See Shleifer and Summers (1990) for a lucid discussion of the noise trader approach.

horizon, this so called noise trader risk will prevent them from driving the price down all the way to fundamentals. In this way, persistent deviations of asset prices from their fundamental values are possible despite the existence of rational arbitrageurs. Scharfstein and Stein (1990) present a model in which two managers sequentially have to decide whether or not to invest in a project. Both managers receive private information about whether the investment's payoff will be positive or negative. The crucial point in Scharfstein and Stein's model is that because the labour market uses the observed investment decisions to update its belief about the ability of the managers, the second manager will (under some conditions) mimic the decision of the first *regardless* of his own private information. This latter feature of their model has since become known as 'rational herding'. Although the model is framed in terms of corporate investment, Scharfstein and Stein argue forcibly why rational herding might also be possible in asset markets.[8]

Whereas these papers analyse what factors might be responsible for non-fundamental movements in asset prices in general, two recent papers apply similar methods to address problems specific to the foreign exchange market. Firstly, Mark and Wu (1998) revisit the forward discount puzzle (which is that forward rates are biased predictors of future spot exchange rates) and show that a noise trading model à la De Long et al. (1990) can better explain this puzzle than the intertemporal asset pricing model. Secondly, Jeanne and Rose (1999) also build on the model by De Long et al. (1990) to explain why fixing exchange rates may reduce exchange rate volatility without causing an increase in the volatility of other macroeconomic variables. The key result is that under a flexible exchange rate regime there exist multiple equilibria of exchange rate volatility for the same amount of fundamental volatility: One equilibrium with few noise traders and low volatility and one with many noise traders and high volatility. What a target zone does in Jeanne and Rose's model is to pin down the economy on the equilibrium with low volatility.

Whilst this theoretical research, which explicitly takes into account the existence of market participants who do not form their expectations on the basis of fundamental variables, has already yielded some new insights and promises to yield more, it relies on somewhat artificial

[8]Scharfstein and Stein's (1990) paper was followed by a great number of other papers examining rational herding. See in particular Avery and Zemsky (1998), as well as the references contained therein.

assumptions concerning how non-fundamentalist traders actually behave. Moreover, the models are in general intended to explain puzzling empirical phenomena ex-post and thus do not tend to provide empirically testable predictions. For these reasons, the aim of the present investigation is to complement the theoretical research on the nature and consequences of non-fundamentalist trading with a detailed empirical analysis of the one type of non-fundamentalist trading for which there exists abundant evidence of its widespread use in foreign exchange markets, namely technical analysis.

Even though a great number of articles has been published on technical analysis, research in this area of economics is still highly fragmented and there is a marked lack of agreement as regards even the most basic questions of, for instance, how to measure technical trading rule profitability and how to adjust trading profits for risk. For this reason, the first objective of this investigation is to carry out a systematic analysis of the evidence for the profitability of technical trading rules for flexible exchange rates. The central issues in this context are the statistical and the economic significance of technical trading rule returns. We find strong evidence for technical trading rule profitability and show that for several ways of adjusting trading rule returns for risk, technical trading rules tend to be less, rather than more risky than buying and holding currencies or stock market indices. Moreover, we show that there are no indications of a decrease in technical trading profits during our sample period. Finally, we present evidence that technical trading rules systematically exploit deviations from uncovered interest parity, and that there is a clear day-of-the-week pattern in the profitability of technical trading rules, with by far the greatest returns accruing over weekends.

Having provided evidence of technical trading rule profitability for flexible exchange rates, the question of who or what is responsible for it arises immediately. One candidate, which has often been suspected as a source of technical trading profits, is foreign exchange interventions by central banks. The second objective of this investigation is to elucidate the relationship between technical trading rule profitability and central bank interventions. The starting point of our analysis is a puzzling empirical regularity which is that technical trading rules are extremely profitable when trading against central banks and that at the same time central banks make profits with interventions. We will show that this seeming contradiction is due to the fact that exchange rates only move in accordance with the intentions of central banks in the long run. In the short run, in contrast, the opposite tends to be the case. Moreover, we will argue that the reason why technical trading rules are so profitable on days when central banks intervene is that since interventions are aimed at turning technically motivated trends (or, more metaphorically, are aimed at pricking speculative bubbles) they are a more reliable signal of the existence of a trend than a technical trading rule by itself. Nevertheless, we show that central bank interventions can be profitably exploited by following technical trading rules whenever it becomes known that interventions have occurred and we argue that this may be a reason why central banks sometimes intervene secretly.

The third and final objective of this investigation is to examine the influence of the exchange rate regime on technical trading rule profitability. For this purpose, we examine the profitability of applying technical trading rules to several exchange rates which belong(ed) to the Exchange Rate Mechanism (ERM) of the European Monetary System (EMS). We find that for these exchange rates technical trading rules are unprofitable and that this was the case even during the early 1980s, when the exchange rate parities in the ERM were non-credible and subject to frequent realignments. Moreover, examining the cases of the Italian lira and the British pound reveals that leaving/joining the ERM had a significant effect on the profitability of technical trading rules, which shows that technical trading rule profitability is regime-specific. Assuming that technically motivated trading increases volatility, we argue that the above-mentioned regime-specific difference in the volatility of exchange rates may be due to the presence/absence of technical trading.

The structure of this investigation is as follows: Technical analysis is defined and explained in Chapter 2. In Chapter 3 we present the evidence of technical trading rule profitability for flexible exchange rates. Chapter 4 analyses the relationship between technical trading profits and central bank interventions. In Chapter 5 we turn to the question of the influence of participating in a fixed exchange rate regime on technical trading rule profitability. Chapter 6 concludes.

Chapter 2

Technical Analysis Defined

In one of the standard manuals of technical analysis, Pring (1991) defines technical analysis as follows:[1]

> The art of technical analysis - for it is an art - is to identify trend changes at an early stage and to maintain an investment posture until the weight of the evidence indicates that the trend has reversed.

In order to understand this definition it is necessary to clarify what a trend is, what is meant by 'weight of the evidence' and in what sense one can speak of technical analysis as an art.

2.1 Definition of a Trend

Murphy (1986) defines an uptrend as a sequence of successively higher peaks and troughs, where the exchange rate at (t) is described as a peak if it is higher than both the exchange rate at $(t-1)$ and $(t+1)$, and where a trough is defined analogously.[2] Let s_t be the natural logarithm of the exchange rate at date (t) and let $Peak_t$ ($Trough_t$) be the set of peaks (troughs) in the exchange rate time series, which occur before (t). This can be written as follows:

$$Peak_t \quad = \quad \{s_i : s_{i-1} < s_i, s_{i+1} < s_i, i \leq t-1\}$$

[1] Pring (1991), p. 16.
[2] Murphy (1986), p. 55.

$$= \{p_t^1, p_t^2 ... p_t^n\} \tag{2.1}$$

$$\begin{aligned} Trough_t &= \{s_i : s_{i-1} > s_i, s_{i+1} > s_i, i \leq t-1\} \\ &= \{tr_t^1, tr_t^2 ... tr_t^m\} \end{aligned} \tag{2.2}$$

Given these conventions, a sufficient condition for the existence of an uptrend at date (t) is:

$$(p_t^n > p_t^{n-1}) \text{ and } (tr_t^m > tr_t^{m-1}) \tag{2.3}$$

Analogously, a sufficient condition for the existence of a downtrend at date (t) is:

$$(p_t^n < p_t^{n-1}) \text{ and } (tr_t^m < tr_t^{m-1}) \tag{2.4}$$

In order to make a definition out of these conditions, it is necessary to assume in addition that once a trend has been identified, it is taken to be in effect until there are clear indications of a reversal.[3] For any date (t) these assumptions allow us to decide whether the exchange rate at this date is on an upward or downward trend.[4] Note that depending on the frequency of the data used, this definition applies to both short and long term trends.

Regarding the concept of the weight of the evidence, Pring (1991) states:[5]

> The weight of the evidence is the objective element in technical analysis. It consists of a series of indicators or techniques that work well most of the time in the trend-identification process.

The range of such indicators and techniques mentioned is extremely wide. It varies from the trivial to the exotic. In the following section we present a selection of indicators and trading rules which have been popular with technical analysts on the one hand and with researchers investigating the effectiveness of technical analysis on the other.

---

[3] Murphy (1986), p. 30.

[4] N.B. Strictly speaking this is only true for dates after a succession of higher or lower peaks and troughs has been observed for the first time.

[5] Pring (1991), p. 17.

2.2 Examples of Technical Trading Rules

Peak-and-Trough-Progression[6]

Probably the most basic trading rule that exists is known as peak-and-trough-progression. It is based on the definition of a trend given in (2.3) and (2.4) and states that one should go long in a currency during uptrends and go short during downtrends.[7] If ϕ_t denotes the signal of the trading rule with (+1) referring to a long and (−1) to a short position and if s_t denotes the natural logarithm of the exchange rate S_t, then this trading rule can be written as follows:[8]

$$\phi_t = \begin{cases} 1 \quad if \begin{cases} \phi_{t-1} = 1 \ and \ s_t \geq tr_t^m \\ \phi_{t-1} = -1 \ and \ s_t > p_t^n \end{cases} \\ -1 \quad if \begin{cases} \phi_{t-1} = 1 \ and \ s_t < tr_t^m \\ \phi_{t-1} = -1 \ and \ s_t \leq p_t^n \end{cases} \end{cases} \tag{2.5}$$

Momentum

Murphy (1986) defines the momentum as the difference between the current exchange rate and the exchange rate, say L_m days ago.[9] If the momentum crosses the zero line, a trading signal is given. Put in another way, the momentum rule states that one should go long in a currency while natural logarithm of the current exchange rate is greater than that L_m days ago. Formally:

$$\phi_t = \begin{cases} +1 & s_t > s_{t-L_m} \\ -1 & s_t < s_{t-L_m} \\ \phi_{t-1} & s_t = s_{t-L_m} \end{cases} \tag{2.6}$$

Some authors also use a so-called filter when applying this trading rule, which means changing from a short to a long position (a long to a short position) only if the current exchange rate

[6] Schulmeister (1987) refers to this trading rule as the point-and-figure rule. We follow instead the terminology used in Pring's (1991) manual of technical analysis.

[7] Pring (1991), p. 16f.

[8] N.B. Trading rules are defined in terms of the natural logarithm of the exchange rate, because this ensures that defining exchange rates in quantity rather than in price notation will simply lead to a change of sign of trading rule signals. We only deviate from this procedure in the case of the more exotic head-and-shoulders-rule.

[9] Murphy (1986), p. 277f. It needs to be pointed out that this definition is different from that given in Pring (1991), p. 127ff. We use Murphy's definition since this is the one used in a number studies of the effectiveness of technical analysis (e.g., Schulmeister 1987, Menkhoff and Schlumberger 1995).

has risen (fallen) by a certain percentage below (above) its lagged value. Filters are intended to avoid so-called whiplash signals, which are instances when the exchange rate crosses a critical value (like in this instance its lagged value) but then reverses strongly, thus causing heavy losses. Whilst introducing another dimension to the problem of deciding on a parametrisation of the trading rule, it is sometimes unclear whether filters have any real value.

Moving Average

Moving averages are used to smooth erratic movement in exchange rates in order to reveal the underlying trend. When a short moving average penetrates a longer moving average, a trend is considered to be initiated. The moving average trading rule states that one should go long in a currency as long as a shorter moving average of the natural logarithm of past exchange rates is greater than a longer moving average. Formally:

$$\phi_t = \begin{cases} +1 & \frac{1}{L_s}\sum_{i=0}^{L_s-1} s_{t-i} > \frac{1}{L_l}\sum_{i=0}^{L_l-1} s_{t-i} \\ -1 & \frac{1}{L_s}\sum_{i=0}^{L_s-1} s_{t-i} < \frac{1}{L_l}\sum_{i=0}^{L_l-1} s_{t-i} \\ \phi_{t-1} & \frac{1}{L_s}\sum_{i=0}^{L_s-1} s_{t-i} = \frac{1}{L_l}\sum_{i=0}^{L_l-1} s_{t-i} \end{cases} \qquad with \;\; L_s < L_l \tag{2.7}$$

There exist a number of more sophisticated versions of the moving average trading rule. These include using linearly weighted and exponentially smoothed moving averages, both of which attach greater value to more recent observations.[10] Moreover, like for the Momentum rule, filters are often used in an attempt to avoid whiplash signals, too.

Trading Range Break-Out[11]

As the name suggests, the trading range break-out rule signals a change in position whenever the exchange rate moves out of its previous trading range, which is defined by the local maxima and minima over a fixed horizon of past exchange rates. The rule states that one should go long if the natural logarithm of the exchange rate is greater than its previous local maximum and stay long until the exchange rate penetrates its local minimum. Formally:

[10] See, e.g., Pring (1991), p. 119ff.

[11] N.B. This trading rule is sometimes also referred to as the channel rule (e.g. Lukac et al. 1988, Taylor 1992). Our terminology follows Brock, Lakonishok and LeBaron (1992).

$$
\phi_t = \begin{cases} 1 \quad if \begin{cases} \phi_{t-1} = 1 \; and \; s_t \geq \underset{1 \leq i \leq K}{Min} \, (s_{t-i}) \\ \phi_{t-1} = -1 \; and \; s_t > \underset{1 \leq i \leq K}{Max} \, (s_{t-i}) \end{cases} \\ -1 \quad if \begin{cases} \phi_{t-1} = 1 \; and \; s_t < \underset{1 \leq i \leq K}{Min} \, (s_{t-i}) \\ \phi_{t-1} = -1 \; and \; s_t \leq \underset{1 \leq i \leq K}{Max} \, (s_{t-i}) \end{cases} \end{cases} \tag{2.8}
$$

Filter Rules

Filter rules were first used by Alexander (1961, 1964) in the context of testing the randomness of movements in stock prices. Interestingly, whilst being investigated extensively in the academic literature on technical trading rule profitability, filter rules appear not to be in wide use amongst practising technical analysts. Indeed, no mention of them is found in either of four technical analysis manuals consulted.[12] Filter rules state that one should go long in the foreign currency if the exchange rate rises by at least $x\%$ beyond its previous local minimum and keep that long position until the exchange rate falls by at least $y\%$ below its previous local maximum.[13] The range over which the local maxima (minima) are determined starts with the date of the most recent sell signal, $t_{\text{last sell}}$, (buy signal, $t_{\text{last buy}}$) and ends at the present date. The filter rule can be written formally as follows:

$$
\phi_t = \begin{cases} 1 \quad if \begin{cases} \phi_{t-1} = 1 \; and \; s_t \geq \underset{t_{\text{last buy}} \leq i \leq t}{Max} (s_i) - y \\ \phi_{t-1} = -1 \; and \; s_t > \underset{t_{\text{last sell}} \leq i \leq t}{Min} (s_i) + x \end{cases} \\ -1 \quad if \begin{cases} \phi_{t-1} = 1 \; and \; s_t < \underset{t_{\text{last buy}} \leq i \leq t}{Max} (s_i) - y \\ \phi_{t-1} = -1 \; and \; s_t \leq \underset{t_{\text{last sell}} \leq i \leq t}{Min} (s_i) + x \end{cases} \end{cases} \tag{2.9}
$$

Head-and-Shoulders

The last and already somewhat exotic indicator to be considered is the famous head-and-shoulders pattern. It is defined as a sequence of three peaks, of which the second one (the head) is higher than the other ones (the shoulders). The line through the local minima between head

[12] These are Pring (1991), Murphy (1986), Edwards and Magee (1997) and Luca (1997).

[13] Note the difference between the trading range break-out rule and the filter rule. The former gives a buy signal depending on how the current exchange rate compares to a local maximum, whilst the latter gives a signal depending on how it compares to a local minimum.

and left shoulder and between head and right shoulder is called neckline. Once the exchange rate crosses the neckline after forming the right shoulder, the pattern is considered complete and is interpreted as an indication of an imminent downtrend.[14]

Clearly, the formalisation of the head-and-shoulders trading rule requires a definition of peaks and troughs. Osler and Chang (1995) use the trading signals of a filter rule with parameters $x = y = \chi$ to generate a sequence of alternating peaks and troughs. Peaks are defined as the local maxima between the dates of buy signals of the filter rule and subsequent sell signals and troughs are defined as the local minima between the dates of sell signals and subsequent buy signals.

Let $Peak(0)$, $Peak(1)$, $Peak(2)$ and $Peak(3)$ denote four successive peaks and $Trough(0)$, $Trough(1)$ and $Trough(2)$ the troughs between the peaks. According to Osler and Chang, the head-and-shoulders pattern is then characterised by the following 5 conditions:[15]

1. The head must be higher than the shoulders:

$$Peak(2) > Peak(1) \text{ and } Peak(2) > Peak(3)$$

2. The pattern must be preceded by an uptrend, i.e. by rising peaks and troughs:

$$Peak(1) > Peak(0) \text{ and } Trough(1) > Trough(0)$$

3. Vertical symmetry - the left shoulder must be at least as high as the midpoint between right shoulder and its preceding trough and vice versa:

$$\begin{aligned} Peak(1) &\geq \frac{1}{2}(Peak(3) - Trough(2)) \\ & \quad and \\ Peak(3) &\geq \frac{1}{2}(Peak(1) - Trough(1)) \end{aligned}$$

4. Horizontal symmetry - the time between left shoulder and head must not be more than

[14] N.B. There also exists an analogous trading rule (the inverse head-and-shoulders pattern) signalling imminent upturns. See Pring (1991), p. 66f.

[15] Osler and Chang (1995), p. 11f.

2.5 times the time between the head and the right shoulder (and vice versa):

$$\begin{aligned}(t_{Peak(2)} - t_{(Peak(1))}) &< 2.5(t_{Peak(3)} - t_{(Peak(2))}) \\ & \quad and \\ (t_{Peak(3)} - t_{(Peak(2))}) &< 2.5(t_{Peak(2)} - t_{(Peak(1))})\end{aligned}$$

5. Since the head-and-shoulders pattern is supposed to signal an imminent trend reversal, a time limit is imposed:

$$t < t_{Peak(3)} + (t_{Peak(3)} - t_{Peak(1)})$$

If these conditions are fulfilled, the head-and-shoulders patterns gives a sell signal when the exchange rate, S_t, crosses the neckline, i.e. if:

$$S_t < Trough(1) + \frac{(t - t_{Trough(1)})}{(t_{Trough(2)} - t_{Trough(1)})}(Trough(2) - Trough(1))$$

2.3 Technical Analysis as an Art

Given the above patterns and trading rules, Pring (1991) explains what he means by calling technical analysis an art as follows:[16]

> The art consists of combining these indicators into an overall picture and recognizing when that picture resembles a market peak or trough.

As Pring admits, this is the subjective element of technical analysis, which clearly does not lend itself to rigorous empirical analysis. Other authors are even more contentious concerning the prerequisites of carrying out technical analysis. For instance, Luca (1997) writes:[17]

> [Chart interpretation] requires an individual, innate talent mixed with a personally acquired and refined research and understanding.

Given such remarks and the open admission of subjectivity it may not be altogether surprising that technical analysis has been treated in the best case with ridicule and in the worst

[16] Pring (1991), p. 17.
[17] Luca (1997), p. 9.

with outright contempt by the academic community. The following remark by Malkiel (1990) sums up well the consensus concerning the usefulness of technical analysis in the early 80's:[18]

> Technical analysis is anathema to the academic world. We love to pick on it. Our bullying tactics are prompted by two considerations:
>
> (1) The method is patently false and
>
> (2) its easy to pick on.

Recently, however, evidence supporting the profitability of technical trading rules has been mounting.[19] The next chapter adds to this body of evidence by showing that even without 'artistic' predisposition or 'innate talent' technical trading rules are useful at predicting future movements of floating exchange rates.

[18] Malkiel (1990), p. 132.

[19] See Menkhoff and Schlumberger (1995) for a survey of the literature.

Chapter 3

Technical Trading Rule Profitability for Floating Exchange Rates

3.1 Methodology

When addressing the issue of the profitability of technical trading rules the first question that needs to be answered is which trading rules in particular should be looked at. Most papers investigating the profitability of technical analysis evaluate the returns from following trading rules that are either said to be popular in practice or assumed to be representative in some other sense. For instance, Brock, Lakonishok and LeBaron (1992) note:[1]

> The most popular moving average rule is 1-200, where the short period is one day and the long period is 200 days. While numerous variations of this rule are used in practice, we attempted to select several of the most popular ones: ...

As Neely et al. (1996) point out, the problem with this approach is that "there remains some doubt as to whether the reported excess returns could have been earned by a trader who had to make a choice about what rule or combination of rules to use at the beginning of the sample period."[2] Another problem is that the popularity of a trading rule today may not be independent of its performance during the sample period that is being investigated.

[1] Brock, Lakonishok and LeBaron (1992), p. 1735. See also Menkhoff and Schlumberger (1995), who look at 33 Moving Average trading rules "with round numbers preferred" (p. 196).

[2] Neely et al. (1996), p. 2.

There are two studies which avoid (at least to some extent) this so-called 'ex-post selection bias'. The first is Taylor (1992) which examines filter, trading range break-out and moving average trading rules as well as a more sophisticated ARIMA-based decision rule using daily currency futures prices from 1977 to 1987. Starting from June 1982, every half year the parameters that would have maximised returns over the entire space of past prices are determined and are then used for the next six months to yield ex-ante trading rule returns. Taylor finds that the trading rules produced significant excess returns and that "[m]uch to the author's surprise the sophisticated rules would have been outperformed by the simple filter and technical rules."[3] Taylor's procedure can be criticised on the grounds that whilst the choice of individual trading rules has been endogenised, there still exists a potential bias arising from the choice of trading rule classes considered.

The paper by Neely et al. (1996) tries to avoid this problem by using genetic programming to endogenise the choice of trading rule to be examined. Their approach involves splitting up the sample into three sub-periods: A training, a selection and a validation period. During the training period, trading rules are randomly generated from a very wide range of trading rule classes and recombined depending on how profitable they were in this period. After each step of recombination, the most profitable rule is examined in the selection period. After 50 recombinations the most profitable rule during the selection period is determined and then analysed in the validation period (unless it produced a negative average return, in which case it is discarded).[4] Neely et al. generated and examined 100 rules in this manner for six daily exchange rate time series from 1975 to 1994.

The average annual excess return across the 100 rules is 6.05% for the DEM/USD exchange rate, whilst the average excess return across rules and exchange rates is 2.87%. Moreover, almost all of the rules produced positive excess returns. The authors argue that this performance is good given that the validation period was very long, lasting from 1981 to 1995. Whilst this may be true, it would nevertheless have been interesting to see if the results would have improved if the training and selection periods had been lengthened. It is also worth noting that the returns found by Neely et al. (1996) are considerably smaller than those found by almost all authors

[3] Taylor (1992), p. 108.

[4] N.B. To be precise, a rule can also be chosen after less than 50 recombinations if the last 25 recombinations did not produce a rule with as high a return during the selection period. Cf. Neely et al. (1996), p. 11.

analysing simple technical trading rules and that this is despite the fact that "it takes 81 days of computing time to replicate the results of [their] paper on a 120 MHz Pentium".[5] This confirms Taylor's (1992) finding that the use of sophisticated econometric and computational techniques does not necessarily lead to better results than following simple technical trading rules.[6]

The approach adopted in this study is to examine the profitability of simple technical trading rules exhaustively, which is to say, for *all reasonable* parameter values. In order to avoid a potential ex-post selection bias concerning the trading rule class to be investigated, we only consider classes which are known to have been in wide use at the time our sample starts. Since our datasets start in 1979, it seems unproblematic to examine the moving average rule, which Cornell and Dietrich (1978) claim to be in wide use on foreign exchange markets.[7] Because of its fundamental role in technical analysis, we also look at the peak-and-trough-progression rule, which is said to date back to the writings of Charles Dow in the beginning of this century and which is also discussed in Kaufman (1978). The latter rule makes the job of carrying out an exhaustive analysis of its profitability easy, since it does not have any parameters. From the class of moving average trading rules we consider its simplest member, which is the moving average rule with the length of the shorter moving average fixed to unity.[8] The highest value for the length of the longer moving average discussed in either the academic literature or technical analysis manuals is 250. In order to be sure that really all reasonable parameter values are considered, we examine the profitability for all lengths of the longer moving average between 2 and 500 days.

The justification for using this extensive rather than Taylor's (1992) intensive approach is that trading rule returns evaluated according to Taylor's approach are extremely sensitive to changes in the length of periods over which the return-maximising trading rule parameters are chosen (this will be shown in Section 3.5.2).[9] Simply reporting trading rule returns for all parameter values is therefore likely to give a more reliable reflection of technical trading rule

[5] Neely et al. (1996), fn 6.

[6] Note also in this context the recent study by Gençay (1999), which finds that the use of moving average trading rule signals in nearest-neighbour and feedforward network regressions significantly improves their forecasting accuracy.

[7] Cornell and Dietrich (1978), p. 115.

[8] N.B. These rules are also known as single moving average rules.

[9] We suspect, but did not ascertain, that a similar sensitivity problem also affects the use of genetic programming to examine technical trading rule profitability.

profitability.

This chapter is structured as follows: After discussing the data used in our calculations in Section 3.2, we use conventional t-tests as well as bootstrapping simulations to show that both the peak-and-trough-progression and the moving average rule produce returns that are significantly greater than zero in Section 3.3. In Section 3.4 the question whether trading rule returns can be interpreted as a compensation for risk is examined considering a whole range of methods for adjusting returns for risk. The issue of the stability of technical trading rule returns is addressed in Section 3.5. In Section 3.6 we analyse the relationship between trading rule profits and other empirical anomalies concerning flexible exchange rates. Finally, Section 3.7 contains a discussion of our results.

3.2 Data Summary

For the calculations that follow we use daily exchange rate and one-week eurorate data for the US dollar (USD), the German mark (DEM) and the Japanese yen (JPY). The sample runs from January 2, 1979 to December 31, 1992.[10] Figures 3-1 and 3-2 show the frequency distributions and some summary statistics of log first differences of daily DEM/USD exchange rates and of the corresponding daily interest differentials.[11] [12] Exchange rate changes appear to have little drift, there is little evidence of skewness but strong evidence of excess kurtosis (fat tails). Compared to interest differentials, exchange rate changes have a substantially higher variance. The interest differential between the US and Germany was positive on average, yet the negative tail of the distribution is relatively thick, as reflected in negative skewness. Moreover, there is evidence of fat tails for interest differentials, too. Using the Jarque-Bera test we can reject the Null-hypothesis of normality of exchange rate changes and of interest differentials at all conventional significance levels.

All three exchange rate time series were also checked for non-stationarity. Table 3.1 contains

[10] Exchange rates are the London close from Nat West Bank, interest rates are the London close from the Financial Times and Nat West Bank.

[11] Daily interest differentials were calculated by dividing annual rates of interest by 260. Dividing the interest differential by 360 would lead to an understatement of the influence of interest differentials because of weekends. We divide by 260 (=52*5 working days per year) so that interest differentials are correctly accounted for *on average*.

[12] The corresponding Figures for the JPY/USD and JPY/DEM exchange rates can be found in Appendix A.

Figure 3-1: Frequency Distribution and Summary Statistics of Daily DEM/USD Exchange Rate Changes

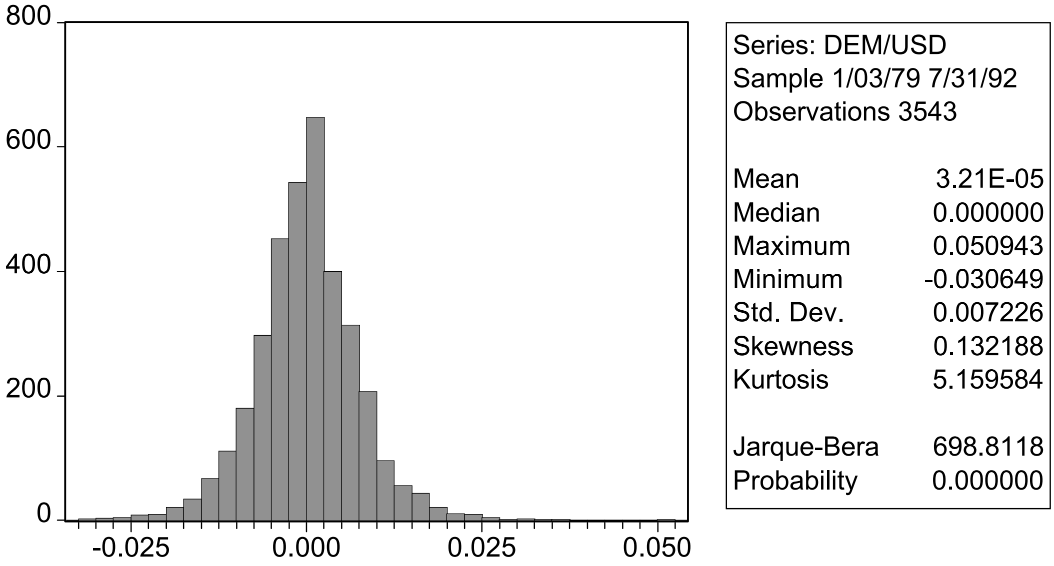

Figure 3-2: Frequency Distribution and Summary Statistics of Daily USD-DEM Interest Differentials

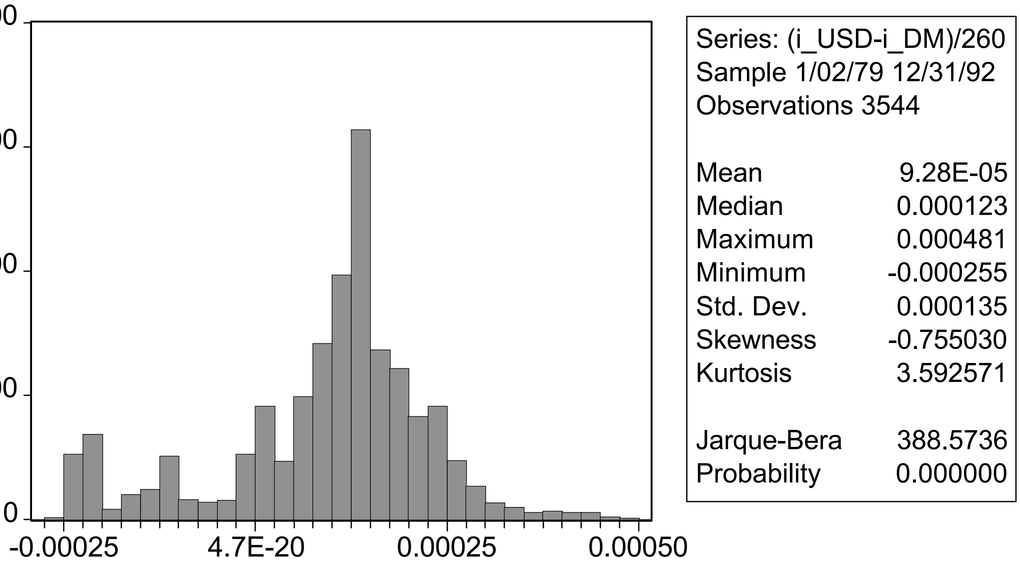

Table 3.1: Phillips-Perron Unit Root Test for Daily Changes of Flexible Exchange Rates

Equation	$\Delta s_t - \Delta s_{t-1} = \alpha + \beta \Delta s_{t-1} + u_t$		
	DEM/USD	JPY/USD	JPY/DEM
$\widehat{\alpha}$	3.72E-05	1.30E-04	8.41E-05
S. E.	1.21E-04	1.10E-04	9.54E-05
t-value	0.3070	1.158	0.8814
$\widehat{\beta}$	-0.9884	-0.9849	-0.9199
S. E.	0.0168	0.0168	0.0168
t-value	-58.87	-58.67	-54.91
PP Test Statistic	-58.91**	-58.79**	-54.96**
Durbin-Watson	2.001	2.000	2.000

** indicates significance at the 1% level. Truncation lag = 8.

the results of a Phillips-Perron test. The Null-hypothesis of a unit root is strongly rejected for the first differences of all exchange rates. It is well known that high frequency exchange rate time series do not have a constant variance. Table 3.2 contains the results of Engle's (1982) ARCH-test of conditional heteroskedasticity. The Null-hypothesis of homoskedasticity is also strongly rejected for all exchange rates. We also examined the pattern of autocorrelation of exchange rate changes. Table 3.3 shows the autocorrelation coefficients (AC), the partial autocorrelation coefficients (PAC), as well as Ljung Box Q-statistics (Q-stat) with corresponding p-values for lags from 1 to 20 for the DEM/USD exchange rate.[13] Clearly, there is no evidence of autocorrelation in daily exchange rate changes.[14]

In particular in the context of analysing the usefulness of extrapolative technical trading rule it suggests itself to analyse time dependencies in exchange rate changes not only by looking at their autocorrelation structure but also by examining their spectrum. This is because the spectrum reveals the importance of periodic cycles at different frequencies, which, one might naively think, can be exploited using extrapolative technical trading rules. The population spectrum of a covariance-stationary process, whose autocovariances, γ_i, are absolutely summable, is defined as follows:[15]

[13] The corresponding tables for the JPY/USD and JPY/DEM exchange rates are in Appendix B. See Tables B.1 and B.2.

[14] N.B. For the JPY/DEM exchange rate there is some evidence of positive first order autocorrelation. This might incline one to expect higher technical trading profitability for this exchange rate. As it turns out, however, the trading rules tend to be less rather than more profitable. This suggests that the connection between autocorrelation and trading profits is less close as one might think.

[15] Hamilton (1994), p. 152f.

Table 3.2: ARCH-Test for Daily Changes of Flexible Exchange Rates

Equation	$(s_t - \overline{s})^2 = \alpha + \beta(s_{t-1} - \overline{s})^2 + u_t$		
	DEM/USD	JPY/USD	JPY/DEM
$\widehat{\alpha}$	4.58E-05	4.00E-05	2.71E-05
S. E.	1.98E-06	1.71E-06	1.19E-06
t-value	23.16	23.37	22.73
$\widehat{\beta}$	0.1214	0.0621	0.1653
S. E.	0.0167	0.0168	0.0166
t-value	7.2833	3.7070	9.97
F-statistic	53.05	13.74	99.46
p-value	0	0.0002	0
$Obs * R^2$	52.29	13.69	96.79
p-value	0	0.0002	0
Durbin-Watson	2.0211	2.007	2.048

Table 3.3: Autocorrelation Function of Daily DEM/USD Exchange Rate Changes

Laglength	AC	PAC	Q-Stat	p-value
1	0.012	0.012	0.4782	0.489
2	0.001	0.001	0.4815	0.786
3	0.03	0.03	3.5792	0.311
4	-0.009	-0.01	3.8945	0.42
5	0.029	0.029	6.8872	0.229
6	-0.009	-0.01	7.1605	0.306
7	0.013	0.014	7.7323	0.357
8	0.031	0.028	11.057	0.198
9	0.025	0.026	13.278	0.15
10	-0.017	-0.019	14.284	0.16
11	0	-0.001	14.284	0.218
12	-0.002	-0.004	14.304	0.282
13	-0.001	-0.001	14.311	0.352
14	0.019	0.018	15.596	0.339
15	0.009	0.01	15.899	0.389
16	-0.001	-0.003	15.905	0.46
17	-0.023	-0.025	17.855	0.398
18	-0.012	-0.011	18.357	0.432
19	0.031	0.031	21.695	0.3
20	0.007	0.007	21.846	0.349

N.B. The asympt. S.E. of the corr. coeffs. is 0.0168.

$$s_Y(\omega) = \frac{1}{2\pi}\left(\gamma_0 + 2\Sigma_{j=1}^{\infty}\gamma_j \cos(\omega j)\right)$$

Rather than using the sample analogue of the population autocovariances, we use a modified Bartlett kernel to estimate the spectral density function.[16] The estimate of the spectrum is thus:[17]

$$\widehat{s}_Y(\omega) = \frac{1}{2\pi}\left(\widehat{\gamma}_0 + 2\Sigma_{j=1}^{q}(1 - j/(q-1))\widehat{\gamma}_j \cos(\omega j)\right) \tag{3.1}$$

Evaluating (3.1) requires that one decides on a bandwidth q. We let $q = 4(\frac{T}{100})^{\frac{2}{9}}$ as in Newey and West (1987), which yields $q = 9$. Figure 3-3 shows the spectral density of DEM/USD exchange rate changes as a function of j, where $\omega_j = \frac{2\pi j}{T}$.[18 19]

Whilst it is possible to make out three peaks, these are rather small. In order to check whether they are likely to have come about by chance, we repeatedly sampled with replacement from the empirical distribution of exchange rate changes and estimated the spectral densities of these i.i.d processes. We found that peaks as large as the ones for the original series were the norm rather than the exception. This leads us to conclude that there is no evidence of periodic cycles and that thus neither in the time, nor in the frequency domain are there any indications of exploitable time dependencies in exchange rate changes. Note, however, that this does not rule out a potential use for technical trading rules because trading rules like those reviewed in the previous section are non-linear functions of past observations. For this reason, a separate analysis of the statistical significance of technical trading rule returns is necessary.

[16] The problems of using the sample analogue of the population autocovariances are discussed in Hamilton (1994), p. 163f.

[17] Hamilton (1994), p. 167.

[18] Note that this implies that if we find a peak at $j = j_0$ then this corresponds to a cycle with a period of $\frac{T}{j_0}$ days.

[19] The corresponding figures for the other two exchange rates are, again, in Appendix A.

Figure 3-3: Spectrum of Daily DEM/USD Exchange Rate Changes

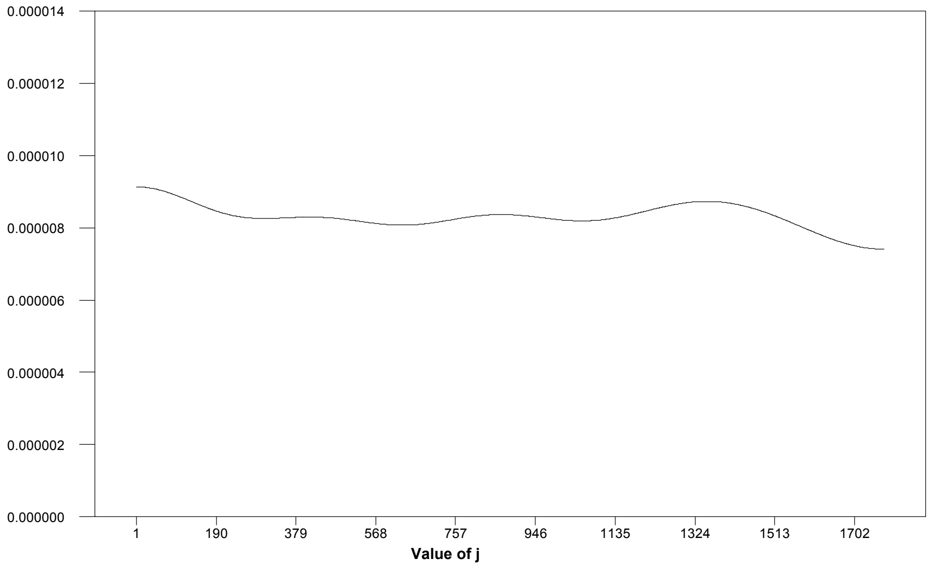

3.3 Statistical Significance of Technical Trading Rule Returns

3.3.1 Significance Assuming that Trading Rule Returns are Independently and Identically Distributed

Using definitions of trading rule signals, ϕ_t, from equation (2.5) for the peak-and-trough-progression rule and from equation (2.7) for the moving average trading rule, daily trading rule returns, r_t, are evaluated as follows:

$$r_t = \phi_{t-1}(s_t - s_{t-1} - \frac{1}{260}(i_{t-1} - i^*_{t-1})) - c * \frac{n}{T} \tag{3.2}$$

i_t (i^*_t) are the domestic (foreign) 1-week eurorates;[20] n is the number of times the trading rule indicates a change in position over the entire sample, T is the number of days in the sample and c is the cost of a round-trip on the markets for major currencies. We let c equal 0.05%,

[20] Ideally, one would want to use overnight eurorates. The differences are likely to be small, however. Note also that (3.2) assumes that exchange rates are written in price notation, i.e. s_t is the price of foreign currency in terms of domestic currency.

Table 3.4: Profitability of the Peak-and-Trough-Progression Rule

	DEM/USD	JPY/USD	JPY/DEM
Mean	0.1188	0.0846	0.0789
Std. Dev.	0.1173	0.1064	0.0929
Std. Err.	0.0318	0.0288	0.0252
t-value	3.7379	2.9354	3.1328

which is approximately the transaction cost faced by a large institutional investor.[21] Annual returns are derived from daily returns by multiplying by 260.

Table 3.4 contains evidence of the profitability of the peak-and-trough-progression rule for the bilateral exchange rates between USD, DEM and JPY. Average annual returns lie between 7 and 12%. Moreover, higher mean returns seem to go along with higher standard deviations of returns. If one assumes that trading rule returns come from a population with constant mean and finite variance, one can use the central limit theorem to test the hypothesis of zero population mean against the alternative hypothesis of returns greater than zero. The last column in Table 3-4 gives the values of the test statistic which is approximately standard normally distributed under the assumptions made. For all exchange rates the Null-hypothesis can be rejected at conventional significance levels.

Figures 3-4 to 3-6 display the returns from following moving average trading rules for the three exchange rates. Average annual returns are drawn as a function of the parameter chosen for the length of the moving average. The figures also show the lower bounds of one-sided 5% and 1% confidence intervals, which are derived making the same assumptions as before. If a lower bound is greater than zero, this means that the Null-hypothesis of zero returns can be rejected at the 5%/1% level. We find that for very short lengths of the moving average trading rule returns are dominated by transaction costs for all three exchange rates. Then, for lengths up to around 200, returns are significant. For longer lengths returns tend to fall and become insignificant for very large values of the moving average.[22] Thus, for a very wide

[21] Neely et al. (1996), p. 12. This value for the size of transaction costs is the median of the values considered in the literature.

[22] This pattern of movement is less clear for the DEM/USD. In contrast, for some range of parameter values returns even appear to increase. This result should be interpreted bearing in mind that moving average rules with very large parameters are useful only at picking up very large swings in the exchange rate and that the USD/DEM exchange rate was characterised by some unusually large swings during the sample period (i.e. especially in the early 1980's). For this reason it seems reasonable to interpret the decrease of trading rule returns for very large

Figure 3-4: Annual Returns of Moving Average Trading Rules for the DEM/USD Exchange Rate

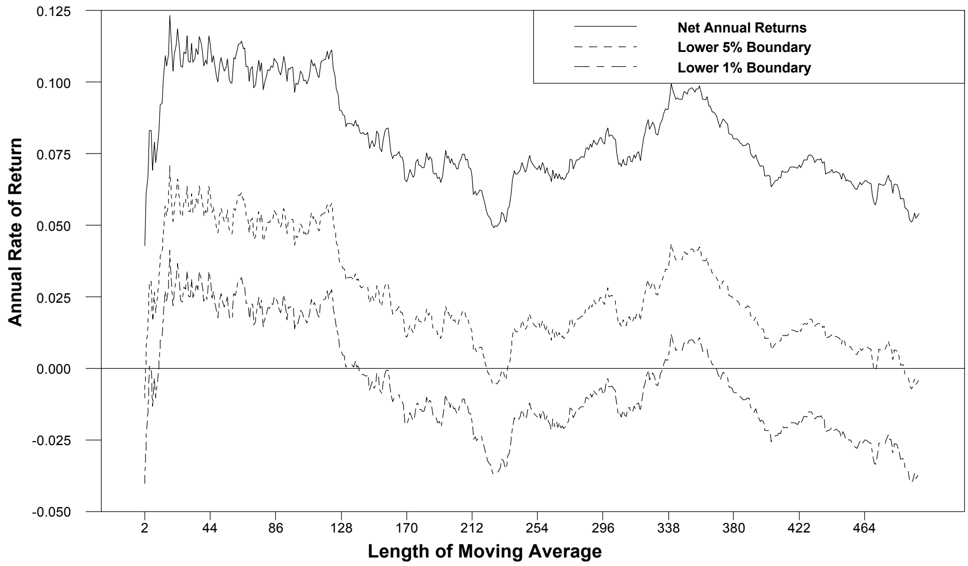

range of parameter values (between about 5 to at least 200), moving average trading rules yield significant returns.

The problem with the above inference concerning the significance of trading rule returns is that significant trading rule returns indicate that exchange rate changes (net of interest differentials) are not i.i.d. and that thus trading rule returns are not i.i.d. either.[23] It follows that one of the assumptions used to derive the test-statistics and the confidence intervals is invalid, which makes them, strictly speaking, meaningless. One way to deal with this problem is to make explicit assumptions concerning the process generating exchange rate changes and use bootstrapping simulations to assess how likely it is to find as large or larger trading rule returns for the assumed data generating processes.

parameter values as the 'normal' behaviour, and to consider the results for the USD/DEM exchange rates as unusual.

[23]For discussion of this issue see Levich and Thomas (1993), p. 452, and Kho (1996), p. 260f.

Figure 3-5: Annual Returns of Moving Average Trading Rules for the JPY/USD Exchange Rate

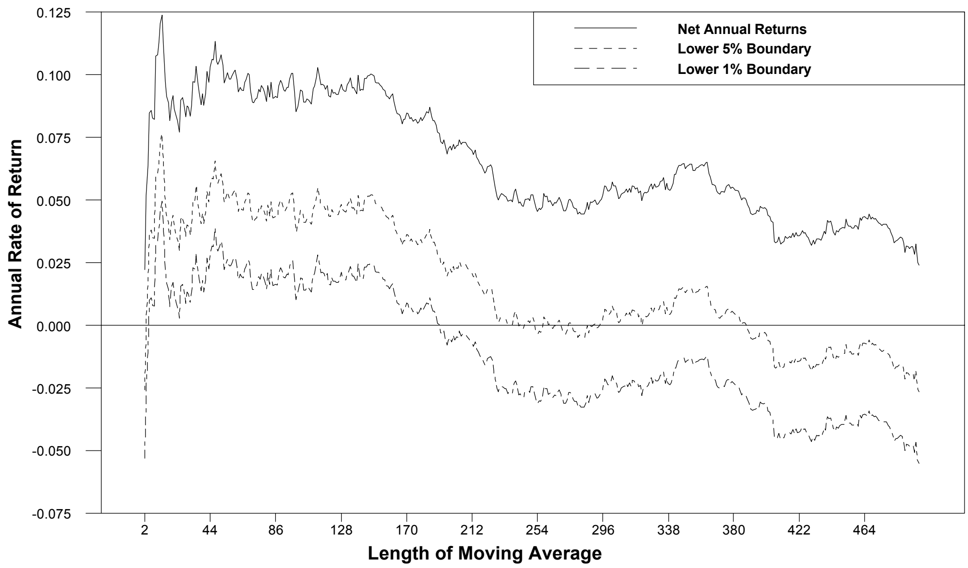

Figure 3-6: Annual Returns of Moving Average Trading Rules for the JPY/DEM Exchange Rate

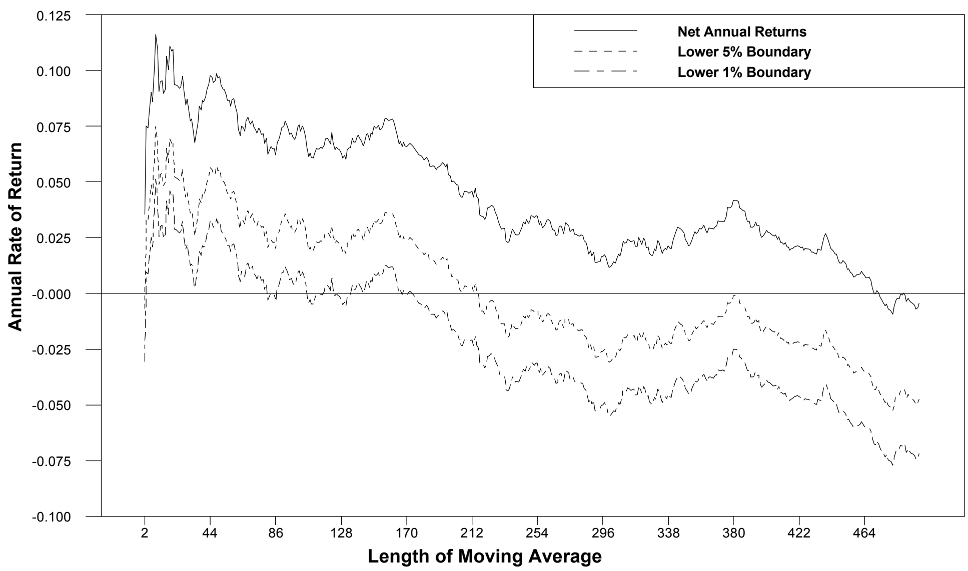

3.3.2 Significance Tests Using GARCH-Bootstraps

The use of bootstrapping methodology to test the significance of technical trading rule returns was pioneered by Brock, Lakonishok and LeBaron (1992), who used daily DJIA data from 1897-1986 to examine whether returns from following moving average and trading range break-out rules are compatible with a number of popular time series models of index returns, including the random walk, a stationary AR(1) and two GARCH models. They find that none of the models can explain the profitability of the technical trading rules and suggest that technical trading rules can be seen as a means to test the specification of time series models of exchange rates (and more generally of any asset pricing model).

LeBaron (1991) carries out an analogous investigation for weekly USD/DEM, USD/GBP and USD/JPY exchange rates between 1974 and 1991. He examines the profitability of moving average trading rules and considers a Random Walk, a GARCH(1,1) and a Markov regime switching model as null models. Like Brock, Lakonishok and LeBaron (1992) he finds that none of the models considered can account for the size of the trading rule returns found for the original series. Levich and Thomas (1993) replicate some of LeBaron's (1991) results using daily currency futures prices from 1976 to 1990. They look at moving average and filter rules and carry out Random Walk bootstrapping simulations. They find strong evidence against the Random Walk model. Similar results have been found by Osler and Chang (1995) for the head-and-shoulders trading rule and by Neely et al. (1996) for returns from using genetic programming.

The intuition behind using bootstrapping simulations to examine the significance of trading rule returns is extremely simple.[24] To start off with, one needs to estimate a time series model for the exchange rate under consideration. The estimation yields parameter estimates as well as a set of residuals. With the estimated time series model it is possible to generate 'pseudo exchange rate time series', by means of a computer random number generator. There are two ways one can do this. On the one hand, one can make an assumption about the distribution of the residuals and use the estimated model parameters and the generated residuals to generate a new 'time series'. On the other hand, one can let the computer draw with replacement a sample

[24] For an excellent introduction to bootstrapping see Efron and Tibshirani (1993). See Brock et al. (1992) for a good exposition of bootstrapping methodology as applied to trading rule profitability.

from the set of residuals of the original estimation. Sampling with replacement from the set of residuals is what is meant by bootstrapping. Compared to the aforementioned approach, bootstrapping has the advantage that no assumptions need to be made about the distribution of the residuals.[25] Once such a 'pseudo time series' has been generated, one can examine how profitable a certain technical trading rule would be if the 'pseudo time series' were to realise. The significance of technical trading rule returns can be assessed by generating a large number of 'pseudo time series', applying the trading rule examined to these series and determining in which proportion of simulations returns as high or higher than those for the original time series are found. If this proportion is small, observed returns are evidence against the data generating process assumed for exchange rate changes.

The first step to testing the significance of returns from peak-and-trough-progression and moving average trading rules is thus to develop a good model for exchange rate changes. One very popular class of times series models for exchange rates is the class of autoregressive conditionally heteroskedastic models. From this class we chose the GARCH model as proposed by Bollerslev (1986).[26] The following GARCH(1,1) model was estimated for each exchange rate.[27]

$$\begin{aligned} \Delta s_t &= \alpha_0 + \Sigma_{i=1}^{n} \alpha_i \Delta s_{t-i} + u_t \\ h_t &= \beta_0 + \beta_1 h_{t-1} + \beta_2 u_{t-1}^2 \\ u_t &= \sqrt{h_t} v_t \\ v_t &\sim N(0,1) \end{aligned} \tag{3.3}$$

[25] It might be argued that since one needs to make a distributional assumption for the purpose of estimation, one might as well use the same assumption when doing the simulations. This argument overlooks, however, that the parameter estimates from the estimation of the time series model may still be consistent even if the assumptions concerning the residuals are violated. Take for instance the (realistic) case where the residuals are not normally distributed, but rather have fat tails. Whilst this deviation from the distributional assumption may not harm the consistency of the parameter estimates, it should be taken into account when 'pseudo time series' are generated.

[26] As an alternative we also tried the Exponential GARCH model, which allows for asymmetry between positive and negative innovations. We found evidence of asymmetric effects for only two of the three exchange rates, and even then it was only marginal. For this reason we decided to stick to the basic GARCH model.

[27] The number of ARCH and GARCH terms was determined using the Schwarz (1978) criterion. The lagged endogenous variables to be considered in the regression equation were determined as follows: For each exchange rate the model was first estimated with only a constant term. Using the Ljung-Box Q-statistic, the standardised residuals (v_t) were then checked for autocorrelation. When evidence of autocorrelation was found, more lagged endogenous variables were added successively until signs of autocorrelation disappeared. After this, lagged endogenous variables whose parameter estimates were not even marginally significant were removed, at each step checking again for autocorrelation in the residuals.

Table 3.5: Parameter Estimates for GARCH(1,1) Model of Daily DEM/USD Exchange Rate Changes

	Coeff.	Std. Error	t-value	p-value
α_1	0.0110	0.0180	0.6086	0.5428
α_3	0.0259	0.0170	1.5276	0.1266
α_5	0.0430	0.01439	2.9839	0.0028
β_0	2.05E-06	2.38E-07	8.611	0.0000
β_1	0.8524	0.0114	74.57	0.0000
β_2	0.1119	0.0094	11.91	0.0000

Table 3.5 contains the results of the maximum likelihood estimation of (3.3) for the DEM/USD exchange rate.[28] It might seem peculiar that a_1 was included although its t-statistic is clearly insignificant. The reason for this is that if it is excluded, there is evidence of autocorrelation in the standardised residuals. Moreover, it might be objected that the assumption of normality of the standardised residuals is unlikely to hold in practice. Whilst this is true, Bollerslev (1986) showed that as long as the expected value of the standardised residuals is zero, their variance is one and their fourth moment is finite, then the parameter estimates are strongly consistent even if the normality assumption is violated.[29] Moreover, Monte Carlo studies by Bollerslev and Wooldridge (1992) and Lumsdaine (1995) indicate that already for samples as 'small' as 500 observations the biases in the quasi maximum likelihood estimators are small.[30] Given that our sample contains more than 3500 observations, it seems fair to assume that the non-normality of the standardised residuals will not substantially bias our results.

Using the estimated parameters along with the residuals, 2000 'pseudo time series' were generated for each exchange rate.[31] The peak-and-trough-progression and moving average trading rules with parameters 50, 100, 150, 200 and 250 were then applied to these series to derive an empirical distribution of returns for each trading rule. Table 3.6 contains the annualised mean and standard deviation of trading rule returns as well as the proportion of simulated time series for which returns were higher than for the original time series. We find

[28] Estimation was done with the RATS software package, employing the Bernd, Hall, Hall and Hausman (1974) algorithm. The corresponding results for the other two exchange rates are in Appendix A.

[29] Bollerslev (1986), p. 516. We checked these conditions for each model that was estimated. The results were unequivocal.

[30] Bollerslev and Wooldridge (1992), p. 165.

[31] N.B. Interest differentials were scrambled alongside residuals and were taken into account in the calculations of trading rule profitability.

Table 3.6: GARCH-Bootstrapping Results

		P.&Tr.-Prog.	MA(50)	MA(100)	MA(150)	MA(200)	MA(250)
DEM/USD	Mean Orig.	0.1188	0.1000	0.0995	0.0774	0.0731	0.0743
	Mean Boot.	0.0244	-0.0243	-0.0121	-0.0025	0.0055	0.0106
	S. D. Boot.	0.0343	0.0362	0.0366	0.0368	0.0367	0.0365
	Boot. p-val.	0.0050	0.0015	0.0025	0.0165	0.0330	0.0400
JPY/USD	Mean Orig.	0.0846	0.1041	0.0851	0.0995	0.0699	0.0499
	Mean Boot.	0.0158	-0.0290	-0.0167	-0.0072	0.0007	0.0058
	S. D. Boot.	0.0307	0.0310	0.0313	0.0315	0.0314	0.0312
	Boot. p-val.	0.0125	0.0000	0.0015	0.0000	0.0160	0.0785
JPY/DEM	Mean Orig.	0.0789	0.0965	0.0704	0.0750	0.0501	0.0323
	Mean Boot.	0.0197	-0.0292	-0.0170	-0.0075	0.0005	0.0058
	S. D. Boot.	0.0291	0.0297	0.0300	0.0301	0.0300	0.0299
	Boot. p-val.	0.0230	0.0000	0.0040	0.0045	0.0515	0.1775

that returns from the peak-and-trough-progression rule and from moving average trading rules with parameters smaller or equal to 200 remain significant even if conditional heteroskedasticity is taken into account. It is interesting to note that mean returns for the simulated time series increase with increasing length of the moving average parameter. The reason for this is that moving average rules with larger parameters smooth time series more, which leads to less frequent trading and therefore less transaction costs incurred. The fact that mean returns for the MA(250) rule are greater than zero suggests that the autoregressive terms in (3.3) introduce time series dependencies which can be exploited by means of technical trading rules. However, given the comparatively small magnitude of the mean of simulated returns, it is clear that the assumed model captures only a small part of those time dependencies in exchange rates, which are the root of technical trading rule profitability.

3.3.3 Significance Tests Using Markov Switching Bootstraps

As a second model class for exchange rate changes, which might be capable of accounting for the observed trading rule returns, we want to consider Markov switching models.[32] In Markov switching models the existence of an unobserved variable, referred to as the 'state' or 'regime', is postulated. In the simplest case, this variable can assume two values or states. The evolution

[32] This class of models was first introduced in Hamilton (1989) and applied to foreign exchange markets in Engel and Hamilton (1990).

of states over time is governed by a Markov chain and the distribution of exchange rate changes at any date depends on which state prevails. As in Engel and Hamilton (1990) we will assume that exchange rate changes are independently and identically distributed and come from either of two normal distributions with (potentially) different means and variances. The model we will consider can thus be written as follows:

$$\Delta s_t \sim \begin{cases} N(\mu_0, \sigma_0^2) & if\ Z_t = 0 \\ N(\mu_1, \sigma_1^2) & if\ Z_t = 1 \end{cases}$$

(3.4)

	$Pr(Z_t = 0)$	$Pr(Z_t = 1)$
$Z_{t-1} = 0$	p_0	$(1 - p_0)$
$Z_{t-1} = 1$	$(1 - p_1)$	p_1

Z_t denotes the state prevailing at date t, and $\mu_.$ and $\sigma_.$ denote the parameters of the distributions governing exchange rate changes.

A priori, the decision to try to replicate trading rule returns using Markov switching models is motivated by a number of considerations. As outlined in Chapter 2, technical traders believe in and try to exploit the existence of trends in exchange rates. A simple Markov switching model like the one above is capable of generating trends if the means of the distributions generating exchange rate changes are of opposite sign and if states tend to persist (p_0 and p_1 are large). Moreover, in Markov switching models the current state of the unobserved variable depends on past prices, which is consistent with one of the basic assumptions of technical analysis. Apart from the fact that Markov switching models have been found to fit exchange rate data in sample quite well, Engel's (1994) finding that Markov switching models are only better at forecasting the direction of exchange rate changes than the random walk model seems also to fit in well with the fact that technical trading rules only give binary signals.

(3.4) was estimated by maximum likelihood, using the EM algorithm described in Hamilton (1989).[33] Table 3.7 gives the estimation results for each of the three exchange rates. We find evidence of a difference between states regarding the sign of the means of exchange rate changes

[33] Estimation was carried out with the GAUSS software package using the programs that Engel and Hamilton used for their 1990 paper.

Table 3.7: Parameter Estimates of Markov Switching Model of Daily Exchange Rate Changes

		μ_0	μ_1	p_0	p_1	σ_0	σ_1
DEM/USD	Estimate	0.0023	0.0004	0.9769	0.9732	0.0434	0.1633
	S.E.	0.0059	0.0123	0.0054	0.0064	0.0026	0.0082
	t-value	0.3824	0.0312	181.51	151.50	16.84	19.84
JPY/USD	Estimate	-0.023	0.0359	0.9105	0.9033	0.0284	0.1355
	S.E.	0.0051	0.0103	0.0162	0.0185	0.0023	0.007
	t-value	-4.49	3.492	56.35	48.90	12.29	19.23
JPY/DEM	Estimate	-0.005	0.0131	0.9730	0.9722	0.0226	0.1005
	S.E.	0.0041	0.0077	0.0061	0.0066	0.0013	0.0044
	t-value	-1.200	1.700	158.7	148.44	17.50	22.85

Table 3.8: Markov Switching Bootstrapping Results

		P.&Tr. Prog.	MA(50)	MA(100)	MA(150)	MA(200)	MA(250)
DEM/USD	Mean Orig.	0.1188	0.1000	0.0995	0.0774	0.0731	0.0743
	Mean Boot.	-0.0183	-0.0387	-0.0264	-0.0171	-0.0094	-0.0045
	S. D. Boot.	0.0325	0.0336	0.0340	0.0341	0.0341	0.0340
	Boot. p-val.	0.0000	0.0000	0.0000	0.0030	0.0070	0.0090
JPY/USD	Mean Orig.	0.0846	0.1041	0.0851	0.0995	0.0699	0.0499
	Mean Boot.	-0.0176	-0.0375	-0.0252	-0.0159	-0.0081	-0.0032
	S. D. Boot.	0.0294	0.0322	0.0325	0.0327	0.0326	0.0324
	Boot. p-val.	0.0005	0.0000	0.0015	0.0005	0.0115	0.0555
JPY/DEM	Mean Orig.	0.0789	0.00965	0.0704	0.0750	0.0501	0.0323
	Mean Boot.	-0.0178	-0.0384	-0.0262	-0.0168	-0.0091	-0.0042
	S. D. Boot.	0.0260	0.0267	0.0272	0.0274	0.0273	0.0272
	Boot. p-val.	0.0000	0.0000	0.0000	0.0005	0.0165	0.0880

only for the JPY/USD exchange rate. For the other two exchange rates the mean rates of change of the exchange rate are not significantly different from zero. In contrast, for all three exchange rates the volatility differs greatly between regimes. Since the probabilities of remaining in a state are very large for both states, this means that if volatility is high (low), it is likely to remain high (low). Thus the estimation results provide further evidence for the clustering of volatility rather than for the existence of segmented time trends in exchange rate changes. Although this does not make the attempt to replicate observed trading rule returns very promising, we nevertheless carried out 2000 bootstrap simulations to analyse the empirical distribution of trading rule returns under the Markov switching model. The results are contained in Table 3.8. Clearly, the Markov switching model considered above is, if anything, even less capable of accounting for the observed profitability of technical trading rules.

Our results both concerning the evidence of segmented time trends as well as the ability of a simple Markov switching model to account for technical trading rule returns are very similar to those of LeBaron (1991), which were based on weekly data. In contrast, Engel and Hamilton (1990) found very encouraging results concerning the existence of segmented time trends or 'long swings' using quarterly data. It seems as if the evidence of time trends for some reason disappears as the frequency of the data used increases. Recently, Dewachter (1997) has been more successful at replicating technical trading rule returns for weekly exchange rates, by considering a latent variable both for the mean and for the variance assuming that they are mutually independent. It is a challenge for further research to examine whether his results remain valid if daily data are used.

In this context it is worth noting that even if one succeeds in finding a Markov switching model that can generate trading rule returns as large as those observed, it is not at all clear what one can deduce from this, since the segmented trends model distinctly lacks a foundation in economic theory. In other words, the question 'why do technical trading rules work' would simply become 'why do Markov switching models fit'. One might have discovered a better time series model for exchange rates, but this does not help one to explain the profitability of technical trading rules.

In conclusion, the evidence presented in this section shows that neither a GARCH(1,1) nor a simple Markov switching model can account for trading rule returns of the size found above. Whilst it is thus unlikely that the observed trading rule returns have come about by chance, it may still be the case that returns from trading rules can be explained as a compensation for bearing substantial risk. This question will be addressed in the following section.

3.4 Economic Significance of Technical Trading Rule Returns

The fundamental problem with attempting to adjust trading rule returns for their risk is that there does not exist a generally agreed upon procedure for doing so. As a consequence, a great variety of methods have been used in the literature on technical trading rule profitability in order to find out if, and if so, in what sense trading rule returns can be interpreted as a compensation for bearing risk. For this reason we will consider five different ways of measuring risk-adjusted profits, including both the most widely used methods as well as some less well known ones.

3.4.1 X-Statistic

One of the earliest attempts of testing for the economic significance of technical trading rule returns is Sweeney (1986). As a benchmark for trading rule returns, Sweeney uses the return from buying and holding foreign currency throughout the sample period. Under the assumption of a constant risk-premium Sweeney uses a static Capital Asset Pricing Model to derive a test-statistic (the so-called X-statistic) for returns in excess of returns from following a Buy-and-Hold strategy. Sweeney's X-statistic is defined as follows:[34]

$$X = \frac{1}{N}\Sigma_{t=1}^{N}\phi_{t-1}(s_t - s_{t-1} + i^*_{t-1} - i_{t-1}) - (1-f)\frac{1}{N}\Sigma_{t=1}^{N}(s_t - s_{t-1} + i^*_{t-1} - i_{t-1}) \tag{3.5}$$

s_t denotes the natural logarithm of the exchange rate S_t, ϕ_t represents the signal of the trading rule examined, which is defined as in the previous chapter with the only difference that ϕ_t assumes either the value +1 or 0 (no short sales) and i_t (i^*_t) are domestic (foreign) interest rates. $(1-f)$ is the proportion of days on which the trading rule signals a long position and $(1/N)\Sigma_{t=1}^{N}(s_t - s_{t-1} + i^*_{t-1} - i_{t-1})$ is the average return from buying and holding the foreign currency. The intuition behind (3.5) is that trading rule returns must only compensate for the foreign currency risk for the proportion of days on which the trader is actually exposed to it. If the variance of exchange rate changes exists and is constant, the sampling distribution of the X statistic will be normal.

Sweeney analyses the profitability of filter rules for the exchange rates of 10 currencies versus the US-Dollar between 1975 and 1980 and finds evidence of significant risk-adjusted returns for each exchange rate. His results are confirmed and extended in Surajaras and Sweeney (1992) using exchange rate data for 15 countries from 1974 to 1986.

It is worth noting that the fact that short sales are not allowed for implies that for each exchange rate trading rule returns can be evaluated in two ways depending on which currency

[34] Sweeney (1986), p.166.

Table 3.9: Results of Sweeney's (1986) X-test of the Significance of Risk-adjusted Returns for the Peak-and-Trough-Progression Rule

	DEM/USD		JPY/USD		JPY/DEM	
Perspective	DEM	USD	JPY	USD	JPY	DEM
X-statistic	2.24E-04	2.93E-04	1.60E-04	1.56E-04	1.52E-04	8.06E-05
S. E.	8.44E-05	8.78E-05	7.71E-05	7.90E-05	6.66E-05	6.96E-05
t-value	2.6606	3.3372	2.0752	1.9741	2.2907	1.1579

is assumed to be the domestic currency. Sweeney (1986) only concentrates on the perspective of a US investor examining whether it is a good idea to use technical trading rules to invest in non-US currencies. Of course, the same calculations could also be carried out from the point of view of a non-US investor considering whether or not to buy US dollars. For the sake of completeness, we will report the values of risk-adjusted returns for both perspectives in our own calculations.

Table 3.9 shows the results of adjusting trading rule returns using Sweeney's method for all three exchange rates for the Peak-and-trough-progression rule, whilst Figure 3-7 shows the corresponding results for the moving average trading rule for the DEM/USD exchange rate.[35] Clearly, there is strong evidence that trading rule returns remain significant even if they are adjusted for risk. These results do, however, rely on the assumption of a constant foreign exchange risk premium, which is unlikely to be true.[36] For this reason it is necessary to consider also some other measures of risk-adjusted returns.

3.4.2 Sharpe Ratio

The second measure of risk-adjusted returns to be considered is the Sharpe ratio. It is defined as mean returns in excess of the risk-free rate of interest divided by the standard deviation of returns. Customarily, Sharpe ratios are analysed at the annual frequency. Given the mean and standard deviation of daily returns (μ_r, σ_r), the annual Sharpe ratio can be approximated as

[35] The corresponding figures for the other two exchange rates are in Appendix A.

[36] See, e.g., Frankel and Froot (1987).

Figure 3-7: Results of Sweeney's (1986) X-test of the Significance of Risk-adjusted MA Trading Rule Returns for the DEM/USD Exchange Rate

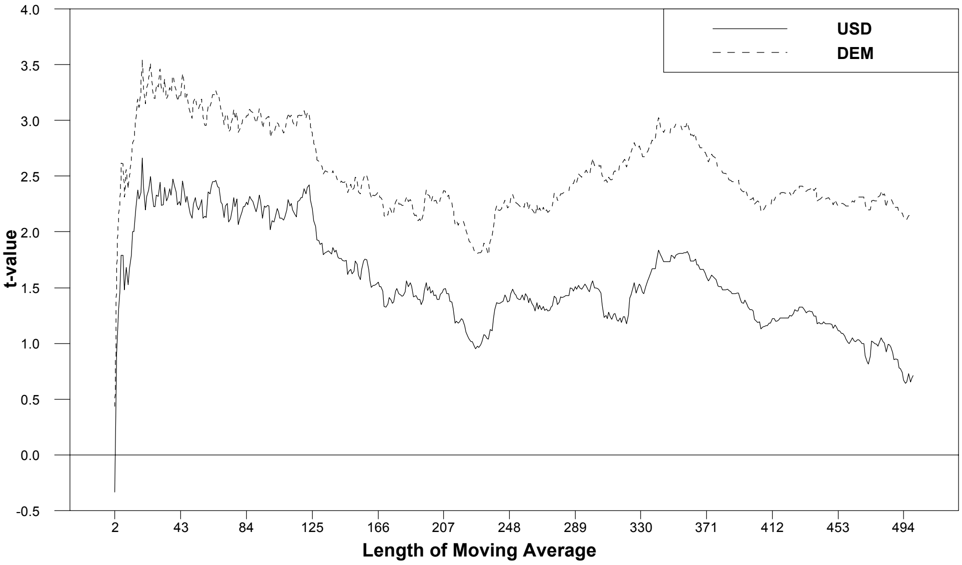

follows:[37] [38]

$$SR = \frac{260\mu_r}{\sqrt{260\sigma_r^2}} = \sqrt{260}\frac{\mu_r}{\sigma_r} \tag{3.6}$$

The question arises with which benchmark the Sharpe ratio for the technical trading rule returns should be compared. Most authors use the Sharpe ratio of buying and holding well diversified stock portfolios as a benchmark. The reason for not considering returns for buying and holding a currency is that these are not well-defined, since the return from buying and holding for instance DEM measured from the perspective of a US investor is the negative of that from buying and holding the USD from the perspective of a German investor. Across

[37] N.B. One need not subtract the risk free rate of interest because "spot speculation of the sort described can be conducted using lines of credit or explicit margin secured by our speculator's initial wealth" (Levich and Thomas (1993), p. 454).

[38] This is only an approximation since it is implicitly assumed that daily rates of return are i.i.d., which is unlikely to be the case (in particular if technical trading rules are profitable). An alternative approach (suggested by LeBaron 1991) to estimate the annual Sharpe ratio would be to calculate the mean and standard deviation of the annual rates of return of all one year (i.e. 260 day) subsamples from the original series of daily rates of return. Since this creates the problem that the first and the last 259 observations are under-represented, we prefer to stick to our approximation even though it is based on an unrealistic assumption.

'perspectives' the average Sharpe ratio is therefore zero. Given that testing whether a Sharpe ratio is greater than zero is equivalent to testing whether mean returns are zero (because of the non-negativity of the standard deviation), risk-adjustment by means of the Sharpe ratio thus makes little sense.

The benchmark annual Sharpe ratio for stock indices and aggregate stock portfolios lies between 0.3 and 0.4.[39] Osler and Chang (1995) calculate a figure of 0.32 for the SP500 over a sample ranging from 1973 to 1994. Many authors have found evidence that the Sharpe ratios of technical trading rule returns are considerably larger than this benchmark value.[40] No attempt has been made so far, however, at trying to determine whether Sharpe ratios of trading rule returns are in some sense significantly greater than the benchmark. If one assumes that excess returns from buying and holding stocks are i.i.d. (with finite variance) with a ratio of mean to standard deviation of 0.32, the central limit theorem can be used to show that the probability of finding an annual Sharpe ratio greater than

$$SR_\alpha = 0.32 + z_\alpha \sqrt{\frac{260}{N}} \tag{3.7}$$

in a sample of size N (for large N) is approximately α (where z_α is such that $\Pr(z > z_\alpha) = \alpha$, for z standard normally distributed).[41] For small α one can interpret finding that for a given sample period returns from following a trading rule are greater than SR_α as indicating that one is unlikely to get as favourable a risk-return relationship from buying and holding a well diversified stock portfolio over a period as long as this sample period. In this narrow sense it is possible to speak of economically significant returns.

Table 3.10 shows the results of calculating Sharpe ratios for the Peak-and-trough-progression rule whilst Figure 3-8 shows the corresponding results for the moving average rule for the DEM/USD exchange rate.[42] For the peak-and-trough-progression rule all Sharpe ratios are significant, even if only marginally so in the case of the JPY exchange rates. As regards the moving average rule we find that up to a length of the moving average of about 150 trading rule

[39] LeBaron (1996), p. 6.

[40] E.g. LeBaron (1996), Osler and Chang (1995, 1998) and Menkhoff and Schlumberger (1995).

[41] A proof for this statement is given in Appendix C.

[42] The figures for the other two exchange rates are to be found in Appendix A. Note that the line of the critical Sharpe-ratio is upwards-sloping since for increasing lengths of the moving average the available sample size decreases.

Table 3.10: Sharpe Ratio of Returns from the Peak-and-Trough-Progression Rule

	DEM/USD	JPY/USD	JPY/DEM	SP500	5%-Crit. Val.
Sharpe Ratio	1.0126	0.7952	0.8487	0.32	0.7656

Figure 3-8: Sharpe Ratios of Moving Average Trading Rules Returns for the DEM/USD Exchange Rate

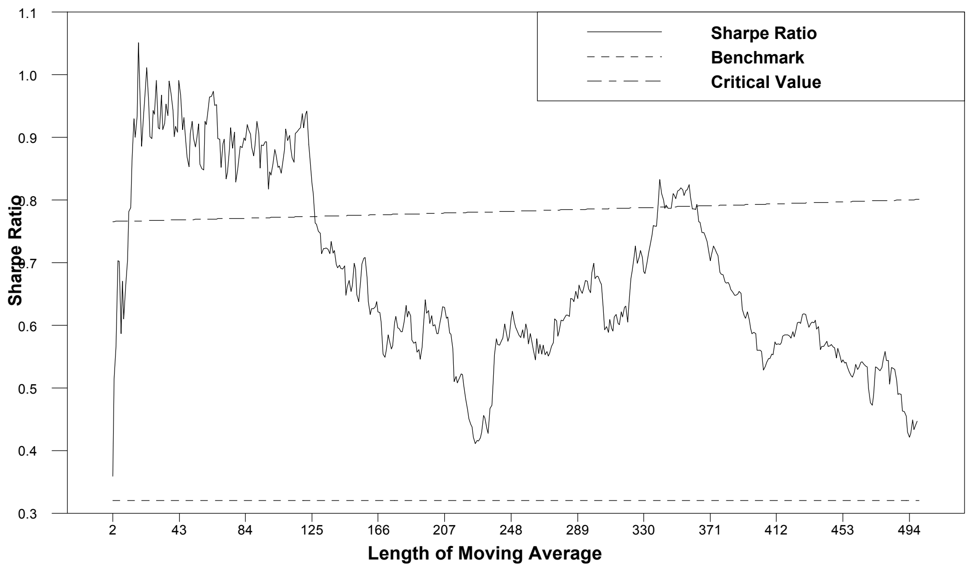

returns are at least marginally significant. For longer lengths there is a general tendency for Sharpe ratios to decrease, with the DEM/USD exchange rate bucking the trend somewhat as was the case for unadjusted trading rule returns. Our results thus indicate that risk-adjusted returns from following the technical trading rules considered are of a magnitude that one is unlikely to realise by buying and holding a well diversified stock portfolio.

3.4.3 Expected Utility

Dacorogna et al. (1991) criticise the use of Sharpe ratios on the grounds that it is a numerically unstable measure for values of the standard deviation near zero. They suggest an alternative measure which is derived assuming normally distributed trading rule returns and constant absolute risk aversion. Under these conditions, maximising expected utility of wealth is equivalent to maximising

$$V = \mu - \frac{C}{2}\sigma^2, \tag{3.8}$$

where μ and σ are the mean and standard deviation of returns and C is the degree of risk aversion. We want to use (3.8) in order to compare trading rule returns with returns from buying and holding a well diversified stock portfolio.[43] Before doing so, we need to decide on a value for C. Dacorogna et al. (1991) argue that C should lie between 8 and 16. This is in broad agreement with Lucas (1987), p. 26: "A value of unity means logarithmic preferences; people seem to be more risk averse than this. No available estimates are as large as 20 , but some do exceed 10." To be on the safe side, we let C vary from 0 to 20. Figure 3-9 shows the values of V both for the peak-and-trough-progression rule for the DEM/USD exchange rate as well as for SP500 and German stock index (DAX) returns over the same period of time.[44] The results for the DEM/USD exchange rate imply that any risk-averse investor, who has constant absolute risk aversion and who assumes that returns are normally distributed, will prefer following the trading rule rather than buying and holding either index. The results for the other two exchange rates are less strong (see Appendix A). Nevertheless, we find that any investor with a risk-aversion of at least 2 will prefer following the trading rule. Note that this means that more rather than less risk-averse investors should use technical trading rules rather than buying and holding a stock index.

3.4.4 Risk of Large Annual Losses

The results in the previous section may be criticised on the grounds that they were derived within an unrealistic theoretical framework. A wholly different approach was used by LeBaron (1991), who estimates the probability of getting a return smaller than 5% below the risk-free rate. Whilst this measure is not firmly grounded in financial theory, LeBaron argues that it is nevertheless of interest since it is one about which investors appear to be concerned.[45] LeBaron

---

[43] N.B. Buying and holding a currency is not used as a benchmark for the same reasons as in the case of the Sharpe ratio.

[44] Source of Data: Datastream.

[45] LeBaron (1991), p. 18. One could also motivate examining this measure with reference to the existence of institutional constraints according to which a trader might lose his job if the value of his portfolio decreases by more than a certain percentage.

Figure 3-9: Expected Utility Comparison between Returns from the Peak-and-Trough-Progression Rule and Index Returns for the DEM/USD Exchange Rate

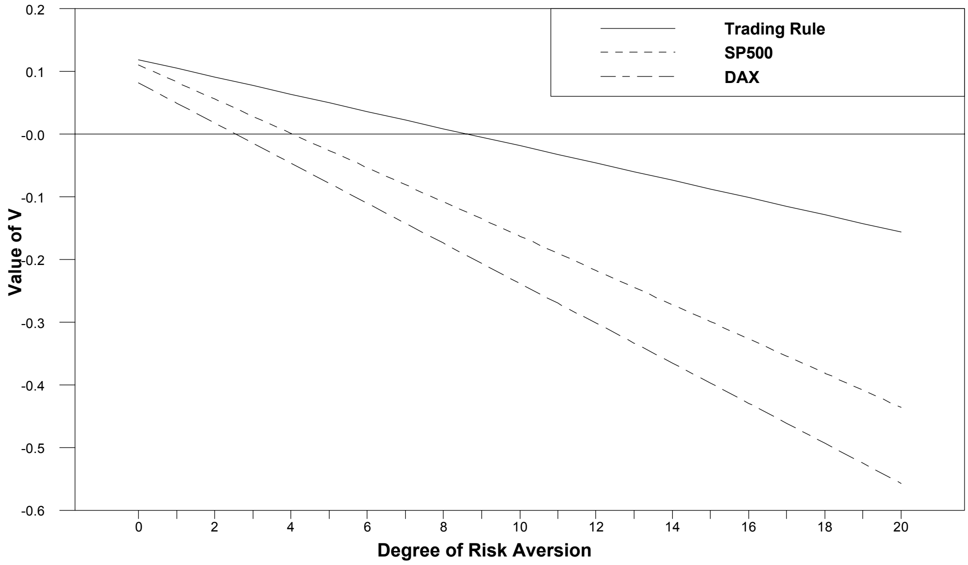

(1991) estimates this probability by randomly sampling one year periods from the trading rule return series. These estimated probabilities are then compared with corresponding estimates for buying and holding the currencies concerned or stock market indices.

Our approach will deviate slightly from LeBaron's. Firstly, we look at the probability of the value of the traders portfolio decreasing by more than 10%.[46] Secondly, rather than carrying out simulations, we will determine limiting probabilities by determining the proportion of one year (260 days) intervals in our sample during which cumulative trading rule returns are smaller than -10%.[47] As a benchmark we will use the corresponding probabilities for buying and holding the relevant national stock market indices as well as those for buying and holding the underlying currencies.[48]

To complement the probability estimates, we also report the maximum annual loss incurred

[46] N.B. The choice of a specific percentage level is, of course, in the last instance always arbitrary. We chose a larger value than LeBaron in order to capture the risk of incurring large losses.

[47] N.B. The influence of interest differentials is taken into consideration.

[48] Thus, for instance, for the DEM/USD exchange rate, we look at the probability of losing more than 10% when holding DEM from the point of view of an US investor and the probability when holding USD from the point of view of a German investor.

Table 3.11: Risk of Large Losses for the Peak-and-Trough-Progression Rule

		Prob.<-10%	Max. Loss
DEM/USD	Pk&Tr-Prog.	0.0016	-0.1146
	Buy-and-Hold USD	0.2396	-0.4074
	Buy-and-Hold DEM	0.3007	-0.4265
JPY/USD	Pk&Tr-Prog.	0.0011	-0.1231
	Buy-and-Hold JPY	0.2495	-0.4499
	Buy-and-Hold USD	0.2454	-0.2757
DEM/USD	Pk&Tr-Prog.	0.0000	-0.0907
	Buy-and-Hold JPY	0.2681	-0.4268
	Buy-and-Hold DEM	0.1506	-0.3192
Indices	SP500	0.0991	-0.2700
	DAX	0.1558	-0.5662
	Nikkei	0.1558	-0.4097

from following the trading rule and from buying and holding indices/currencies. Table 3.11 shows the results for the peak-and-trough-progression rule whilst Figures 3-10 and 3-11 give the results for the moving average rule for the DEM/USD exchange rate.[49] The striking result is that on account of either measure of risk, following the trading rules seems to be less risky than keeping ones funds invested in either currency. This confirms LeBaron's (1991) results and raises the question whether it may be this feature of technical trading rule returns in particular that makes technical analysis so popular amongst financial market practitioners. At any rate, it is clear that our measures of 'drawdown risk' do not allow the identification of trading rule returns as compensation for bearing risk.

3.4.5 Trading Rule Returns as Compensation for Bearing Systematic Risk

The previous analyses have revealed that following either peak-and-trough-progression or moving average rules yields returns which compare favourably with returns from buying and holding well diversified stock portfolios when returns are adjusted for risk in a variety of ways. A number of authors have examined the relationship between trading rule returns and returns from certain proxies for the market portfolio.[50] The point of this exercise was to find out whether at least part of the observed trading rule returns could be interpreted as a compensation for

[49] The corresponding figures for the other two exchange rates are contained in Appendix A.
[50] E.g. Sweeney (1986), Taylor (1992), Neely et al. (1996) and LeBaron (1991).

Figure 3-10: Probability that the Annual Return from Following Moving Average Rules is Smaller than -10% for the DEM/USD Exchange Rate

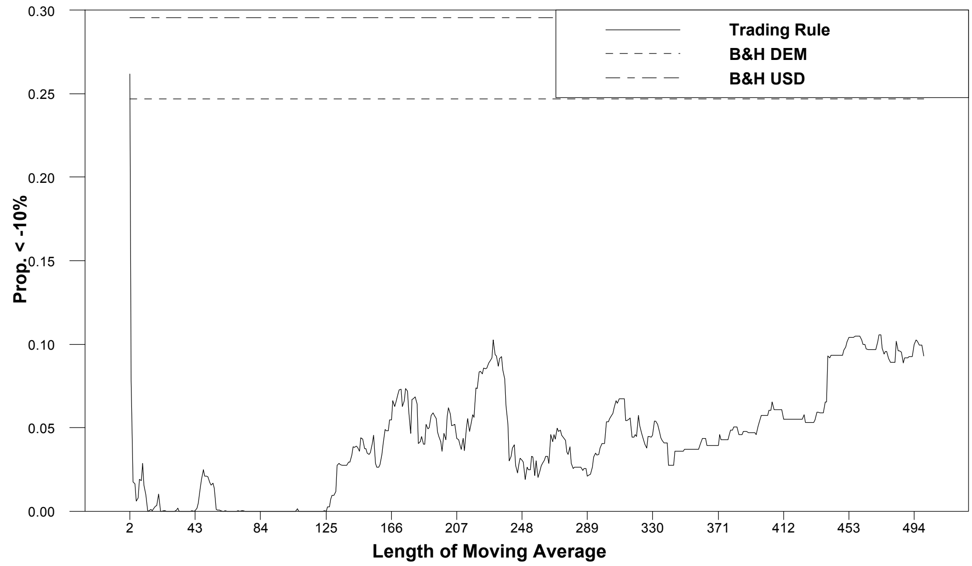

Figure 3-11: Maximum Annual Loss of Moving Average Trading Rules for the DEM/USD Exchange Rate

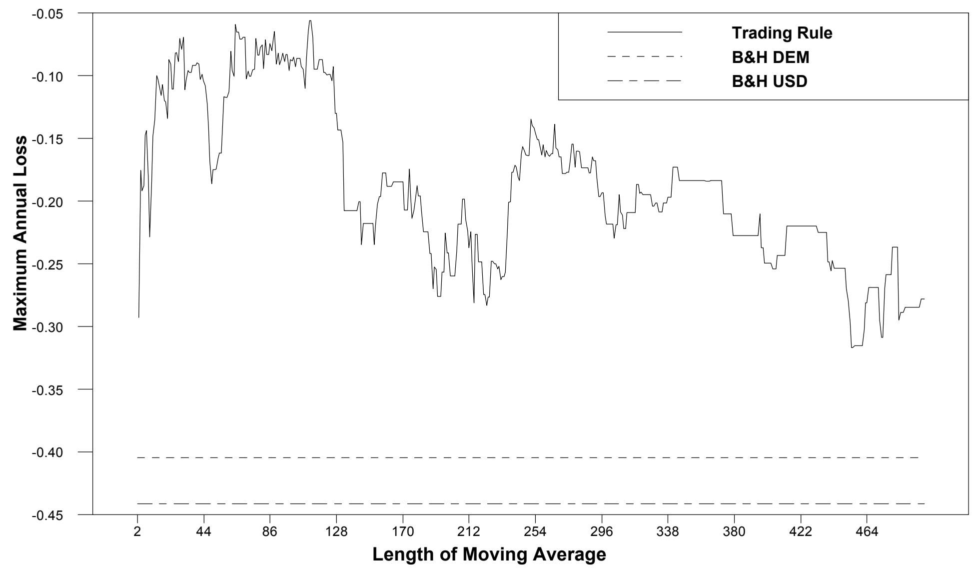

Table 3.12: Estimates of Static CAPM Betas for Peak-and-Trough-Progression Rule Returns

		Estimate of β	S. E.	t-value
DEM/USD	SP500	-0.0315	0.0135	-2.3321
	DAX	-0.0667	0.0157	-4.2506
JPY/USD	SP500	-0.0228	0.0123	-1.8483
	Nikkei	-0.0134	0.0119	-2.7685
JPY/DEM	DAX	-0.0163	0.0100	-1.625
	Nikkei	-0.0123	0.0950	-1.297

bearing systematic risk. We carried out a similar analysis, using for each exchange rate the stock market index returns from the countries concerned as the proxy for the market portfolio. CAPM betas were calculated by estimating the following equation by OLS:[51][52]

$$r_t = \alpha + \beta \widehat{index_t} + u_t \tag{3.9}$$

r_t denotes the trading rule return at t, as defined in (3.2), and $\widehat{index_t}$ denotes the relevant index return. Table 3.12 contains the results of estimating CAPM betas for the peak-and-trough-progression rule for each of the three exchange rates. The corresponding results for the moving average trading rule for the DEM/USD exchange rate are contained in Figure 3-12.[53]

All estimates are negative and many of them are even significantly negative.[54] This confirms the results of LeBaron (1991), Taylor (1992) and Neely et al. (1996). Kho (1996) is somewhat more successful at explaining technical trading rule profitability in terms of systematic risk. He estimates three bivariate GARCH models of weekly foreign exchange futures and MSCI world equity market index returns and uses these to carry out bootstrapping simulations. Average simulated trading rule returns lie between a fifth and a half of the actual returns for the various models. The difference between simulated and actual return is significantly greater than zero

[51] N.B. Instead of the individual stock indices for each country, we also tried a number of variants of combining the indices and of regressing trading rule returns on measures of the relative performance of the stock indices. The results only changed marginally, however.

[52] N.B. The standard errors of the estimates of beta were calculated using White's (1980) autocorrelation and heteroskedasticity consistent covariance estimator.

[53] Results for the other two exchange rates are contained in Appendix A.

[54] It is worth noting, however, that there exists a potentially serious simultaneity bias in the estimates of beta, which might arise from exchange rate changes having an influence on the stock market. Whilst our results (in particular the finding of significant coefficients) should thus be interpreted with caution, they nevertheless indicate strongly that there does not exist a positive relationship between trading rule profits and stock market index returns.

Figure 3-12: Estimates of CAPM Betas for Moving Average Trading Rule Returns for the DEM/USD Exchange Rate

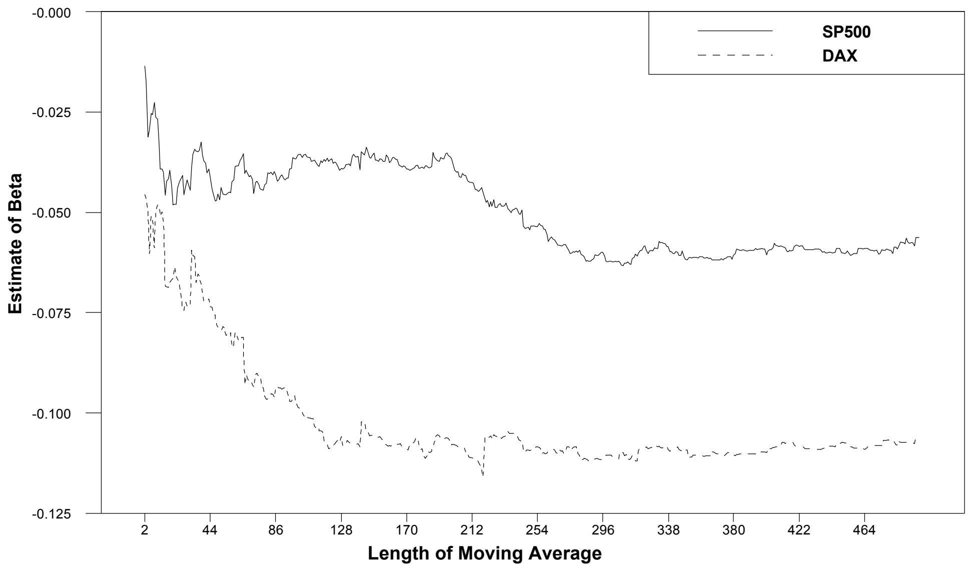

in 20 out of 24 cases (4 exchange rates times 6 trading rules) for the worst model and for only 4 out of 24 cases for the best. Kho concludes: "... my analysis of the risk premia implicit in the measured profits does not support the conclusion that the profits are unusual compared to risk."[55]

This conclusion is slightly premature, however. Even if we consider only his best fitting model, it is important to notice that although trading rule returns are only significant in 4 of the 24 cases, in almost all other cases the bootstrapping p-values lie around 10%. It would have been interesting to see the results of a joint test of the significance of trading rule returns across trading rules. Kho asserts that: "... a joint test across trading rules within each currency can be implemented visually as shown in Fig. 2."[56] This figure, however, only allows to see where actual trading rule returns lie in relation to bootstrapping confidence intervals for each trading rule. The implicit assumption seems to be that a joint test consists in checking whether average actual trading rule returns lie inside or outside the confidence intervals (which do not differ very

[55] Kho (1996), p. 287.
[56] Kho (1996), p. 284f.

much across trading rules). However, this approach will bias the results unless technical trading rule returns are perfectly correlated. Since this is not the case, a correctly specified joint test would tend to lead to stronger results.[57] Given the fact that most trading rule returns are found to be marginally significant, it seems likely that a joint test would indicate the significance of trading rule returns for all four exchange rates considered. It remains to be seen whether this conjecture is true.

On a more general note, one wonders whether the right direction for future research is to try out more and more intricate asset price models with the aim of being able to interpret technical trading rule returns as a compensation for bearing systematic risk. Technical analysis is sometimes criticised on the grounds that if one tortures any time series long enough, one will sooner or later come up with a profitable trading strategy. Given the extensive evidence in favour of the profitability of technical trading rules provided in this and the previous section, maybe this argument can be turned around, in the sense that if one simply tries enough asset pricing models, one will sooner or later find one which identifies trading rule returns with risk premia.[58] Such a finding might allow one to hold on to a naive belief in the efficiency of financial markets. The question is, however, whether it is not more sensible to try to uncover the mechanisms which make technical analysis profitable.

3.5 Stability of Trading Rule Returns

3.5.1 Evidence of Decreasing Returns

The last two sections provided evidence that following peak-and-trough-progression and moving average trading rules yields statistically and economically significant returns. This leads us to the question of whether trading rule returns have been stable over time. In the context of trading rule returns, it is not entirely clear what stability should mean. If we were to interpret stability as the stability of the mean daily returns, we would indubitably find periods with differing means. This would, however, not be a particularly interesting result, as all manuals of technical analysis emphasise that their tools are not infallible, but rather help to tilt the odds

[57] See Section 5.4 for a detailed description of how such a test can be conducted.
[58] See Krugman and Miller (1992), p.20f., for an argument to the same effect.

Table 3.13: Estimates of a Time Trend in Peak-and-Trough-Progression Rule Returns

		α	β
DEM/USD	Estimate	2.95E-04	2.37E-05
	S. E.	2.28E-04	3.04E-05
	t-value	1.296	0.7791
JPY/USD	Estimate	4.56E-04	-1.92E-05
	S. E.	2.23E-04	2.77E-05
	t-value	2.047	-0.692
JPY/DEM	Estimate	4.36E-04	-1.94E-05
	S. E.	1.94E-04	2.46E-05
	t-value	2.248	-0.7879

in one's favour. In a more relevant sense, stability means that trading rule returns should not decrease over time. We will use two measures to examine whether or not this is the case. To start off with, we follow LeBaron's (1991) approach and look at one-year rolling returns, which are defined as follows:

$$r_t^{roll} := \Sigma_{i=t-129}^{t+130} r_{t+i} \tag{3.10}$$

Figure 3-13 shows the results for the peak-and-trough-progression rule for each of the three exchange rates.[59] The figure reveals that whilst there is a great deal of variability in annual returns over the period there is no indication of decreasing returns over time. We also estimated the following equation in order to test for a deterministic time trend in trading rule returns:

$$r_t = \alpha + \beta t + u_t \tag{3.11}$$

Table 3.13 shows the results of estimating (3.11) for the peak-and-trough-progression rule. For no exchange rate are the parameters significantly smaller than zero - for the DEM/USD exchange rate the coefficient is even greater than zero. There thus appears to be no tendency for technical trading rule returns to have decreased during the period examined, which confirms the results of Taylor (1992) and Kho (1996).

[59] For the sake of clarity, we only report the results for the peak-and-trough-progression rule in this section. The corresponding results for the moving average trading rule look very similar; see, for instance, LeBaron (1996), Fig. 2.

Figure 3-13: Rolling One-Year Returns from Following the Peak-and-Trough-Progression Rule

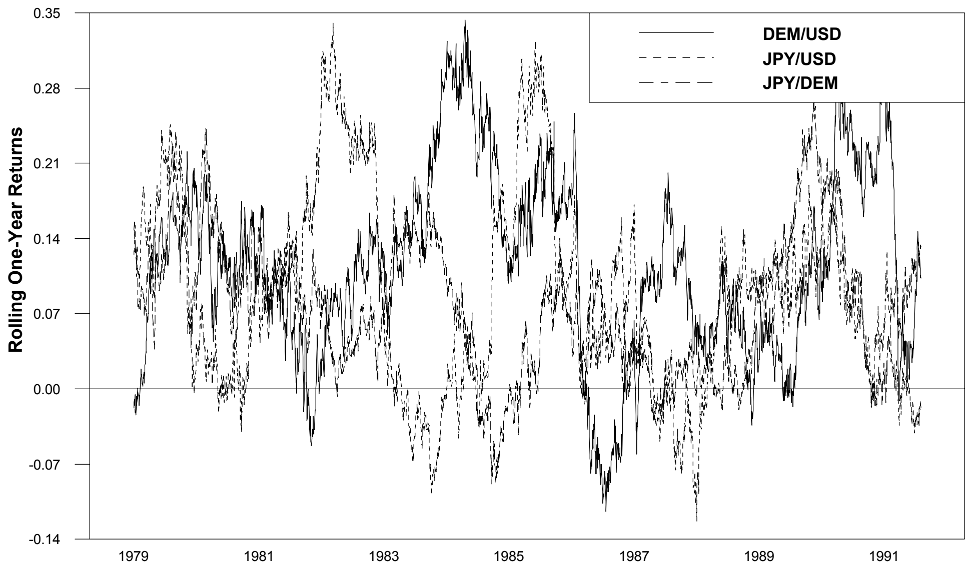

3.5.2 Stability of the Relative Profitability of Technical Trading Rules

Another aspect of the stability of trading rule returns that is of interest concerns the question whether the relative profitability of individual trading rules is stable over time. Both Sweeney (1986) and Menkhoff and Schlumberger (1995) analyse the profitability of those technical trading rules which yielded the highest returns in a first part of their sample in an out-of-sample period. Sweeney (1986) finds that out of 23 trading rules that yielded significant returns in the first sample only 8 produced significant returns in the second sample.[60] Menkhoff and Schlumberger (1995) calculate Spearman rank correlation coefficients between the ranks of the returns of 33 moving average and 10 momentum trading rules for the first and second halves of their sample. The coefficients are never significant and are as often negative as they are positive.[61] These results suggest that the past performance of individual trading rules is not a

[60] N.B. Sweeney's (1986) out-of-sample period is about twice as large as his in-sample period (620 versus 1220 observations). We strongly suspect that the number of significant trading rule returns would have been even smaller if the sample had been split up (more naturally) into two equally large subsamples (if only because of the greater standard error of mean returns).

[61] It is worth noting that Menkhoff and Schlumberger's (1995) test of the significance of the correlation in profitability ranks is not well-defined because it overlooks that returns from following different moving average or momentum trading rules are positively correlated. However, since taking account of this correlation would

Table 3.14: Profitability of MA Trading Rules in the 1st versus 2nd Half of the Sample

	# Sig. in 1st Half (Out of 499)	# Sig. in 2nd Half (Out of 499)	# Sig. in Both Halves	# Sig. in 2nd Half / # Sig. in 1st Half
DEM/USD	346	53	53	0.1532
JPY/USD	266	47	47	0.1767
JPY/DEM	137	43	37	0.2701

good indicator for their future profitability.

As a starting point we will split our sample in two halves and examine which proportion of moving average trading rules that yielded significant returns in the first half of the sample also yielded significant returns in the second half. The results are contained in Table 3.14. The first point to notice is the marked drop in the number of significant trading rules in the second half of the sample.[62] Of these rules practically all had already yielded significant returns in the first half of the sample. However, this would have been of little help for someone who had to make a decision at the end of the first half of the sample (30.6.1986) about which trading rule to employ in the second half. If a rule had been chosen randomly amongst those that had yielded significant returns in the first half, in only between 15% and 30% of the cases would this have led to significant returns in the second half.

Of course, splitting up the sample in two halves is in the last instance arbitrary. Moreover, one might think that it will tend to bias the results downwards because the fictitious investor is given the chance to choose a trading rule only once. In order to analyse the stability of the relative performance of individual trading rules further, we therefore examine whether always following that individual trading rule that has performed best in the past yields higher returns than following randomly chosen rules. Examining such 'optimising' trading rules requires making two decisions: One must decide firstly, how often the relative performance is evaluated and secondly, over which sample of past prices the performance of individual trading rules should be compared. We will assume that the trader optimises (in the sense of determining the best rule) daily. The decision concerning the number of past daily returns, over which the trader optimises is more problematic. If there is a time-varying structure in the relative profitability

tend to increase the critical values of this test, Menkhoff and Schlumberger's (1995) conclusion of no correlation between the profitability ranks still goes through.

[62] N.B. Of course this does not imply that trading rule profitability has decreased over time, as was evidenced in the previous section.

Table 3.15: Profitability of Following Randomly Chosen MA Trading Rules

	DEM/USD	JPY/USD	JPY/DEM
Mean	0.0883	0.0812	0.0503
Std. Dev.	0.0200	0.0169	0.0148

of moving average trading rules, it is just as conceivable that it changes at a weekly as at an annual frequency. In order to allow for all possibilities, we let the length of the selection period vary from twenty days (i.e. approximately one month) to 520 days (i.e. about two years). Since this examination is extremely computer-intensive, we restricted the sample of possible trading rules to moving average rules with a moving average up to 250.[63] As a benchmark we also calculate the profitability of randomly choosing a moving average trading rule each period anew by generating 2000 such rules and determining mean and standard deviation of average annual returns of these random rules. Table 3.15 contains the results of following randomly chosen trading rules for all three exchange rates, whilst Figure 3-14 shows average annual returns from following 'optimal' trading rules as a function of the length of the horizon over which the rules are chosen for the DEM/USD exchange rate.[64]

The results suggest that the strategy of always using the moving average rule that has worked best in the recent past works, if anything, worse than using a random rule. An explanation for this result might be that optimisation is more likely to lead to a choice of very high or very low parameter values (because of overfitting), for which the moving average trading rule does not work very well. Moreover, Figure 3-14 also shows that even small changes in the length of the selection period often lead to very large changes in the return from following this strategy.

When discussing the method to be used to measure technical trading rule profitability at the beginning of this chapter, we mentioned that one way to avoid a potential ex-post selection bias in measuring technical trading rule returns was to endogenise the choice of rule by looking at those rules which have been most successful in some selection period. It is precisely the above instability of trading rule returns with respect to changes in the length of the selection period which made us choose the approach of leaving the parameter values exogenous and simply reporting annual returns for all reasonable parameter values.

[63] N.B. As before, 250 is chosen as the upper limit since it the highest value of the parameter of the rules considered in both the academic and the commercial literature on technical analysis.

[64] The corresponding figures for the other two exchange rates are contained in Appendix A.

Figure 3-14: Profitability of Following Optimally versus Randomly Chosen MA Trading Rules for the DEM/USD Exchange Rate

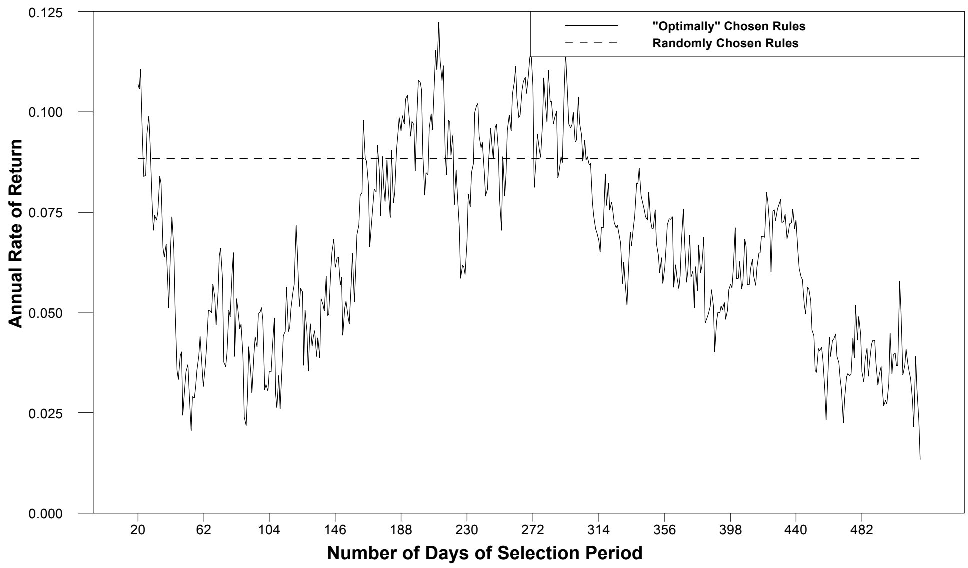

In conclusion, therefore, we found no evidence of a decrease in trading rule profitability over time. Whilst there is a distinct lack of stability concerning the relative performance of individual trading rules, the observed trading rule profitability appears nevertheless to be exploitable by means of following randomly chosen moving average trading rules.

3.6 Further Anomalies

3.6.1 Trading Rule Returns and Uncovered Interest Parity

It is worth noting a connection between our results concerning the profitability of technical trading rules and those of studies examining the validity of uncovered interest parity (UIP). UIP states that if the interest rate in country A is $x\%$ greater than that of country B at a horizon of K months, then the currency of country A should depreciate on average by $x\%$ over the same period. The intuition behind UIP is that systematic deviations from UIP can be exploited by rational risk-neutral speculators. UIP has been tested over an over again and the

Table 3.16: Components of Peak-and-Trough-Progression Rule Returns for Flexible Exchange Rates

	DEM/USD	JPY/USD	JPY/DEM
Gross Returns	0.1288	0.0934	0.0906
S. E.	1.21E-04	1.10E-04	9.55E-05
t-value	4.0938	3.2780	3.6477
Transaction Costs	-0.0140	-0.0144	-0.0145
S. E.	2.61E-06	2.64E-06	2.64E-06
t-value	-20.71	-21.01	-21.10
Interest Effects	0.0040	0.0056	0.0028
S. E.	2.74E-06	2.91E-06	1.51E-06
t-value	5.6163	7.3870	7.0729

results have practically always been negative.[65] Not only does the interest differential predict future changes in the exchange rate poorly, but it does not even tend to predict the direction of the change in the exchange rate correctly (i.e. its predictive accuracy was worse than that of a random walk). In order to analyse the relationship between deviations from UIP and technical trading rule returns, we split up net annual returns from following a trading rule into three components: Gross returns, transaction costs and interest effects. Table 3.16 shows the results of the decomposition for the peak-and-trough-progression rule whilst Figure 3-15 shows the average annual interest effect for the moving average trading rule for all three exchange rates.[66]

With the exception of the case of the JPY/DEM exchange rate for lengths of the moving average greater than 340, the interest differential tends to contribute to rather than to reduce trading rule profitability for all exchange rates. Moreover, the t-values reported for the peak-and-trough-progression rule indicate that the observed positive interest effects are unlikely to have come about by chance. This means that interest differentials not only do not compensate for predictable components in exchange rate changes, but that instead deviations from UIP are systematically exploited by technical trading rules.

If one assumes that deviations from UIP are due to the existence of time-varying risk-premia, it might be tempting to interpret these results as an indication that at least part of the observed trading rule returns are a compensation for bearing currency risk. However, there are two considerations which make it unlikely that this suggestion is true. Firstly, our

[65] See Lewis (1995b) for a survey over the literature.

[66] Gross returns and transaction costs for moving average trading rules are shown in Appendix A.

Figure 3-15: Average Annual Interest Earnings/Losses from Following MA Trading Rules for Flexible Exchange Rates

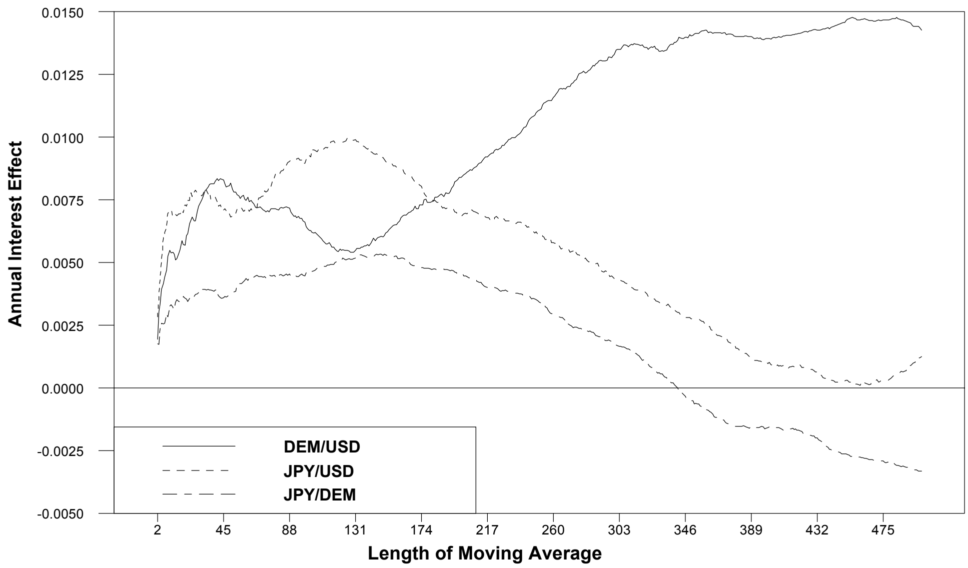

results concerning the economic significance of trading rule returns in Section 3.4 suggest that following technical trading rules on foreign exchange markets tends to be less risky than buying and holding a currency. Secondly, this interpretation is not consistent with the results of a study by Frankel and Froot (1987), which uses survey data on exchange rate expectations to split up deviations from UIP into a risk-premium and an expectational error component. Frankel and Froot (1987) find that expectational errors tend to have the same sign as and tend to be even larger than deviations from UIP.[67] This implies that the risk-premium has the opposite sign of the deviation from UIP and indicates that exploiting deviations from UIP is more likely to entail holding the currency that is less rather than more risky. In addition, the results of Frankel and Froot (1987) suggest that both deviation from UIP and technical trading rule profitability are caused by systematic expectational errors.

[67] Frankel and Froot (1987), Tables 1-3. See also Blake et al. (1986).

3.6.2 Buy/Sell-Asymmetries and Day-of-the-Week Effects

When carrying out their analysis of the profitability of technical trading rules for daily Dow Jones Industrial Index data, Brock Lakonishok and LeBaron (1992) found strong evidence of an asymmetry between the trading rule returns depending on whether the trading rule signalled a long or a short position. In contrast, LeBaron (1991) found no such evidence for weekly exchange rate data. In order to find out whether the reason for this disagreement is the use of exchange rate rather than stock market data or instead is the use of weekly rather than daily data, we report trading rule returns on days on which the peak-and-trough-progression rule signalled long and short positions in Table 3.17. For none of the currencies is the difference significantly different from zero at a 5% significance level. Moreover, there does not seem to be a large difference in the volatility between long and short periods. There thus appears to be a difference in the structure of technical trading rule returns between stock- and foreign exchange markets. We suspect that this difference might have to do with the fact that stock markets are characterised by a much more pronounced long term trend than are foreign exchange markets.

Szakmary and Mathur (1997) report that there exist large differences between the profitability of technical trading rules on different days of the week. They find that trading rule returns are significantly positive on Mondays and Fridays and either insignificant or significantly negative on all other days. Table 3.17 also contains the results of examining returns separately for each day of the week. Clearly, Szakmary and Mathur's (1997) result is confirmed. Note also that the volatility on Mondays and Fridays is larger than on the other days of the week.

In order to see whether these results are robust, we carried out an analogous examination for the moving average trading rule.[68] Figure 3-16 shows average annual returns for each day of the week for the USD/DEM exchange rate.[69] Whilst there is some evidence that Monday returns decrease for very long lengths of the moving average, for lengths of 200 and less (which are allegedly most popular in practice) exceptionally high trading rule returns on Mondays are also found for the moving average trading rule. In contrast, the result that returns are also large on Fridays does not turn out to be robust.

[68] N.B. We only report results concerning the day-of-the-week effect. Examining the difference in mean returns between long and short periods for moving average trading rule yielded no new insights.

[69] The corresponding pictures for JPY/USD and JPY/DEM are in Appendix A.

Table 3.17: Long/Short- and Day-of-the-Week Effects of the Peak-and-Trough-Progression Rule

		Long	Short	Diff.	Mon	Tue	Wed	Thu	Fri
DEM/USD	Mean	0.1308	0.1075	0.0232	0.5460	-0.156	0.0827	0.0385	0.1146
	S. D.	0.1171	0.1175	n.a.	0.1283	0.1170	0.1138	0.1040	0.1189
	S. E.	0.0456	0.0443	0.064	0.0801	0.0702	0.0682	0.0626	0.0718
	t-val.	2.8652	2.4295	0.3652	6.8192	-2.230	1.2125	0.6149	1.5968
	# Obs.	1711	1832	121	668	723	724	716	713
JPY/USD	Mean	0.1129	0.0552	0.0577	0.3319	-0.038	0.0283	-0.012	0.1313
	S. D.	0.0993	0.1134	n.a.	0.1175	0.1038	0.1035	0.0980	0.1080
	S. E.	0.0377	0.0439	0.058	0.0733	0.0622	0.0620	0.0591	0.0652
	t-val.	2.9984	1.2583	0.9986	4.5289	-0.615	0.4563	-0.196	2.0139
	# Obs.	1806	1737	69	668	723	724	716	713
JPY/DEM	Mean	0.0917	0.0666	0.0251	0.2763	-0.008	0.0429	-0.021	0.1180
	S. D.	0.0886	0.0968	n.a.	0.1009	0.0906	0.0929	0.0863	0.0930
	S. E.	0.0344	0.0367	0.0503	0.0630	0.0543	0.0557	0.0520	0.0562
	t-val.	2.666	1.814	0.4994	4.390	-0.141	0.7701	-0.398	2.0102
	# Obs.	1738	1805	67	668	723	724	716	713

Mean stands for mean annual return, S. D. for standard deviation of annual returns, S.E. for standard error of annual returns and z-value gives the values of the test-statistic $\frac{Mean-0}{S.E.}$.

* The standard error of the difference in means is estimated as $\sqrt{S.E._{Long}^2 + S.E._{Short}^2}$.

Figure 3-16: Moving Average Trading Rule Returns for Different Days of the Week

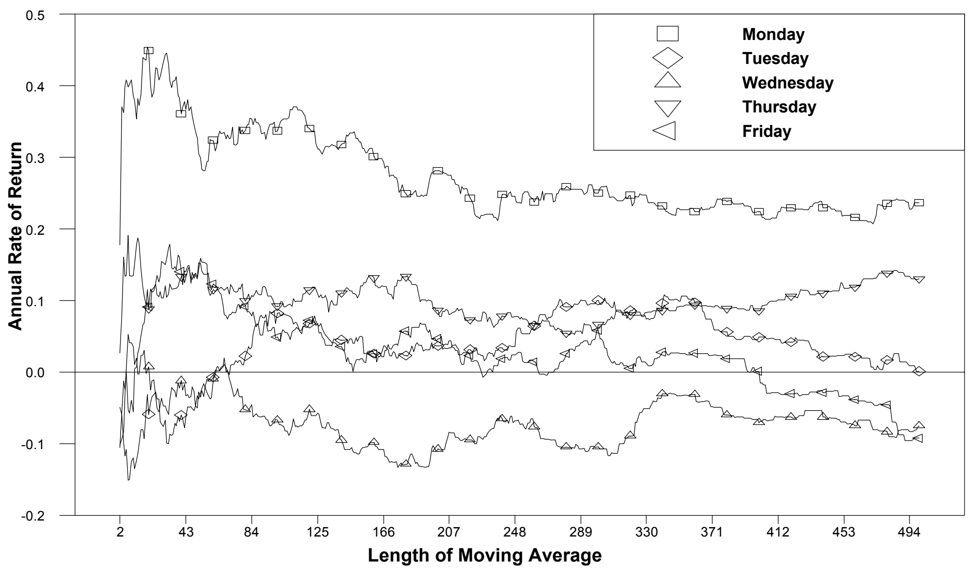

Given these results, one question which immediately arises is whether such a day-of-the-week effect can be found in the original exchange rate data. Cornett et al. (1995) carry out a very thorough analysis of day-of-the-week effects and other seasonalities in the foreign exchange futures market. When regressing daily close-to-close exchange rate changes on day-of-the-week dummies, they find no evidence of patterns in day-of-the-week returns. A replication of their results with our dataset did not yield any new insights.[70] However, in the context of technical trading rule profitability it is maybe more relevant to examine the autocorrelation in exchange rate changes on different days of the week. The reason for this is that if we regress exchange rate changes on day-of-the-week dummy variables, we only find out whether an exchange rate has a tendency to depreciate or appreciate on a particular day of the week. What we are more interested in, however, is whether trends are more likely to be continued on some days than on others. To investigate this issue we estimated the following equation:[71]

$$\Delta s_t = \alpha + (\beta_1 D_{1t} + \beta_2 D_{2t} + \beta_3 D_{3t} + \beta_4 D_{4t} + \beta_5 D_{5t})\Delta s_{t-1} + u_t, \tag{3.12}$$

where the D_{it} are day-of-the-week dummies. Table 3.18 contains the results.[72] For all exchange rates there is evidence of positive autocorrelation between exchange rate changes from Fridays to Mondays. In contrast, evidence of autocorrelation between Thursdays and Fridays is mixed.[73] In order to assess to what extent the profitability of technical trading rules is due to exploiting this 'week-end' effect, we carried out bootstrapping simulations with (3.12) as a null model.[74] Table 3.19 shows the mean and standard deviation of applying the peak-and-trough-progression rule as well as the MA(50), the MA(100), the MA(150), the MA(200) and the MA(250) trading rules to 2000 simulated exchange rate 'time series'.

The results show that our model can account only for a small proportion of the profitability of technical trading rules in general and of the day-of-the-week pattern in trading rule returns

[70] The results are contained in Table B.5 in Appendix B.

[71] N.B. Standard errors were calculated using the White's (1980) autocorrelation and heteroskedasticity consistent covariance estimator.

[72] N.B. For all exchange rates, the Durbin Watson test statistic was between 1.99 and 2. Moreover, Ljung-Box Q-statistics were in no case significant.

[73] N.B. These results are quite similar to those of Keim and Stambaugh (1984) for stock market index returns.

[74] N.B. Residuals were split up into Monday residuals, Tuesday residuals, etc. in order to allow separate resampling for each day of the week.

Table 3.18: Day-of-the-Week Autocorrelation Patterns in Exchange Rate Changes for Flexible Exchange Rates

	Parameter	Estimate	S.E.	t-value
DEM/USD	α	-3.70E-05	1.21E-04	-0.3050
	β_1	0.1330	0.0532	2.5009
	β_2	-0.0622	0.0407	-1.5256
	β_3	-0.0119	0.0406	-0.2925
	β_4	-0.0373	0.0446	-0.8366
	β_5	0.0693	0.0507	1.3687
JPY/USD	α	-1.24E-04	1.10E-04	-1.1252
	β_1	0.1224	0.0488	2.5101
	β_2	-0.0528	0.0359	-1.4711
	β_3	0.0166	0.0440	0.3774
	β_4	-0.0534	0.0362	-1.4756
	β_5	0.0631	0.0485	1.3004
JPY/DEM	α	-7.92E-05	9.51E-05	-0.8336
	β_1	0.1691	0.0548	3.0883
	β_2	0.0763	0.0391	1.9529
	β_3	0.0469	0.0482	0.9735
	β_4	-0.0125	0.0474	-0.2634
	β_5	0.1278	0.0505	2.5312

Table 3.19: Day-of-Week-Bootstrapping Results

		Pk&Tr-Prog.	MA(50)	MA(100)	MA(150)	MA(200)	MA(250)
DEM/USD	Mean Orig.	0.1188	0.1000	0.0995	0.0774	0.0731	0.0743
	Mean Boot.	-0.0086	-0.0364	-0.0243	-0.0148	-0.0071	-0.0022
	S. D. Boot.	0.0317	0.0327	0.0331	0.0333	0.0332	0.0331
	Boot. p-val.	0.0000	0.0000	0.0000	0.0020	0.0095	0.0115
JPY/USD	Mean Orig.	0.0846	0.1041	0.0851	0.0995	0.0699	0.0499
	Mean Boot.	-0.0113	-0.0377	-0.0255	-0.0161	-0.0084	-0.0036
	S. D. Boot.	0.0294	0.0298	0.0301	0.0303	0.0302	0.0301
	Boot. p-val.	0.0000	0.0000	0.0000	0.0000	0.0035	0.0365
JPY/DEM	Mean Orig.	0.0789	0.0965	0.0704	0.0750	0.0501	0.0323
	Mean Boot.	0.0243	-0.0239	-0.0117	-0.0022	0.0055	0.0105
	S. D. Boot.	0.0263	0.0275	0.0278	0.0280	0.0278	0.0276
	Boot. p-val.	0.0180	5.00E-04	0.0035	0.0045	0.0515	0.2120

in particular. It remains a challenge for future research to find out firstly, who or what causes the day-of-the-week autocorrelation pattern in exchange rate changes and secondly, how this relates to the profitability of technical trading rules.

3.7 Discussion

This chapter provided evidence for the profitability of simple technical trading rules in foreign exchange markets. In order to ensure that trading rule returns could indeed have been realised by a trader who had to make a decision ex-ante about which specific trading rule was to be used, various ways have been proposed in the literature to make the choice dependent on the past performance of trading rules. In contrast, the approach adopted in this chapter was to consider trading rules that are known to have been in wide use at the time our sample starts and to analyse their profitability for all reasonable parameter values. As a justification for this procedure we provided evidence that always following that individual trading rule which has yielded the highest returns in the recent past produces returns that are very sensitive to changes in the length of the period over which the optimal rule was determined. We take this as an indication that there exist few stable time dependencies in exchange rate changes beyond the very general one which is exploited by technical trading rules. The fact that trading rules based on sophisticated econometric methods or genetic programming do not yield higher returns than the simple trading rule considered in this chapter supports this interpretation.

Our analysis of the profitability of peak-and-trough-progression and moving average trading rules showed that both trading rules yield returns which are significantly greater than zero. In particular, we provided evidence that this result holds even when conditional heteroskedasticity and potential regime switches in exchange rate changes were allowed for. In order to ascertain whether the observed trading rule returns can be seen as a compensation for bearing high levels of risk, we analysed trading rule returns using a whole variety of approaches for measuring/adjusting for risk. On account of none of these measure did trading rules emerge as more risky than buying and holding either currency or stock market indices. To the contrary, we found strong indications that following the trading rules is less risky than the Buy-and-Hold alternatives. We also examined the change of trading rule returns over time, but found no indications that the profitability was decreasing or even disappearing. These results immediately raise the question why technical trading rules are profitable. One answer to this question is examined in the next chapter.

Chapter 4

Technical Analysis and Central Bank Interventions

4.1 Introduction

A number of authors have conjectured that one reason for the profitability of technical trading rules in flexible exchange rate markets may be the interventions of central banks, who are assumed to follow other objectives than profit-maximisation with their transactions. In which way precisely central banks would introduce patterns into the movement of exchange rates that can be exploited by means of technical trading rules has never been shown. Sweeney (1986) just notes that "[i]ll-conceived government intervention can create profit opportunities such as those found above." The most explicit formulation of the connection between central bank interventions and technical trading rules can be found in Szakmary and Mathur (1997), p. 514:

> If central banks smooth out changes in exchange rates and delay adjustment to underlying fundamental forces by leaning against the wind, it may be expected that trend-following forecasters profit from interventions.

The story seems to run as follows: After an exogenous shock to fundamentals, the exchange rate would, without central bank interventions, jump to a new equilibrium level (e.g., Dornbusch overshooting). Wishing to reduce volatility, central banks try to prevent the exchange rate from jumping by leaning against the wind. Thereby they delay the adjustment of the exchange rate.

If adjustment is delayed, exchange rates will display a trend during the phase of adjustment. This trend may then be picked up and exploited by trend-following forecasters, who utilize trading rules of the type considered in the previous chapters. This story yields three empirically testable predictions:

1. Technical trading rules are only profitable during intervention periods.
2. Technical trading rules tend to bet against the central banks.
3. Central bank interventions are unprofitable.

There are three studies that address the connection between interventions and trading rule profits empirically. Szakmary and Mathur (1997) regress monthly MA trading rule returns on a constant and on changes in foreign exchange reserves (as the proxy for interventions), differentiating between leaning-against-the-wind- and leaning-with-the-wind interventions. They find that only the coefficient of leaning against the wind intervention is significantly positive and interpret this as saying that interventions explain mean profitability of MA trading rules.[1] Neely and Weller (1997) use genetic programming to search over a large space of possible technical trading rules (using both past prices as well as past intervention data as inputs) in order to identify successful rules in a selection period, before going on to examine the selected rules in an out-of-sample period. They find that "trading rules make most of their money from $t-1$ to t when intervention occurs in t." Moreover, "... intervention, when it occurs, is on the opposite side of the market from that taken by the trading rules."[2] Thus there is evidence in favour of prediction 2, and, to some extent, of prediction 1. In a third study, LeBaron (1996) uses daily exchange rate, interest rate and Federal Reserve Bank intervention data. He finds that Moving Average trading rules are extremely efficient at predicting the direction of exchange rate changes on days when the Fed intervened. Moreover, he finds that average returns are dramatically reduced if one restricts the sample to days on which no interventions were carried out. These results confirm prediction 1. On balance, therefore, the first two predictions appear to be confirmed by the data.

[1] Szakmary and Mathur (1997), Table 3 and p. 526ff.
[2] Neely and Weller (1997), p. 9.

Concerning the third prediction, however, there is some recent evidence by Leahy (1995) and Sweeney (1999) suggesting that rather than causing losses, central bank interventions yield substantial profits.[3] There thus appears to be a contradiction between technical trading rules being most reliable when trading against central banks and central bank interventions being profitable.

This chapter has four aims. The first is to examine whether this seeming contradiction is robust to considering not only interventions by the Federal Reserve Bank, but also those of the Bundesbank. Finding that this is the case, the second aim is to resolve this seeming contradiction. We find that the paradox is due to the fact that in the short run exchange rates (adjusted for interest differentials) move in the opposite direction of that in which central banks intend them to go, and that in the long run this effect reverses. This allows a conditional conclusion concerning the effectiveness of interventions: If they are effective at all, they work with a lag.

Whilst this explanation might resolve the seeming contradiction, there remains the question of why technical trading rules are so profitable on days when central banks intervene. It is the third aim of this chapter to rationalise this empirical regularity without putting it down to central banks transferring funds to technical traders. After showing some evidence that central bank interventions react to technical trading, we will present a stylised analysis of the conditions under which moving average trading rules are profitable. We will argue that if central banks lean against technical trends, then restricting the sample to intervention periods comes down to excluding unprofitable periods. The fourth and last aim of this chapter is to examine whether high technical trading rule profitability on intervention days can be exploited by alert traders. Whilst we find that simply betting against central banks is not a profitable strategy, it turns out that even if traders' information sets are restricted to those interventions which were reported in the newspaper, they can still use technical trading rules to make substantial profits when interventions take place.

This chapter is structured as follows: After discussing the intervention data in Section 4.2, we will demonstrate the robustness of the seeming contradiction between the profitability of

[3] Moreover, central banks generally claim that their foreign exchange interventions are profitable; see e.g. Gleske (1982).

technical analysis and central bank interventions in Sections 4.3 and 4.4. In Section 4.5 we will resolve the paradox, give an explanation for high trading rule profitability on intervention days and examine whether central bank interventions can be exploited. In Section 4.6 the results are discussed.

4.2 Data Summary

The intervention series consists of the daily amounts of USD (DEM) purchased by the Fed (Bundesbank) from the beginning of 1979 to the end of 1992.[4] Figure 4-1 shows the interventions of both central banks as well as the USD/DEM exchange rate. Whereas the Bundesbank intervened more or less continuously during the sample period, the Fed hardly intervened at all for substantial parts of the sample (most noticeably during the first Reagan administration in the early 80s).[5] The date of the Plaza agreement is easily recognisable in the Figure of Fed interventions as the biggest one day intervention in the whole period. The Bundesbank, in contrast, does not seem to have intervened as strongly in the aftermath of the Plaza. However, the figures suggest that its very heavy selling of USD in the beginning of 1985 might have started the decline of the dollar.

In order to see the overall dimension of intervention activities, it is also useful to look at cumulative interventions of Fed and Bundesbank. These are shown in Figure 4-2. Between 1979 and 1992 the Bundesbank was a net seller of USD on a scale exceeding 40bn, most of which occurred, it seems, to counteract the strong USD appreciation between 1981 and the beginning of 1985. The Fed appears to have been a more active player than the Bundesbank in the late 70s and early 80s. Since the reintroduction of interventions in 1985 it seems to intervene on a similar scale as the Bundesbank.

Table 4.1 contains some statistics of the intervention series. The first row gives the unconditional means of the interventions of Fed, of the Bundesbank and of the sum of Fed and Bundesbank intervention. Since the sum of the means of Fed and Bundesbank interventions is equal to the mean of the sum of interventions, we can conclude that Bundesbank and Fed

[4] Source: Federal Reserve Bank, Bundesbank.

[5] See Frankel (1990) and Dominguez and Frankel (1993) for a history of intervention activity during this time.

Figure 4-1: Federal Reserve and Bundesbank Interventions on the DEM/USD Exchange Market

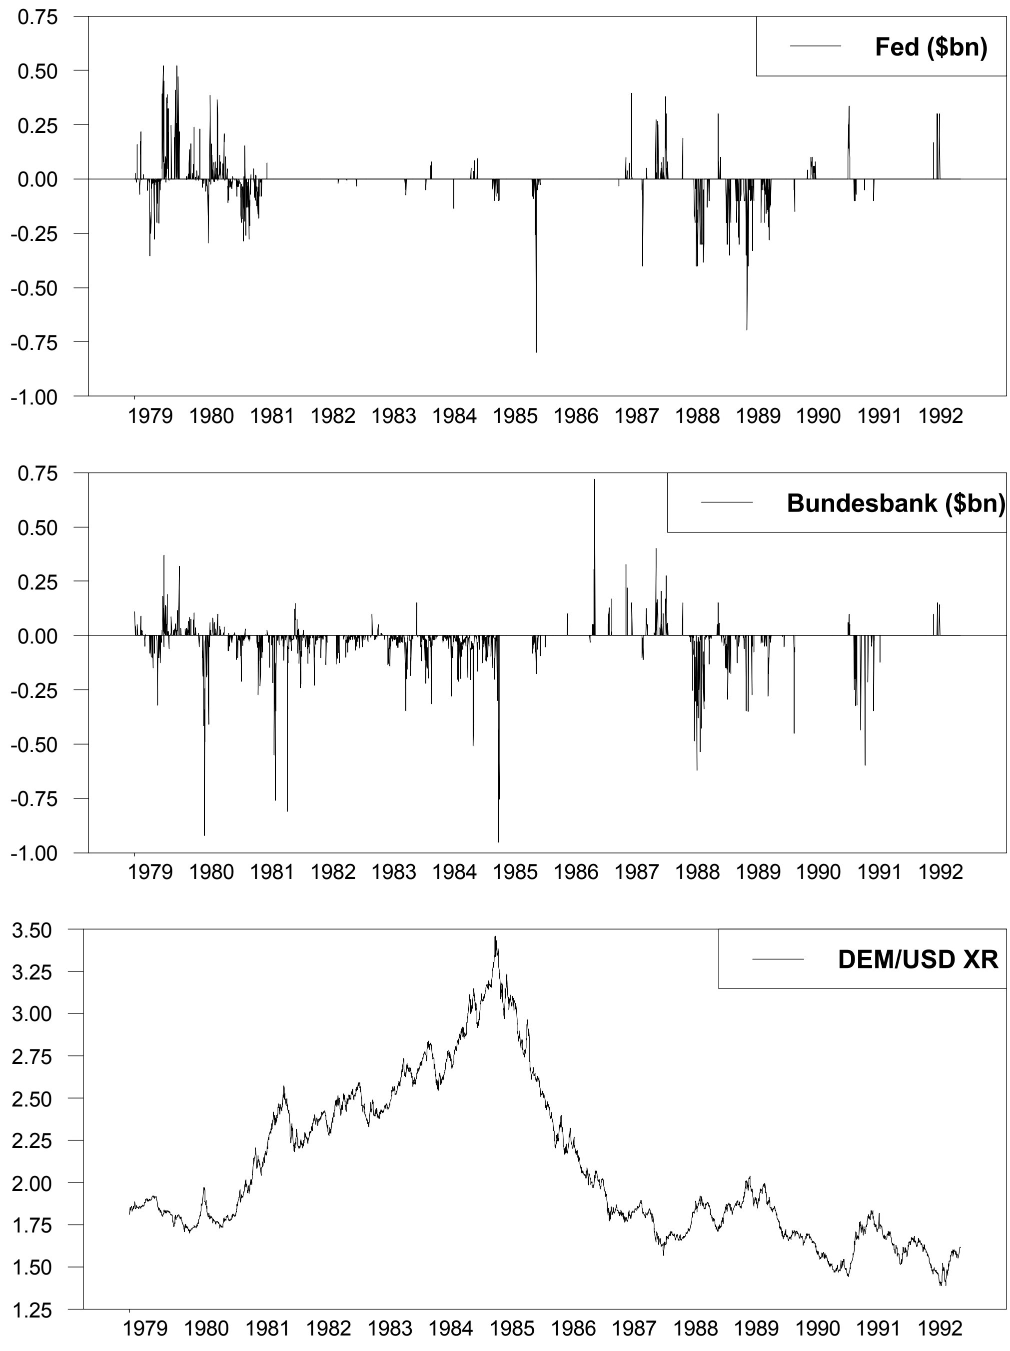

Figure 4-2: Cumulative Interventions of Fed and Bundesbank on the DEM/USD Exchange Market

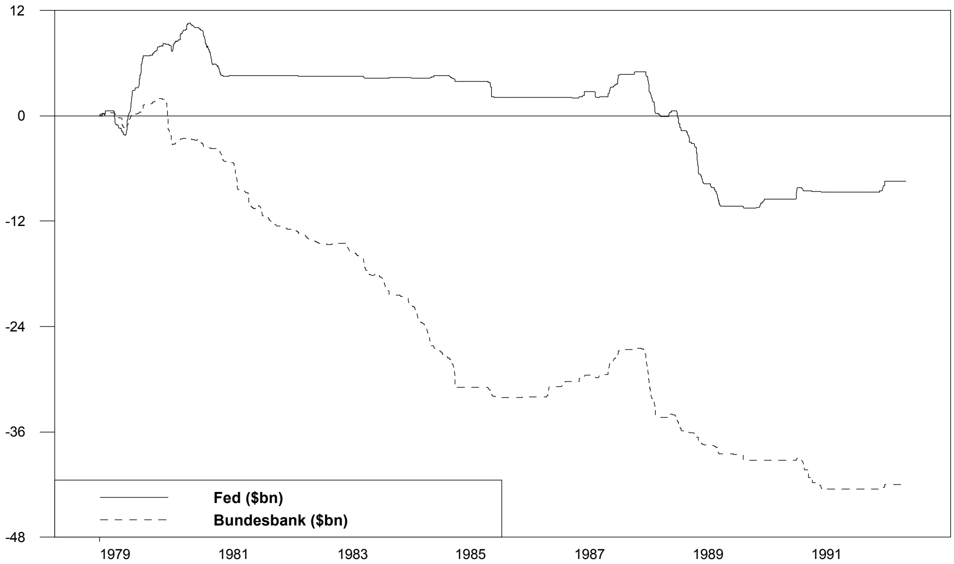

never intervened at cross purposes.[6] Moreover, the average size of interventions was roughly around USD 100m per day on which an intervention took place. It is worth noting that this compares to an estimated average daily turnover of about USD 1200bn in 1995.[7] One may wonder why central banks should be able to have any influence on the behaviour of exchange rates, given that daily interventions account for at most 1% of turnover. However, in this context it should not be forgotten that the dimension of daily turnover on foreign exchange markets is in itself a stylised fact in need of explanation. One theory that attempts to explain it is the 'Hot-Potato'-theory. 'Hot Potato' is a metaphor employed by foreign exchange dealers to describe the repeated passing of inventory imbalances from dealer to dealer following a new customer order.[8] For illustrative purposes, consider the following simple example which is due to Lyons (1995), p.2:[9]

[6] Strictly speaking, it is possible that interventions at cross purposes occurred but that they cancelled out. The likelihood of this happening can be considered as remote, however.

[7] BIS (1996), p.5.

[8] Lyons (1995).

[9] Lyons (1997) presents a formal model of hot potato trading.

Table 4.1: Intervention Summary Statistics

	Fed	Buba	Either
Mean (x_t)	-2.10	-11.86	-13.96
Mean $(\|x_t\|, x_t \neq 0)$	112.11	73.75	105.62
Fraction in Market	0.14	0.26	0.32
$P(x_t = 0 \| x_{t-1} = 0)$	0.95*	0.87*	0.85*
$P(x_t \neq 0 \| x_{t-1} \neq 0)$	0.58*	0.61*	0.68*

x_t is the amount of USD purchased at t; * indicates significance at 1%.

> Suppose there are 10 dealers, all of whom are risk averse, and each currently with zero net position. A customer sale of $10m worth of DM is accommodated by one of the traders. Not wanting to carry the open position, the dealer calculates his share of this inventory imbalance - or 1/10th of $10m - calls another dealer, and unloads $9m worth of DM. The dealer receiving this trade then calculates his share of this inventory imbalance - or 1/10th of $9m - calls another dealer, and unloads $8.1m worth of DM. The hot-potato process continues. In the limit, the total inter-dealer volume generated by $10m customer trade is $9m/(1-0.9)=$90m. The resulting share of wholesale trading that is inter-dealer: 90% roughly matches the empirical share.

Clearly, if there is some truth in the 'Hot-Potato'-theory, then the impact of an intervention on turnover is likely to be a multiple of the size of the original intervention. Consequently, the fact that intervention sizes are dwarfed by foreign exchange turnover no longer precludes interventions influencing exchange rates.

Table 4.1 also contains information concerning the frequency of interventions. On average the Fed intervened on one day in 7, the Bundesbank on one day in 4 and either central bank on one day in 3. Last not least, intervention- and non-intervention-periods tend to cluster, as reflected by Markov-switching probabilities significantly greater than one half.

4.3 Technical Trading on Intervention Days

Following LeBaron (1996), we use daily exchange rate and intervention data and compare technical trading rule returns on days when central banks intervened with returns when they did not. We extend LeBaron's results in two ways: Firstly, we consider not only two individual trading rules but all MA rules with moving averages between 2 and 500 as well as the peak-and-trough-progression rule. Secondly, we take into account not only Fed but also Bundesbank intervention. We split up the complete sample of trading rule returns into returns on days when interventions took place and returns on days when no interventions took place. Annual mean returns and standard errors are then evaluated for each subsample in turn.

Table 4.2 contains the results for the peak-and-trough-progression rule for the DEM/USD exchange rate. Whilst there is some evidence that trading rule returns are larger when central banks intervene (in particular for the Bundesbank), LeBaron's (1996) result that the profitability of technical trading rules disappears when intervention days are removed from the sample is not confirmed. Figures 4-3 and 4-4 show the corresponding results for moving average trading rules for the case when Bundesbank and Fed interventions are pooled.[10] For moving average rules there is stronger evidence that trading rules are unprofitable when central banks do not intervene. However, for short lengths of the moving average trading rule returns remain marginally significant. This means that only trading rules that are designed to pick up long term swings in exchange rates (i.e. moving average rules with high parameter values) are unprofitable in the absence of intervention. This is a first indication that there may be other factors responsible for the profitability of technical trading rules. Given the remarkably large trading rule returns on intervention days, however, our results are nevertheless compatible with the notion that central banks interventions contribute to technical trading profits, even if they are not the sole cause.

So far we have not considered the possibility that technical trading rules and central bank interventions tend to trade in the same direction. If this were the case, it would no longer be puzzling that trading rules are highly profitable on intervention days whilst central banks earn money with interventions, too. Table 4.3 contains the results of calculating the proportion of intervention days on which the sign of the trading signal of the peak-and-trough-progression rule did not agree with the sign of the central bank intervention. The proportions are about 75% no matter whether one looks at interventions by the Fed, by the Bundesbank or by either Fed or Bundesbank. Table 4.3 also contains asymptotic standard errors, which were derived assuming

[10] N.B. The corresponding figures for Fed and Bundesbank alone are contained in Appendix A.

Table 4.2: Annual Rates of Return of Peak-and-Trough-Progression Rule on Intervention Days

		Fed	Bundesbank	Either
Intervention	Mean	0.2331	0.2976	0.2319
	S. D.	0.1377	0.1261	0.1224
	S. E.	0.1016	0.0672	0.0585
	t-val.	2.2932	4.4315	3.9613
No Intervention	Mean	0.1009	0.0564	0.0653
	S. D.	0.1138	0.1138	0.1146
	S. E.	0.0331	0.0358	0.0377
	t-val.	3.0477	1.5744	1.7323

Figure 4-3: Annual Rates of Return of MA Trading Rules on Days when either Fed or Bundesbank Intervene

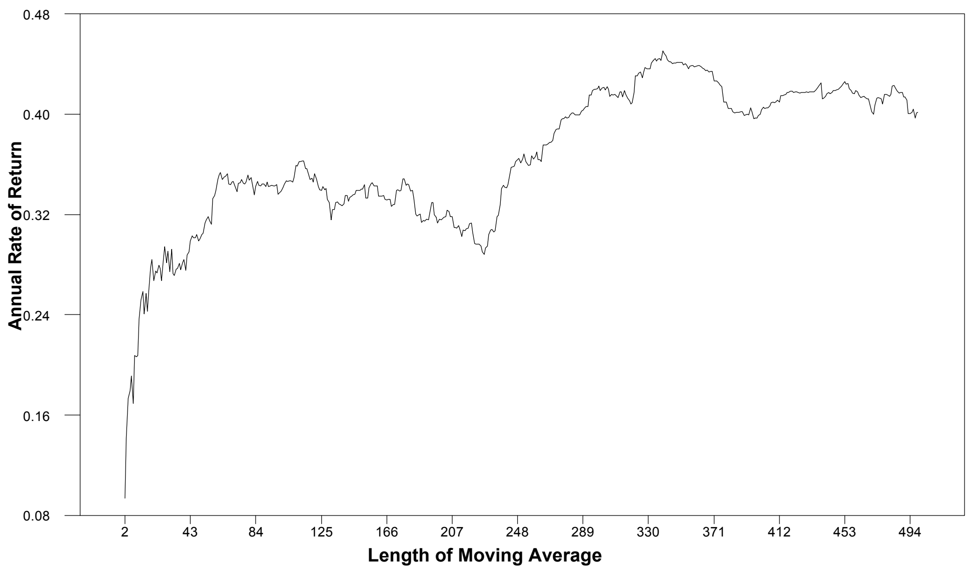

Figure 4-4: Annual Rates of Return of MA Trading Rules when neither Fed nor Bundesbank Intervene

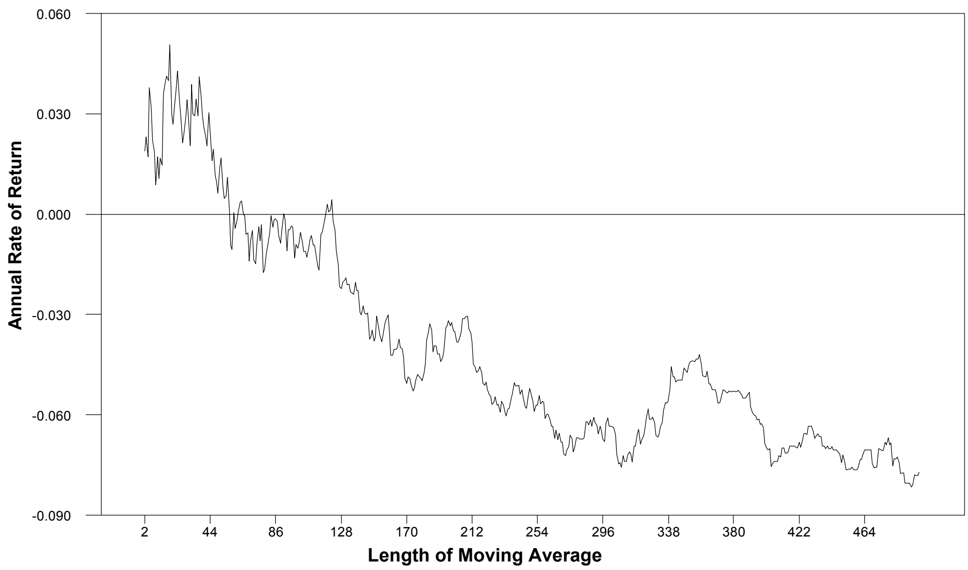

Table 4.3: Proportion of Trades of Peak-and-Trough-Progression Rules against Central Banks

	Fed	Bundesbank	Either
Proportion	0.7736	0.7506	0.7436
Std. Err.	0.0192	0.0143	0.0130

that the proportions are quasi-binomially distributed. Clearly, the observed proportions are significantly greater than 0.5.

In order to see whether this result is robust, we carried out analogous calculations for the moving average trading rule for the case when Bundesbank and Fed interventions are pooled.[11] Figure 4-5 shows the results. Proportions lie between 0.6 and almost 0.9. Moreover, they are significantly greater than 0.5 for all values of the moving average parameter.

It thus follows that on intervention days both the peak-and-trough-progression rule and the moving average rule yield very high returns, whilst trading against central banks in the large majority of cases. This strongly suggests that there exists a connection between technical trading and central bank interventions and suggests that central banks transfer funds to technical

[11] Again, the corresponding figures for Fed and Bundesbank alone are contained in Appendix A.

Figure 4-5: Proportion of Trades of MA Rules against either Fed or Bundesbank

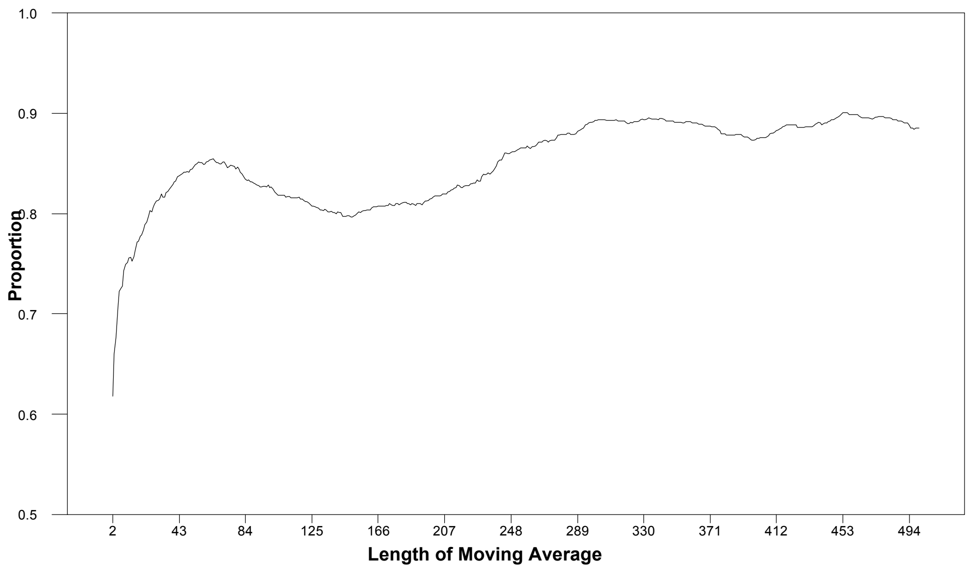

traders. If this were the case, however, we should expect the central bank to make losses with its interventions. Whether this is the case will be examined in the next section.

4.4 Profitability of Central Bank Interventions

Research concerning the question of whether central bank intervention in flexible exchange markets is profitable has for a long time been hampered by the lack of adequate intervention data.[12] Authors had to reconstruct intervention time series from changes in foreign exchange reserves, which were published (in the best case) at monthly frequencies. This has a number of drawbacks. An obvious problem is that intra-monthly changes in reserves are missed. Moreover, reserve changes do not correspond one to one with interventions because of transactions between central banks, forward market interventions and official borrowing of foreign currency.[13]

Concerning the results of calculating of the profitability of interventions using such data, there is a marked lack of consensus. Taylor (1982) found statistically and economically signifi-

[12] See Sweeney (1997) for a recent survey of the literature.
[13] Leahy (1995), p. 824.

cant intervention losses for nine central banks for the period from 1973 to 1979. In the case of the Fed, Taylor's (1982) results are challenged by Jacobson (1983) who finds that the Fed made large intervention profits when interest earnings are included and the sample period is extended until 1982. Moreover, Bank of England (1983) and Fase and Huijser (1989) find evidence for the profitability of interventions by the Bank of England and the Nederlandsche Bank, respectively.

So far only few studies have utilised daily intervention data to examine the profitability of interventions. The most influential paper in this literature is Leahy (1995), who uses intervention data for the Fed from 1973 to 1992. He finds that, the Fed earned \$12.3bn (\$4.2bn) intervening on the DEM/USD (JPY/USD) market.[14] Leahy employs a number of tests to examine whether interventions have any predictive power for subsequent deviations from uncovered interest parity. However, he does not carry out an outright test of the significance of the profits. More recently, Sjöö and Sweeney (1999) and Sweeney (1999) analysed the profitability of interventions by the Sverige Riksbank (1986-1990) and the Fed (1985-1991), respectively. Both use a variant of Sweeney's (1986) X-statistic to test the significance of risk-adjusted interventions returns. Whilst Sjöö and Sweeney (1999) find positive but insignificant intervention profits for Sweden, Sweeney (1999) finds highly significant risk-adjusted profits for the US.

Before carrying out the calculations of intervention profits for Bundesbank and Fed, it is worthwhile to say a word about the relationship between profitability and effectiveness of central bank interventions. In his famous paper on the advantages of having a flexible exchange rate system, Friedman (1953) argued that "destabilizing speculation is largely equivalent to saying that speculators lose money, since speculation can be destabilizing in general only if speculators on the average sell when the currency is low in price and buy when it is high."[15] A corollary of this argument is that a criterion for the success of central bank interventions in smoothing out temporary fluctuations in exchange rates is that central banks earn money with interventions.[16] However, since the time this paper was written, a number of authors have produced counter-examples showing that it is possible for destabilising speculation to be profitable[17] and that speculation can be profitable without affecting the exchange rate.[18] For this reason, one cannot

[14] Leahy (1995), Table 1.
[15] Friedman (1953), p. 175. Note also, however, the qualifications he makes in footnote 8 on the same page.
[16] Friedman (1953), p.188.
[17] E.g., Baumol (1957), Hart and Kreps (1986) and Szpiro (1994).
[18] E.g., Leahy (1995), Dominguez and Frankel (1993).

draw any immediate conclusions concerning the effectiveness of interventions from finding them either profitable or not.

As a starting point, we adopt Leahy's approach to measuring intervention profits with the only major deviation that we use 1-week eurorates rather than 3-month treasury bill rates.[19] Let x_t be the amount of USD bought on day t; as before S_t is the exchange rate and i_t (i_t^*) are the USD (DEM) 1-week eurorates. Let T be the last date in the sample (i.e. the 31.12.1992). The profitability of interventions is analysed by creating a zero-cost portfolio which mirrors intervention activity. It is assumed that whenever the Fed/Bundesbank buys USD, it borrows the necessary funds in the Eurodeutschmark market and invests the purchased USD in the Eurodollar market. When USD are sold the reverse transactions are carried out. It is further assumed that each such intervention position is maintained until the last day of the sample. Under these assumptions, the contribution of the intervention at time t to profits from intervention measured as of date T can be approximated as follows:

$$X_{t,T} = x_t \left[\Pi_{j=t}^{T-1}(1+i_j) - \frac{S_t}{S_T} \Pi_{j=t}^{T-1}(1+i_j^*) \right] \tag{4.1}$$

The first part of the square bracket is the return on the x_t USD purchased and invested every day anew in the Eurodollar market (assuming for the purpose of illustration that x_t is positive). The second part of the square bracket is the cost of borrowing the necessary funds for the intervention:[20] We borrow $x_t S_t$ DEM at t and have to pay it back with interest at the exchange rate prevailing at time T. Total profits from interventions measured as of date T are thus the sum of the contributions of interventions before T:[21]

$$\Pi_T = \sum_{t=1}^{T-1} X_{t,T} = \sum_{t=1}^{T-1} x_t \left[\Pi_{j=t}^{T-1}(1+i_j) - \frac{S_t}{S_T} \Pi_{j=t}^{T-1}(1+i_j^*) \right] \tag{4.2}$$

Table 4.4 contains intervention profits evaluated as of the last day in the sample (i.e. 31.12.1992). According to this measure of intervention profits, the Fed earned around \$12.5bn

[19] As it turns out, it makes practically no difference which of these two interest rates one uses. Compare Figure 4-6 below with Leahy (1995), Figure 1.

[20] Another way of looking at it is as the opportunity cost of the investment (Leahy 1995, p.826).

[21] Note that the term in square brackets can be interpreted as the ex-post deviation from uncovered interest parity between time t and T.

Table 4.4: Profitability of Central Bank Interventions

	Total Profits	Profits for:		Bootstrap p-value
	31.12.1992	USD+20%	USD-20%	Random Walk
Fed	12572.7	10375.7	15256.3	0.0090
Bundesbank	26636.3	8573.5	53141.8	0.0975

whilst the Bundesbank even earned \$28.6bn. Since our method of calculating intervention profits assumes that each intervention position is held until the last day in the sample, the estimates of intervention profits are very sensitive to changes in the value of the exchange rate in T. In order to check whether our results are robust against changes in S_T, we also calculated intervention profits for a 20% higher and 20% lower value of S_T. Whilst profits remain positive, in particular Bundesbank profits react very strongly to changes in the exchange rate. This is not surprising given that the Bundesbank was a very large net seller of USD during the sample period (see Figure 4-2).

It suggests itself to examine whether these profits could have come about by chance. There is, however, considerable disagreement concerning the question of how to ascertain the significance of intervention profits.[22] The basic problem one faces when testing for significance can best be illustrated if one neglects interest rates. In this case (4.2) can be written as:

$$\begin{aligned} \Pi_T &= \sum_{t=1}^{T-1} x_t \left[1 - \frac{S_t}{S_T}\right] \\ &= \sum_{t=1}^{T-1} x_t \left[\frac{S_T - S_t}{S_T}\right] \end{aligned}$$

If we assume that log exchange rates follow a random walk without drift, such that

$$\ln S_t = \ln S_{t-1} + \varepsilon_t,$$

interventions profits can be approximated as follows:

[22] See, for instance, the debate between Taylor (1982, 1989) and Spencer (1985, 1989) in the Journal of Political Economy.

$$\Pi_T \cong \sum_{t=1}^{T-1} x_t \sum_{i=t}^{T-1} \varepsilon_i \tag{4.3}$$

From (4.3) it is apparent, that Π_T contains a variable which is integrated of order 1, which implies that standard significance tests do not apply. Further difficulties are created by the fact that interventions are not independent of past exchange rate changes. Thus, even if interest differentials are ignored and exchange rates are assumed to follow a random walk, it is by no means obvious how one should go about testing the significance of the observed intervention profits.

Corrado and Taylor's (1986) approach is to derive the distribution of intervention profits under the assumption that interventions are negatively linearly related to exchange rate changes, motivating this assumption with the leaning-against-the-wind nature of central bank interventions. If exchange rates follow a random walk and interventions are a linear function of exchange rate changes, then this implies that interventions are independently identically distributed. One problem with this approach is, however, that we know from the Markov switching probabilities reported in Table 4.1 that both intervention and non-intervention days tend to cluster, which is incompatible with the Corrado and Taylor's (1986) assumptions. The same problem also affects the approach chosen by Sweeney (1999), who uses Monte Carlo simulations to derive the critical values for standardised intervention profits under the assumption that exchange rates and interventions are correlated but independently and identically normally distributed.[23] Another problem with Sweeney's (1999) approach is that it seems highly questionable whether the normal distribution is a good approximation for interventions.

There are basically two ways to avoid these problems. One way would be to estimate first an intervention reaction function, which could then be used in simulations to derive the distribution of intervention profits. The problem with this approach is that intervention data are censored. This implies that reaction functions need to be estimated as limited dependent variable models (usually Probit or Logit), which are of little use for simulating actual intervention amounts.

Alternatively, one could take the intervention time series as given and carry out bootstrapping simulations for exchange rates and interest rates with the aim of finding out, how likely

[23] Sweeney (1999), Appendix A.

it is to find as large or larger intervention profits given some null model for exchange rates.[24] Adopting this approach, we carried out 2000 bootstrap simulations for a random walk null model. The results are also contained in Table 4.4. Whilst we found as large or larger intervention profits in less than 1% of the simulations for the Fed, this proportion is almost 10% for the Bundesbank. This again reflects the fact that the Bundesbank was a large net seller of Dollars over the sample period. Whenever a simulated USD series depreciates greatly over the sample period, we are bound to find very large intervention profits. This points to a fundamental problem in this approach to measuring intervention profits.

Our data set indicates that the Bundesbank sold more than $40bn during our sample period. Yet from the Bundesbank's annual reports it emerges that the stock of dollar assets of the Bundesbank changed from DEM 70.57bn in January 1979 to DEM 66.94bn in December 1992.[25] Thus the Bundesbank must have acquired reserves by some other way than the foreign exchange market.[26] Apart from some minor sources of foreign exchange reserves like forward market transactions and official lending of reserves, the most important source is likely to be transactions with the Fed. The true profits the Bundesbank made from intervening thus depend on when such transactions took place and at which rates they were carried out. However, this consideration should only influence the way profits are split up between the Fed and the Bundesbank. Given the size of the intervention profits we found for both central banks, it seems unlikely that either should have made losses.

In Figure 4-6, intervention profits are broken up according to when they originated. The most noteworthy aspect of the figure is that it shows that almost all interventions have contributed positively to intervention profits. Since this seems to stand in clear contrast to our previous results confirming the profitability of taking positions contrary to central banks, further examination of the relationship between trading rule profitability and central bank interventions is warranted.

[24] N.B. The null hypothesis that is being tested is a joint one of zero intervention profits and interventions not influencing exchange rate changes. Without the latter assumption, it would be necessary to specify, how interventions affect exchange rates, which would complicate our task considerably.

[25] Bundesbank (1998), EU8105: Auslandsposition der Bundesbank - Bestand - Devisen und Sorten: US-Dollar Anlagen.

[26] N.B. Different valuations of the USD reserves do not play an important role here.

Figure 4-6: Cumulative Contributions to Intervention Profits

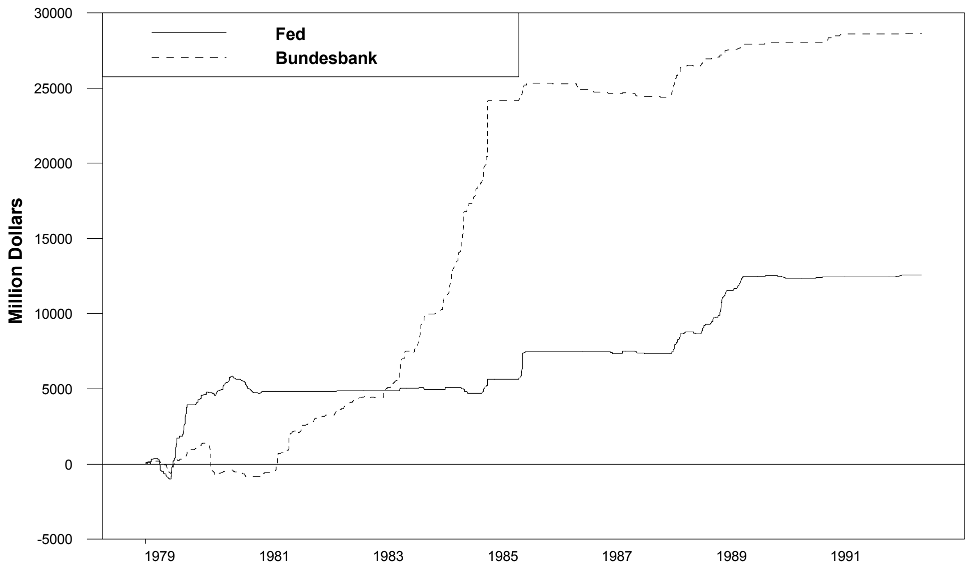

4.5 Squaring the Facts

4.5.1 Short Term versus Long Term Effectiveness of Central Bank Interventions

The key to understanding why technical analysts can earn money by betting against central banks whilst central banks make profits with their interventions, lies already in Figure 4-6. There we examined where profits measured as of 31.12.1992 originated from. Thus we evaluated the contribution of each intervention taking into account the development of the exchange rate and of interest rates from the date of the intervention until the last day in the sample. This means that for all but the last, say 260 observations we looked at the long term profitability of these interventions. Thus, one possible way of resolving our seeming contradiction is that, after interventions, exchange rates (net of interest differentials) move in the opposite direction of central banks' intentions in the short run, but in the long run they move in accordance with interventions.

In order to check this hypothesis, we analyse intervention profits assuming that each intervention position is not held until the end of our sample, but rather held for K days. We assume, for instance, that if the Bundesbank sells USD at t, it will repurchase them directly from the Fed (rather than via the foreign exchange market) at $t+K$. Formally, intervention profits are thus measured as follows:

$$\Pi^*(K) = \sum_{t=1}^{T-K} x_t \left[\Pi_{j=t}^{t+K-1}(1+i_j) - \frac{S_t}{S_{t+K}} \Pi_{j=t}^{t+K-1}(1+i_j^*) \right] \tag{4.4}$$

If our hypothesis is correct, we should expect that $\Pi^*(K)$ is negative for small values of K and that it changes sign as K increases. Figure 4-7 shows intervention profits for K ranging from 1 to 520 days. Clearly, the hypothesis is confirmed by the data. There exists, however, a marked difference between Fed and Bundesbank interventions concerning the time lag before exchange rates (net of interest differentials) start to move in the direction in which central banks intend them to go. Whilst this lag equals 19 days (on average) for the Fed, it is greater than one year for the Bundesbank. An explanation for this might be that the results for the Bundesbank are dominated by the Bundesbank's efforts to stop the massive appreciation of the USD between 1981 and 1984, which was followed by a reversal of the trend only in February 1985. In order to check whether this is the case, we carried out the same calculations for the subsample 1986 to 1992. As Figure 4-8 shows, the difference in lags is now much smaller, equalling 4 days for the Fed and 40 days for the Bundesbank. Nevertheless, a substantial difference remains.

It is worth noting that the term in square brackets in equation (4.4) can be interpreted as ex-post deviations from uncovered interest parity (UIP). Thus, the fact that $\Pi^*(K)$ becomes large and positive suggests that central bank interventions are good predictors of long term deviations from UIP. In order to examine the relationship between interventions and UIP more directly, we regress ex-post deviations from UIP on (lagged) interventions.[27] Thus we estimate:

$$\left[\Pi_{j=t}^{t+K-1}(1+i_j) - \frac{S_t}{S_{t+K}} \Pi_{j=t}^{t+K-1}(1+i_j^*) \right] = \alpha + \beta x_{t-1} + u_t \tag{4.5}$$

$K > 1$ induces u_t to have an "overlapping observation" structure. In order to take account of this and because of heteroskedasticity in deviations from UIP, we use Newey West (1987) autocorrelation and heteroskedasticity consistent covariance estimators. Figures 4-9 and 4-10 show parameter estimates of α and β, as well as Newey-West standard errors when interventions

[27] N.B. Interventions need to be lagged to avoid a simultaneity bias.

Figure 4-7: Intervention Profits Measured for Varying Horizons (Full Sample)

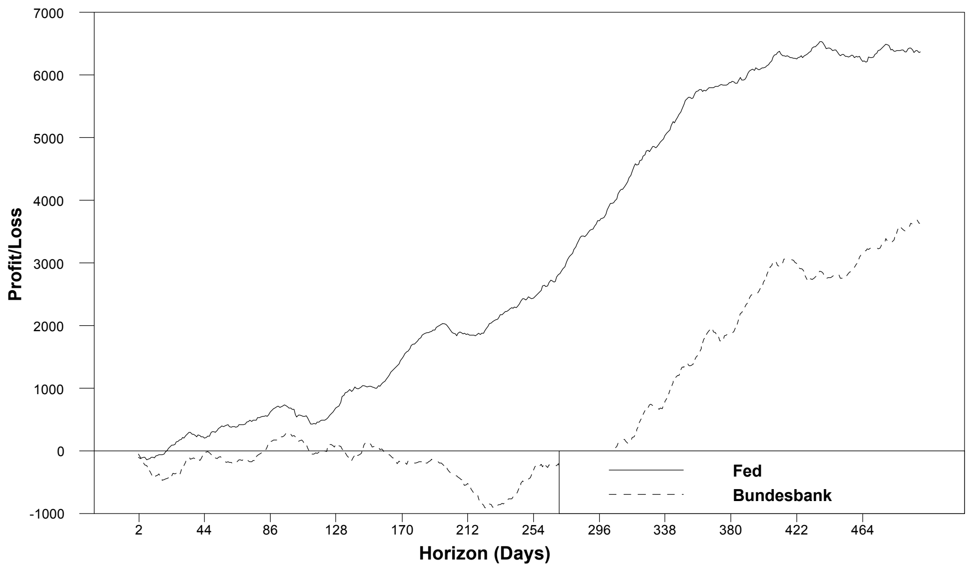

Figure 4-8: Intervention Profits Measured for Varying Horizons (1986 onwards)

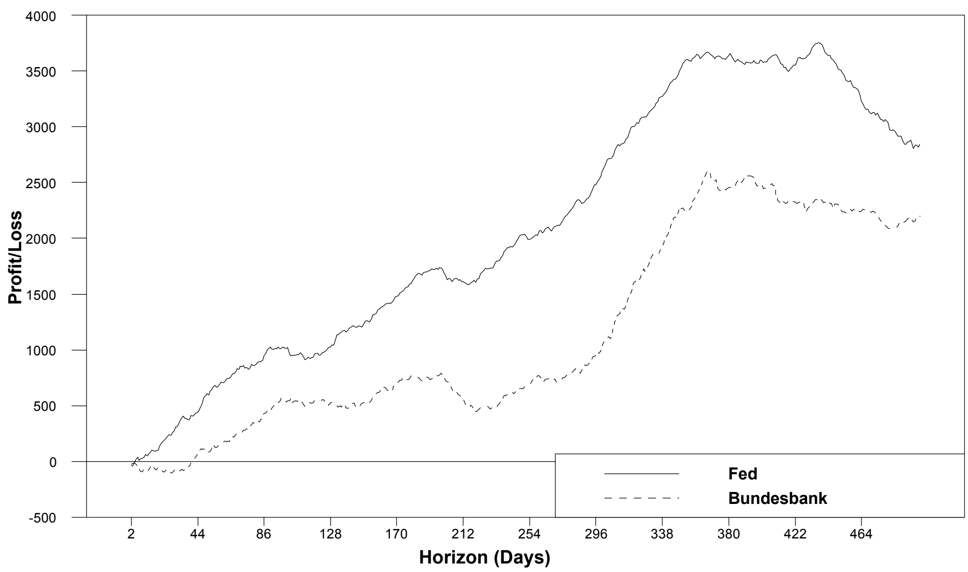

Figure 4-9: Regressing Interventions on Future Deviations from UIP: Estimates of Alpha

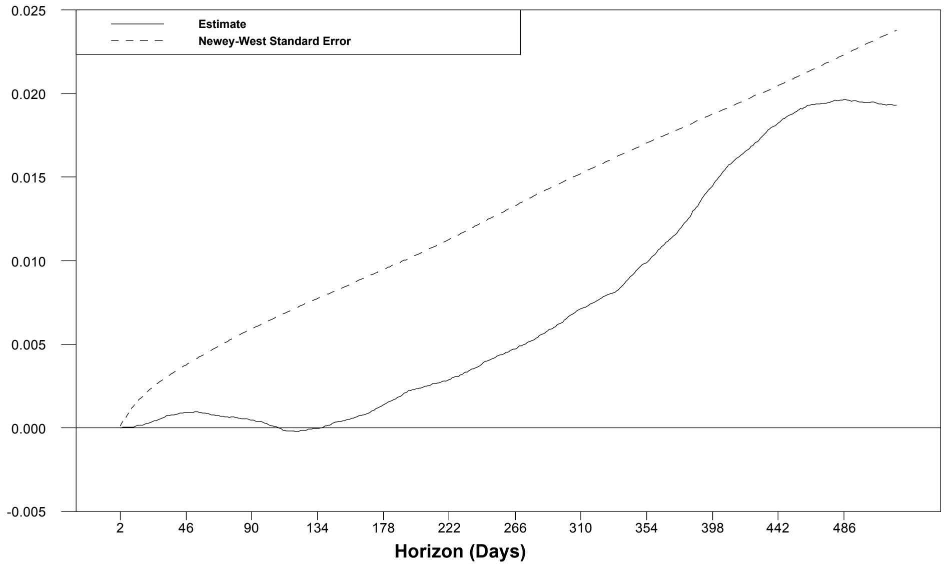

of Fed and Bundesbank are pooled.[28] As before, K ranges from 1 to 520 days.

Parameter estimates for α (which might also be interpreted as estimates of a time invariant risk premium) are never significantly different from zero. The estimates of β are negative (but not significant) for horizons shorter than 27 days, then become positive and end up being significantly greater than zero for horizons greater than 320 days. This confirms statistically that central bank interventions are good predictors of long term deviations from UIP. Moreover, it shows that interventions are profitable because they exploit systematic deviations from UIP.

Our results are in close agreement with a finding of Goodhart and Hesse (1993), who used tick-by-tick exchange rate data and Reuters reports on interventions by the Fed, Bundesbank and the Bank of Japan for a 12 week period in 1989. They found that:[29]

> One-hour gains [of interventions by the Bank of Japan] are pretty randomly distributed, with an overall average loss of 0.07 JPY per USD sold. Again, over 2, 3, 6, 12 and even 24 hours the pattern of gains is unclear, each horizon involving an

[28] The corresponding results when Fed and Bundesbank interventions are not pooled are given in Appendix A.
[29] Goodhart and Hesse (1993), p.384. See in particular Figures 7a-h.

Figure 4-10: Regressing Interventions on Future Deviations from UIP: Estimates of Beta

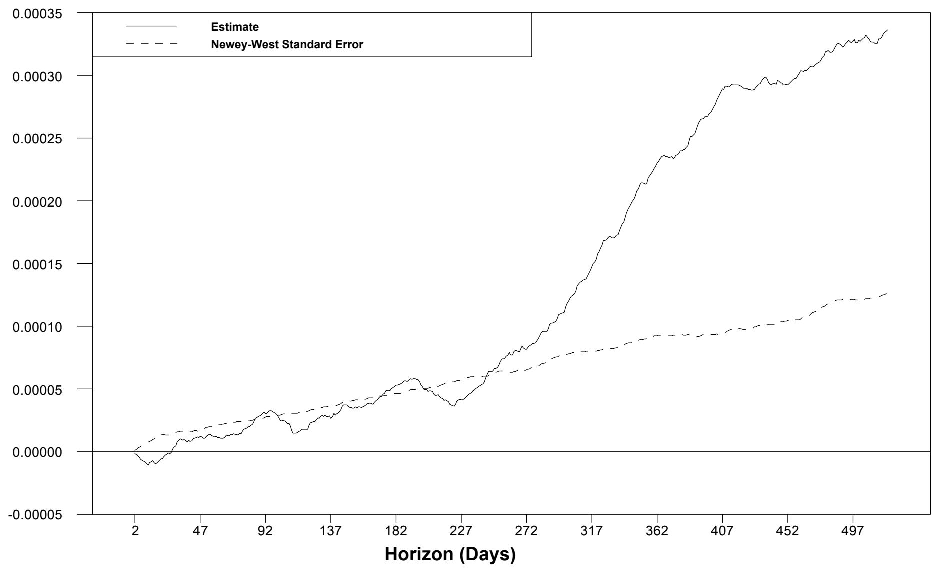

> average loss. But on a one-week horizon, we get a clear, time-dependent success record, with losses in interventions building up until and including those of June 9th, but then turning into gains. ... Looking at Fed interventions produces similar insights: no pattern recognizable on the short horizons (up to 24 hours), but a clear time-dependent pattern on the longer horizon. By the end of the year, the USD had pushed down sufficiently for all Fed interventions in DEM to be profitable, ... Gains patterns for the Bundesbank are similar to those for the BoJ and the Fed.

Therefore, the seeming contradiction between the fact that technical trading rules earn very large returns when trading against central banks and central banks nevertheless making interventions profits is due to two considerations:

- The profitability of technical trading rules and of central bank interventions was calculated over different horizons.
- Exchange rates (net of interest differentials) move in a manner consistent with central banks' intentions only in the long run. In the short run the opposite tends to be the case.

This answer immediately raises the question of why exchange rates behave in this way. Before suggesting an answer, however, it is useful to try to explain first why technical trading rules are so remarkably profitable on days when central bank intervene. So far we can only draw a conditional conclusion concerning the question of the effectiveness of central bank interventions: If they are effective at all, they work with a lag.

4.5.2 Explaining High Trading Rule Profitability on Intervention Days

For the purpose of understanding why technical trading rules are so profitable on days on which central banks intervene, it is necessary to examine the motives that lie behind central bank interventions. A number of authors have estimated central bank intervention reaction functions. Dominguez and Frankel (1993) regress interventions on deviations of spot exchange rates from Purchasing Power Parity (PPP) values and find evidence that central banks lean against deviations from PPP.[30] Almekinders (1995) estimates a Tobit model for interventions, including deviations of the spot exchange rate from its 3, 5 and 7 day moving average as an explanatory variable.[31] He finds strong evidence that deviations from moving averages prompt interventions. Almekinders (1995) also finds evidence that exchange rate volatility causes interventions.[32] However, both Bonser-Neal and Tanner (1996) and Dominguez (1998), find evidence that rather than reacting to volatility, central bank interventions increase volatility. Baillie and Osterberg (1997), finally, estimate a probit model, including deviations from target exchange rate levels identified by Funabashi (1988) as regressors. They find that deviations from target levels tend to increase the probability of intervention. Concerning the question whether exchange rate volatility causes intervention, however, Baillie and Osterberg's (1997) results are inconclusive.[33]

The empirical evidence therefore suggests that it is fair to describe central bank interventions as leaning against deviations from target levels. A slightly different way of interpreting the evidence is to say that central banks intervene to counteract technically motivated trends (as opposed to trends caused by diverging fundamentals like, for instance, persistent inflation

[30] Dominguez and Frankel (1993), p. 78ff.
[31] Almekinders (1995), p. 109ff.
[32] Almekinders (1995), p.103ff.
[33] N.B. Recent evidence using intraday exchange rate data suggests that interventions have a positive effect on volatility. See Chang and Taylor (1998).

differentials). Whilst it is difficult to distinguish between these two interpretations empirically, in particular Almekinders' (1995) result that interventions react to deviations from moving averages fits in nicely with this view of the intentions of central banks: Technical traders using moving average rules start buying a currency if its exchange rate moves above its moving average. The further it moves away, the clearer are the indications (for technical traders) that an upward trend has commenced and thus the more technical traders will enter the market. Thus the fact that central banks tend to intervene the more the spot exchange rate deviates from its moving average means that interventions are more likely to take place / are likely to be more intensive when there is a preponderance of technical trading. This is very suggestive of the fact that interventions are aimed at counteracting technical trading.

In order to make sure that this interpretation rests on a sound empirical basis, we estimate a Probit model for interventions, using deviations of the exchange rate from its one week (5 day), from its one month (21 day) and from its three month (65 day) moving average as explanatory variables. Lagged interventions are also included to serve as a proxy for other non-observable factors that influence the intervention decisions. We follow Baillie and Osterberg (1997) in estimating the model separately for USD purchases and USD sales.[34] If we define x_t^+ as interventions involving USD purchases, and x_t^- as interventions involving USD sales, the models to be estimated can be written as follows:[35]

$$\ln\left(\frac{\Pr(x_t^+\neq 0)}{1-\Pr(x_t^+\neq 0)}\right) = \alpha_1 - \beta_1(s_{t-1} - \tfrac{1}{N}\Sigma_{i=1}^N s_{t-i}) + \gamma_1 x_{t-1}^+ + u_{1t}$$
$$\ln\left(\frac{\Pr(x_t^-\neq 0)}{1-\Pr(x_t^-\neq 0)}\right) = \alpha_2 + \beta_2(s_{t-1} - \tfrac{1}{N}\Sigma_{i=1}^N s_{t-i}) - \gamma_2 x_{t-1}^- + u_{2t}$$

Table 4.5 shows the results of estimating the above equations. For both the Fed and the Bundesbank there is strong evidence that deviations of the spot exchange rate from its moving average prompt interventions in the expected direction. Moreover, the result appears to be stable across various lengths of the moving average.[36] Given these results, it seems reasonable

[34] The reason for doing this is that both positive and negative deviations from target exchange rate levels might increase the probability of intervention, which implies that the relationship between the probability of intervention and deviations from target exchange rates will not be monotonic, let alone linear.

[35] N.B. The signs in front of the parameters are chosen in a way that a priori their estimates should be positive.

[36] It is worth noting that our results contrast with those of Dominguez (1998), who estimates a very similar probit model for the period 1985 to 1994. Dominguez (1998) uses deviations of both the spot exchange rate and of volatility from their respective moving averages as regressors. The crucial difference is, however, that she does not estimate the model separately for USD purchases and sales. As a consequence, it is not particularly

Table 4.5: Estimates for Probit Model of Central Bank Interventions

		x_t^+		x_t^-	
		Estimate	t-val.	Estimate	t-val.
Fed					
$N=5$	α_1, α_2	-1.8509	-44.3403	-1.5666	-45.1543
	β_1, β_2	6.5533	1.2678	12.5727	3.0247
	γ_1, γ_2	0.0104	13.8824	0.0109	-15.6000
$N=21$					
	α_1, α_2	-1.8697	-43.4370	-1.5897	-44.2250
	β_1, β_2	7.2288	3.4124	11.6407	6.7479
	γ_1, γ_2	0.0102	13.7041	0.0104	-14.8497
$N=65$					
	α_1, α_2	-1.8801	-42.8487	-1.6311	-42.6582
	β_1, β_2	4.0813	3.5948	8.9013	9.0228
	γ_1, γ_2	0.0103	13.6217	0.0098	13.9516
Bundesbank					
$N=5$	α_1, α_2	-1.8017	-43.1159	-0.9988	-37.2173
	β_1, β_2	33.4822	6.7695	22.2355	6.8914
	γ_1, γ_2	0.0036	7.9852	0.0044	16.7349
$N=21$					
	α_1, α_2	-1.8582	-40.9549	-1.0068	-36.5941
	β_1, β_2	18.4785	8.5818	16.6929	11.9494
	γ_1, γ_2	0.0033	7.3485	0.0039	14.8464
$N=65$					
	α_1, α_2	-1.8801	-39.9278	-1.0145	-35.7800
	β_1, β_2	9.5392	7.8463	11.5257	15.0449
	γ_1, γ_2	0.0035	7.7509	0.0036	13.6175

to start off from the assumption that central bank interventions are intended to counteract technically motivated trends in foreign exchange rates.

In order to see how this assumption can help to explain why technical trading rules are so profitable on intervention days, it is useful first to discuss some conditions under which trading rules are profitable. Figure 4-11 presents a highly stylised analysis of the profitability of moving average trading rules. The exchange rate moves in cyclical swings around a long term trend in the exchange rate. Whilst this long term trend is reflected in interest differentials, the short term trends carrying the exchange rates above or below the fundamental value are not. Also shown is a moving average of the exchange rate. Phases during which the moving average

surprising that she does not find significant parameters for the coefficient of deviations of the spot exchange rate from its moving average.

Figure 4-11: Stylised Analysis of the Profitability of Moving Average Trading Rules

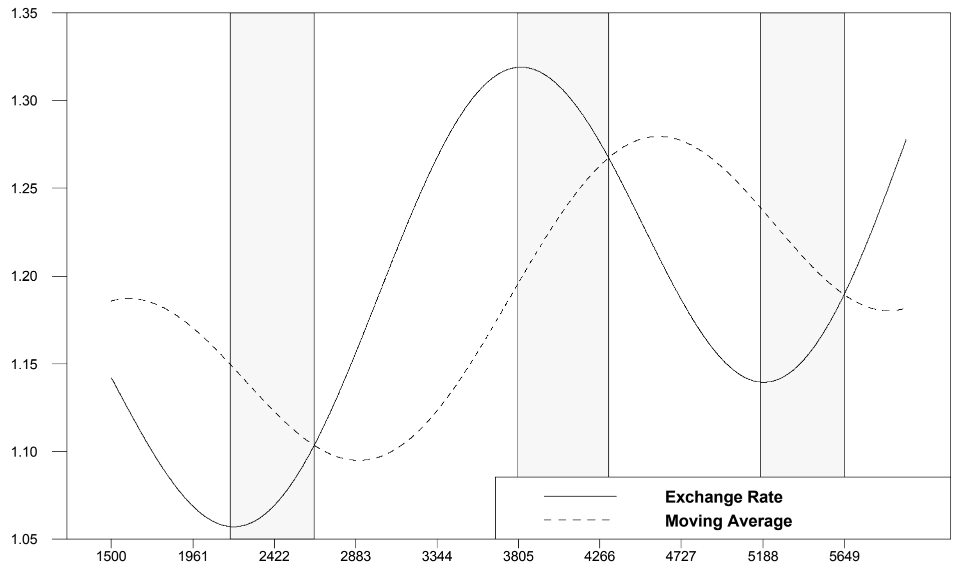

trading rule would have been unprofitable are shaded: They start whenever a (short-term) trend turns and last until the trading rule recognises that the trend has changed.[37]

Figure 4-12 presents a stylised analysis of the timing of central bank interventions. Interventions start after exchange rates have moved sufficiently far away from fundamentals for central banks to recognise an over-/undervaluation and decide to try to counteract it. They last until the trend has turned and the exchange rate moves back towards its fundamental value. In Figure 4-12 periods of central bank interventions are shaded. Comparing Figures 4-11 and 4-12 indicates that periods during which central banks intervene do not coincide with those during which moving average trading rules are unprofitable. The reason for this is that if central banks only intervene to counteract trends away from fundamentals, then they will only start intervening when a trend is already in full swing, and will stop as soon as the trend turns. Thus, technical trading rules are so remarkably profitable on intervention days, because restricting the sample to days on which interventions took place is to a great extent equivalent to excluding periods during which technical trading rules are unprofitable. Put in another way,

[37]To be precise, because of the interest differential the unprofitable periods start slightly earlier at peaks and slightly later at troughs. This is neglected for the sake of clarity.

Figure 4-12: Stylised Analysis of the Timing of Interventions

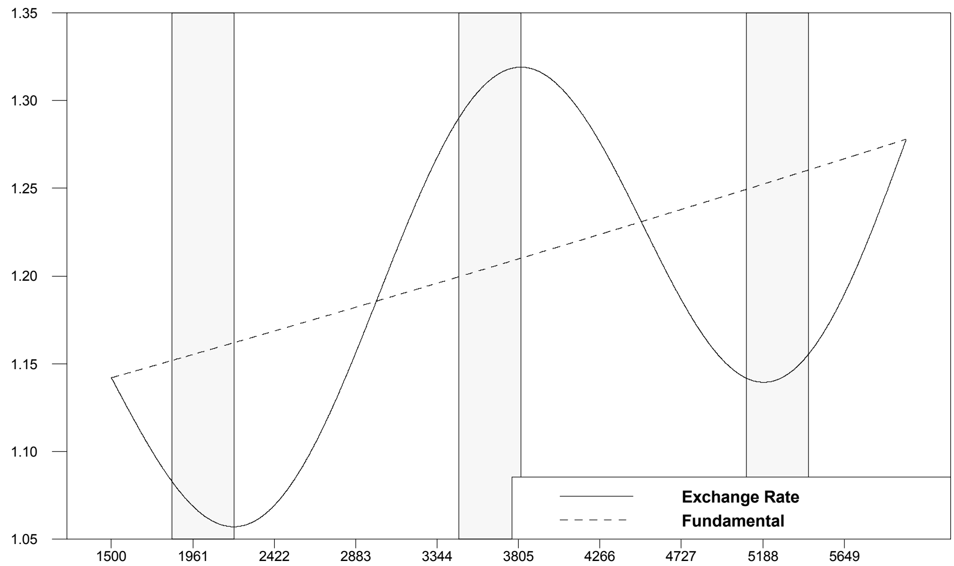

higher trading rule profitability on intervention days could be due to the fact that interventions are a more reliable signal for the existence of a trend than trading rule signals by themselves.

It might be objected that our explanation of high trading rule profitability on intervention days rests on a rather simple analysis of intervention timing and trading rule profitability. Simple though it is, however, it does yield an empirically testable prediction: If central banks stop to intervene as soon as the trend away from the fundamental value has stopped, Figure 4-11 suggests that in the first days after an intervention period technical trading rules should yield losses. The problem with testing this prediction is that our data set does not allow us to identify when precisely intervention periods started and when they ended. This is illustrated in Figure 4-13, which shows the DEM/USD exchange rate between 1985 and 1990 with intervention days shaded. Instead of defining some ultimately arbitrary criterion, according to which the starting and ending dates of intervention periods are to be identified, our approach is to make use of the set of intervention periods which were identified by Catte et al. (1994).[38] These periods are shown in Figure 4-14.

[38] Catte et al.'s (1994) used (confidential) intervention data from 16 central banks when identifying these periods.

Figure 4-13: USD/DEM Exchange Rate and Fed+Bundesbank Interventions

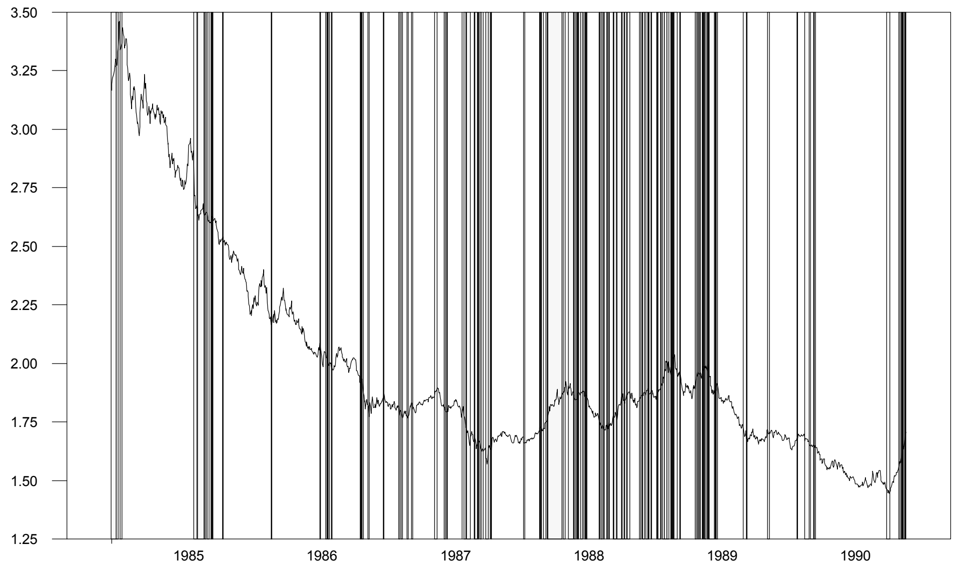

Figure 4-14: Intervention Periods as Identified by Catte et al. (1994)

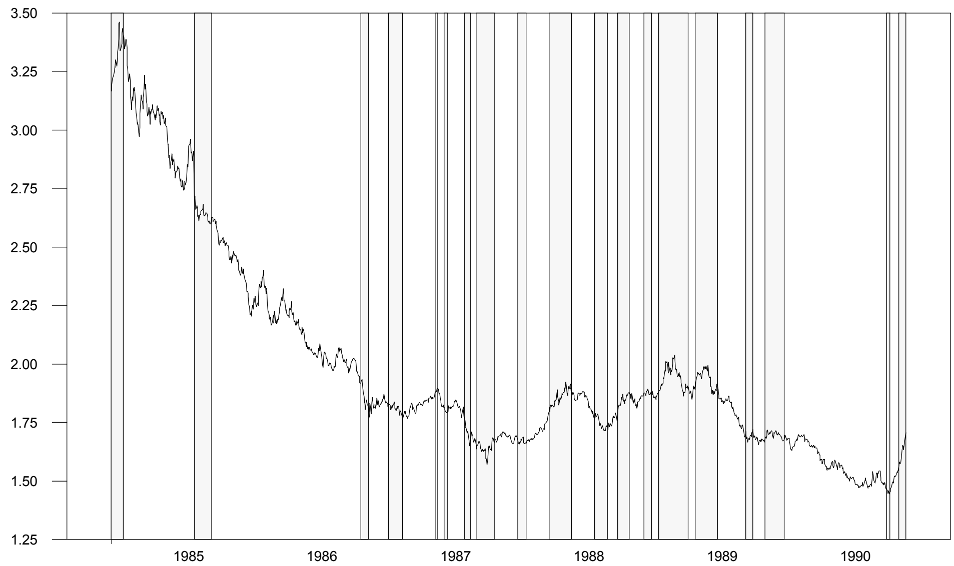

Table 4.6: Peak-and-Trough-Progression Rule Returns during and after Intervention Periods

	Post-Intervention	Other	Intervention	Full Sample
5 Days	-0.4804	0.1271		
10 Days	-0.2407	0.1403	0.1466	0.0959
15 Days	-0.0650	0.1112		

Using these intervention periods, we now split up the sample into three subsamples: Firstly intervention periods, secondly post-intervention periods, which are defined as the first 5, 10 or 15 days after the end of an intervention period and thirdly periods which neither coincide with nor are preceded by interventions. Table 4.6 shows annual rates of return of the peak-and-trough-progression rule for each of these three subsamples and for each definition of post intervention periods.

Clearly, the prediction that trading rule returns should be negative during the first days after interventions is born out in the data. Moreover, our use of intervention periods as identified by Catte et al.'s (1994) only marginally changed annual returns during interventions periods. In order to make sure that these results are not only valid for the peak-and-trough-progression rule, we carried out the same calculations for the moving average rule. Figure 4-15 shows annual rates of return for each of the three subsamples when post-intervention periods are defined as the first 10 days after an intervention episode.[39] Evidently, the same patterns emerge. Another aspect worth noting is that trading rule returns in periods that neither coincide with nor are preceded by intervention periods tend to be as large or even larger than returns for the whole sample period. On the one hand, this strongly suggests that technical trading rules are not only profitable during intervention periods and provides further evidence against the hypothesis that central bank interventions are solely responsible for the profitability of technical trading rules. On the other hand, this shows that the exceptionally high trading rule returns during intervention periods are reduced to average technical trading rule returns during the first few days after interventions. This raises the question, to what extent, if at all, the remarkable profitability of technical trading rules during intervention periods can be exploited.

[39] The corresponding figures for the first 5 days and the first 15 days are, again, given in Appendix A.

Figure 4-15: MA Trading Rule Returns During and After Intervention Periods

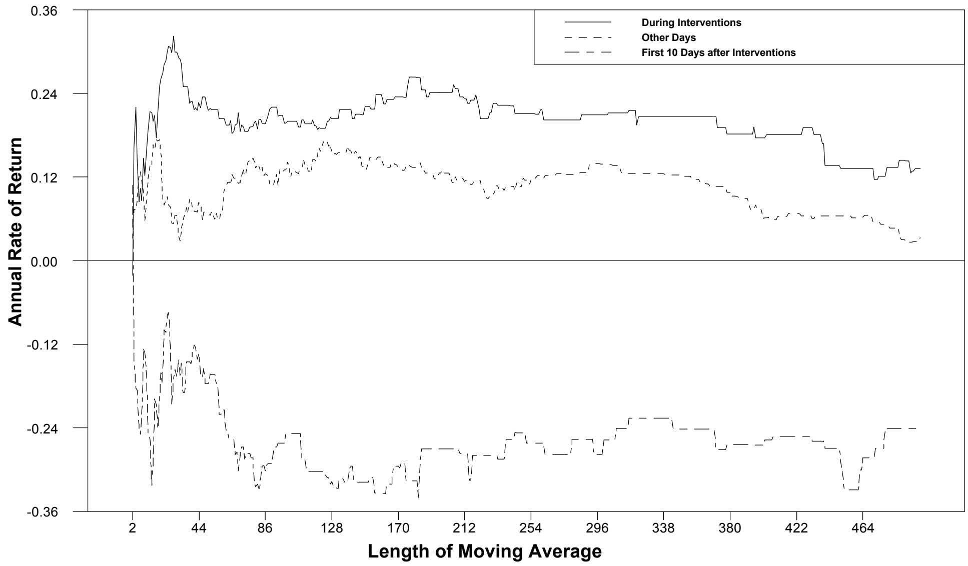

4.5.3 Can Central Bank Interventions Be Exploited?

For the purpose of investigating whether high trading rule profitability on intervention days can be exploited in practice one would ideally like to know the levels of the exchange rate at the exact times at which interventions took place. Alternatively, it would be very useful to have time-stamped Reuters reports on interventions. Since neither are available, we use instead the levels of the exchange rate at the nearest point in time after interventions, which are the levels of the exchange rate at the close on the days on which interventions have taken place. When examining whether knowledge of interventions could have been exploited, we therefore look at the returns from following trading rules between dates t and $t+1$, if intervention has taken place at t. Table 4.7 contains the results for the peak-and-trough-progression rule, differentiating according to which central bank intervened. Annual rates of return lie between 12% and 15%. Standard errors of the returns are such that the null hypothesis of zero returns can be rejected at the 5% significance level (one-sided test). Table 4.7 also contains annual rates of return for the rest of the sample, i.e. the full sample with days after interventions are removed. In tests of a difference in means between the two samples the null hypothesis of no

Table 4.7: Profitability of Trying to Exploit Central Bank Interventions

		Fed	Buba	Either
Peak-and-Trough Progression-Rule	$\mu(r_t \mid x_{t-1} \neq 0)$	0.1502	0.1192	0.1219
	S.E.	0.0855	0.0597	0.0532
	t-value	1.7573	1.9976	2.2908
	$\mu(r_t \mid x_{t-1} = 0)$	0.0674	0.0698	0.0775
	Difference in Means	0.0828	0.0494	0.0443
	S.E. of Difference	0.0920	0.0702	0.0659
	t-value	0.9009	0.7032	0.6722
Bet against Central Bank	Annual Return	0.0215	0.0859	0.0536
	S.E.	0.0805	0.0594	0.0528
	t-value	0.2671	1.4454	1.0152

difference cannot be rejected.[40] It thus seems that the performance of the technical trading rule cannot be improved upon by conditioning on interventions having taken place. Table 4.7 also contains the results of examining the profitability of always betting against the central bank. As the results show, returns are not even significantly greater than zero, let alone greater than returns from following the peak-and-trough-progression rule.

We also analysed whether one could use moving average trading rules to exploit central bank interventions. Figure 4-16 compares annual rates of return from only using moving average trading rules if an intervention by either Fed or Bundesbank has taken place with those on the rest of the days in our sample.[41] Given a standard error of returns of about 0.07, the results appear to be more favourable towards exploitability of central bank interventions.

However, the evidence presented so far rests on the assumption that by the end of each day everybody knows whether an intervention has taken place or not. However, this is unlikely to be the case. Klein (1993) analyses the accuracy of reports of foreign exchange intervention in the New York Times and the Wall Street Journal and finds that the probability that an intervention has taken place, given that it was reported lies between 80 and 90%, whilst the probability of an intervention being reported, given that it has occurred ranged from 70 to 80%.[42]

[40] The test statistic in Table 4.7 is calculated as: $\frac{\overline{X_1}-\overline{X_2}}{\sqrt{S_1^2/N_1+S_2^2/N_2}}$, where $\overline{X}_{\cdot}$, $S_{\cdot}^2$ and $N_{\cdot}$ are the sample mean, the sample variance and the sample size of trading rule returns.

[41] The corresponding figures for interventions by Fed and Bundesbank alone are contained in Appendix A.

[42] Klein (1993), p. 648f. Similar results were found by Bonser-Neal and Tanner (1996), Table 1, Hung (1997), Table 1, and Chang and Taylor (1998), p.194.

Figure 4-16: Profitability of Using MA Trading Rules to Try to Exploit Central Bank Interventions

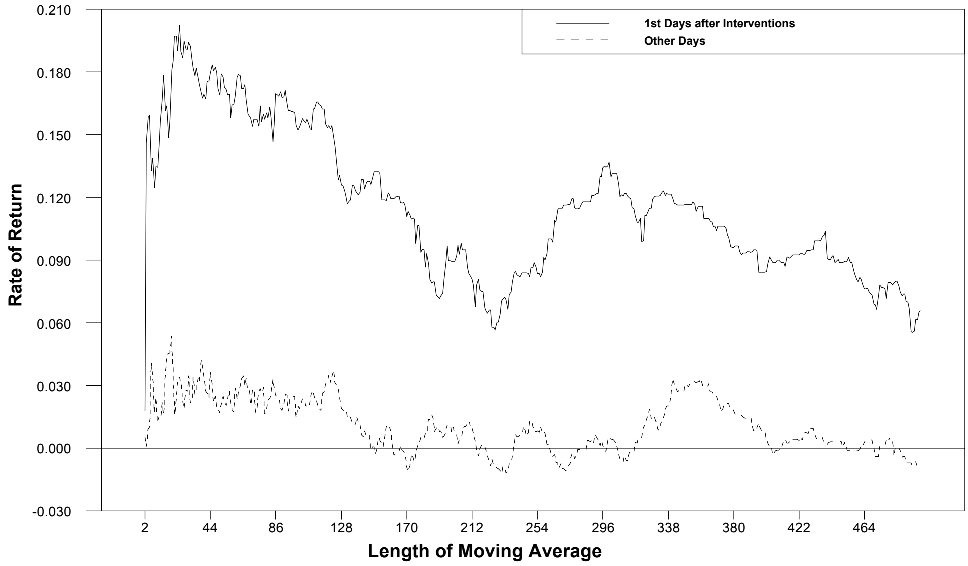

We therefore examine also to what extend it is possible to exploit central bank interventions on the basis of publicly available knowledge. For this purpose we use the data on intervention reports contained in the appendix of Dominguez and Frankel (1993), which spans the period from 1982 to 1990. Intervention reports cited in Dominguez and Frankel (1993) are such that if an intervention is reported in t, it is alleged to have occurred between $t-1$ and t. When examining whether interventions can be exploited, we therefore look at returns from following technical trading rules from t to $t+1$, if intervention was reported in t. Table 4.8 shows the results of trying to use the peak-and-trough-progression rule to exploit intervention reports. Rather than having diminuished, the difference between returns on days after intervention reports and the rest of the sample has grown slightly. Nonetheless, the difference in means is still not significant (though this may be partly due to a decreased sample size of 261 instead of 1137).[43] Figure 4-17 give the corresponding results for the moving average trading rule. As for

[43]The fact that less than a quarter of interventions were reported might seem to stand in contrast to the above mentioned results by Klein (1993). The reason for the difference is that Klein considers only the period 1985 to 1990. Dominguez and Frankel (1993), p. 73, note that before 1985 a very small proportion of interventions was reported.

Table 4.8: Profitability of Trying to Exploit Intervention Reports Using the Peak-and-Trough-Progression Rule

	Bundesbank or Fed
$\mu(r_t \mid x_{t-1}^{rep} \neq 0)$	0.1365
S.E.	0.1163
t-val.	1.1733
$\mu(r_t \mid x_{t-1}^{rep} = 0)$	0.0432
Difference in Means	0.0932
S.E. of Difference	0.1247
t-val.	0.7476

the peak-and-trough-progression rule, returns are larger after reported than after actual interventions.[44] Evidence of a difference in means is slightly weaker than for actual interventions, which is likely to be due to the doubling of the standard error of the difference as a result of the sample size being only a quarter as before. The fact that the value of the test statistic still exceeds 1.2 for all lengths of the moving average between 10 and 170 is very suggestive of the fact that it is indeed possible to exploit central bank interventions using technical trading rules.[45] Whilst this conclusion is an uncomfortable one to draw, our previous findings suggest that the reason for exploitability is that interventions confirm the existence of strong technical trends. Thus the dilemma central banks face is that they want to stop trends away from fundamentals, but that by intervening they signal the existence of such a trend. It is interesting to note that this might explain why market participants are so concerned about the behaviour of central banks. Moreover, it might be one reason why central banks often conduct interventions secretly (and are so unwilling to reveal their intervention data to researchers).

[44] N.B. This increase in the profitability of trading rules when compared with acting on actual interventions is most likely random. However, it is interesting to note that it disappears if small interventions (defined as ones with a volume smaller than $20m) are excluded from the sample of actual interventions. Given that large interventions are more likely to be reported (Klein (1993), Table 6), this (somewhat paradoxically) suggests that it is more profitable to speculate against large interventions than against small ones. Another way of looking at this phenomenon is to say that the larger the interventions are, the stronger is the trend that is intended to be countered and thus the greater are the returns to technical trading.

[45] If might be objected that one should withhold judgement before an analogous examination has been carried out using intraday data and time-stamped intervention reports. However, given that the assumptions concerning the information possessed by traders are weaker, it seems unlikely that such an examination should find less evidence of the exploitability of central bank interventions.

Figure 4-17: Profitability of Using MA Trading Rules to Try to Exploit Intervention Reports

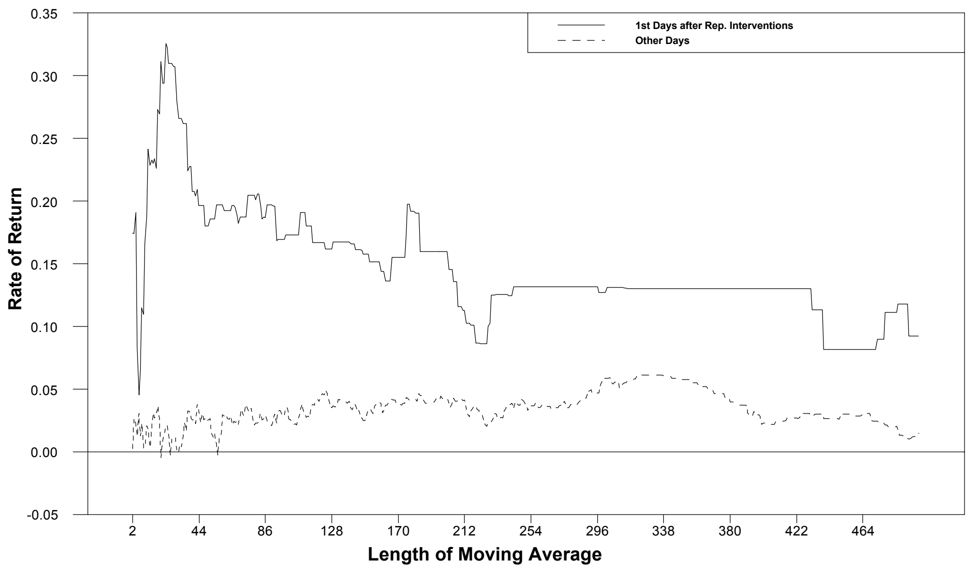

4.6 Discussion

This chapter aimed at clarifying the relationship between central bank interventions and the profitability of technical trading rules. The starting point in this exercise was the puzzling (and so far not satisfactorily explained) empirical regularity that at the same time technical trading rules earn the greatest part of their profits by trading against central banks and that central banks nevertheless appear to make profits with their interventions. After showing that this seeming contradiction was robust to considering a wider range of trading rules and a more complete set of central bank interventions, we went on to suggest that the origin of the puzzle lies in the fact that trading rule profits and intervention profits are measured over different horizons and that exchange rates only move in a manner consistent with central banks' intentions in the long run. This explanation, however, still left us with the question why technical trading rules are so profitable when central banks intervene. For the purpose of answering this question, we first investigated what were the motivations behind central bank interventions and provided some evidence suggesting that central banks react to technical trading. We then argued that if central banks react to technical trends in exchange rates they will only intervene when a trend is already well under way and stop as soon as the exchange rate moves back towards its fundamental value. If this is the case, however, then central bank interventions will tend not to take place when extrapolative technical trading rules are unprofitable (namely when either there is no trend or when a trend has just turned). In order to check this explanation, we looked at post-intervention returns and found that they were highly negative as predicted. Moreover, we found that trading rule returns on days that neither coincide with nor are preceded by interventions are positive and are about as large as trading rule returns for the entire sample, which implies that even if interventions were a cause of trading rule profitability, they would not be the only one. Finally, we examined whether it is possible to exploit the observed high trading rule profitability on intervention days and provided some evidence that this is indeed possible. However, exploitation does not take the form of simply betting against the central bank whenever it intervenes, but rather that of following technical trading rule when finding out about interventions.

Our results concerning the behaviour of exchange rates immediately after interventions are largely consistent with the existing literature on the effectiveness of central bank interventions. Econometric estimates of the influence of interventions on the level of the exchange rate are usually insignificant and even if they are significant, they have the wrong sign.[46] Casual observations, on the other hand, have been more favourable towards the effectiveness of interventions. For instance, Dominguez and Frankel (1993) observe that:[47]

> ... in 10 out of 11 episodes, during the period of intervention, the mark-dollar rate moved in the opposite direction to intervention operations. But ... in 10 of the 11 episodes, in the month following the end of the intervention operations, the mark-dollar rate moved in the same direction as the operations.

Whilst Catte et al. (1994) find that:[48]

> All of the episodes were successful in the sense that interventions inverted the trend of the dollar and ...; in nine cases they were definitively successful in the sense that

[46] E.g., Baillie and Osterberg (1997), Table 1, Weber (1994), p. 16ff, Dominguez and Frankel (1993), p. 113ff. Slightly more successful is Dominguez (1990), p. 143ff. See Almekinders (1995), p. 77ff, for a very thorough survey of the literature.

[47] Dominguez and Frankel (1993), p. 95.

[48] Catte et al. (1994), p. 206.

> in the next episode intervention was in the opposite direction. ... Three were short lived lasting no more than three weeks (...), while the remaining episodes (...) should probably be considered as successful *tout court* because their effects either lasted for several months or were interrupted by minor rebounds that induced central banks to intervene again in the same direction (...).

Of course, claims such as these are often attacked on the grounds that they are not based on an unambiguous methodology.[49] It is nonetheless interesting to note that our finding that exchange rates move in a direction contrary to central banks' intentions in the short run and in the desired direction in the long run describes the same phenomenon. This leads us to the question why exchange rates behave in this manner.

The theoretical literature on the effects of sterilised central bank interventions distinguishes two main channels of influence:[50] The portfolio balance and the signalling channel.[51] The portfolio balance channel refers to interventions influencing exchange rates by changing the relative supply of domestic and foreign assets. However, apart from the fact that empirical evidence suggests that the influence of interventions on exchange rates through this channel is quite weak[52], it predicts that exchange rates move immediately in the direction desired by the central banks. As we have seen, however, the opposite tends to be the case.

Interventions are said to influence exchange rates through the signalling channel if interventions are leading indicators of monetary policy. Suppose, for instance, that the Fed sells USD. If this is an indication that the Fed will lower interest rates in the future, then traders will tend to sell USD today, which leads to an immediate decrease in the value of the USD. Moreover, this depreciation makes it relatively cheap for the Fed to rebalance its reserves. Thus it seems as if both the fact that central banks earn profits from interventions and the fact that exchange rates move in accordance with interventions in the long run can be explained in this way. However, as with the portfolio balance channel, there are two snags: Firstly, the empirical evidence concerning the question whether interventions are leading indicators of monetary pol-

[49] E.g. Almekinders' (1995), p.86, critique of Dominguez and Frankel (1993) and Truman's (1992) comment on the paper by Catte et al. (1994).

[50] N.B. The subsequent arguments assume that both Fed and Bundesbank sterilise their interventions, which is at least what they claim they do. See Dominguez and Frankel (1993), p. 86f, for a discussion of this issue.

[51] See Dominguez and Frankel (1993), ch. 4, for a discussion of these channels.

[52] See, e.g., Ghosh (1992).

icy is, at best, inconclusive.[53] And secondly, the signalling story still leaves open the question why exchange rates move perversely in the short run.

Recently, attempts have been made to approach the question of the effects of central bank interventions from a different direction. Bhattacharya and Weller (1997) and Eijffinger and Verhagen (1997) model interventions in a framework of strategic interaction between central banks and rational speculators. The essential element in these models is that the central bank has an exchange rate target that deviates from the market's assessment of the 'true' value of the exchange rate. Information concerning its target exchange rate is private to the central bank and is partially revealed through interventions. The strength of these models is that they can explain some of the puzzling stylised facts concerning central bank interventions: Bhattacharya and Weller (1997) give an explanation why central banks have an incentive to intervene secretly and why the exchange rate might react perversely to interventions, whilst Eijffinger and Verhagen (1997) in addition explain why interventions might lead to an increase in exchange rate volatility. These models cannot explain all the facts, however. Both models imply that central banks lose money with interventions, which runs counter to the evidence we presented in this chapter. Moreover, adapting the models to take account of the profitability of interventions is likely to be a hard task, because in these models the actions of the central bank are determined as a result of balancing intervention losses and deviations of the exchange rate from its target value. Another shortcoming of these models is that they cannot explain why exchange rates eventually move in accordance with central bank's intentions.

As an alternative explanation, we suggest the following interpretation of the short term behaviour of exchange rates: Suppose that initially the exchange rate behaves like a random walk. At some point, purely by chance, something that looks like a trend in the exchange rate materializes. This is spotted by technical traders, who extrapolate this trend and enter the market. This leads to a continuation of the trend, which in turn attracts more traders to jump onto the bandwagon. Sooner or later a central bank (and maybe also other traders) decides that the exchange rate has moved too much out of line with fundamentals and starts intervening to stop this trend. However, initially there are still more extrapolative traders (maybe attracted by the apparent success of the trading rules in recent history), who enter the market. The trend

[53] See Lewis (1995a), Kaminsky and Lewis (1996) and Ghosh (1992).

continues as long as there is more demand from additional extrapolative traders than there is supply from continued central bank interventions. At some point (though it may be a long time before this point is reached), however, the trend turns and the exchange rate moves back towards its fundamental value. As soon as it starts doing this, interventions stop.

This interpretation can explain why technical trading rules tend to trade against central banks and why this tends to be profitable. At the same time it can account for the fact that central banks nevertheless earn profits with interventions (they buy low and sell high, as Friedman (1953) would say). Last not least, according to this interpretation technical trading rule returns are high and exploitable when central banks intervene because interventions confirm the existence of strong technical trends. However, although our interpretation thus appears to square with all empirical regularities that we presented in this chapter, it clearly is rather vague. It therefore remains a challenge for future theoretical research to formalise at least some of its aspects in order to gain a clearer understanding of the mechanisms underlying the interaction between central banks and technical traders.

Chapter 5

Technical Trading Rule Profitability in the EMS

5.1 Introduction

The results in the previous chapter suggest that central banks intervene on foreign exchange markets to counteract technically motivated trends. Bearing this in mind, one could interpret the decision to participate in a fixed exchange rate regime as a more determined measure introduced (at least partly) for the same purpose. If this is the case, the success of a move to fixed exchange rates can be determined simply by examining whether technical trading rules are profitable in fixed exchange rate regimes. This is done in this chapter for the case of four exchange rates which participated in the Exchange Rate Mechanism (ERM) of the European Monetary System (EMS).

There exists only one paper which studies technical trading rule profitability for fixed exchange rates.[1] Neely and Weller (1998) use genetic programming to investigate whether there exist detectable time dependencies in four EMS exchange rates versus the DEM and the corresponding interest differentials during the period from 1979 to 1996. Whilst the period from

[1] N.B. Lee and Mathur (1996) examine the profitability of Moving Average trading rules for six European cross-rates. These are JPY/DEM, JPY/GBP, JPY/Swiss Franc (SFR), DEM/SFR, DEM/GBP and GBP/SFR. Of these, only the DEM/GBP exchange rate was part of a fixed exchange rate regime and only for less than two years. For this reason we do not consider Lee and Mathur (1996) as a paper on the profitability of technical trading rules in fixed exchange rate regimes.

13/3/1979 to 1/1/1986 is used to generate and select trading rules, the period from 1/2/1986 to 6/21/96 is used to examine the ex-ante profitability of 100 selected trading rules. Taking averages over all 100 rules, they find an average annual rate of return for the Dutch guilder (NLG) of 0.06%, of 0.86% for French franc (FRF), of 2.45% for the Italian lira (ITL) and of 2.75% for the British pound (GBP). For the ITL trading rule returns increased from 0.43% to 6.04% after it stopped being member of the ERM. There is also a slight increase in profitability for the FRF from 0.69% to 1.31% after the fluctuation bands were widened to 15%. In contrast, the GBP trading rules are more profitable when the pound was a member of the ERM than when it was not (4.37% vs. 2.38%). Moreover, Neely and Weller (1998) find that the trading rules did not tend to predict realignments and that the best trading rules rely predominantly on interest differentials. Finally, they compare their trading rule returns with those from following 4 moving average and 4 filter rules and find that these perform less well than the trading rules based on genetic programming.

The first point to note about Neely and Weller's (1998) results is that their claim that 'economically significant' profits could have been made for EMS exchange rates is somewhat overstated. For a start, profits for the FRF and the NLG are tiny. Moreover, the same is true for the ITL whilst it was a member of the ERM. Thus the only case where non-negligible profits were made for an exchange rate which was part of a fixed exchange rate regime was the GBP, for which they found average returns of 4.37%. Not even this is, however, significantly greater than zero.[2] Thus, as regards the profitability of technical trading rules for EMS currencies, Neely and Weller (1998) show little more than that there is some evidence for the ITL after it left the ERM. As regards the alleged superiority of genetic programming based trading rules versus moving average and filter rules a few points are worth making: Firstly, the allegation that conventional trading rules are inferior is not based on any kind of formal test and the choice of rules considered is somewhat arbitrary. Moreover, given that genetic programming based trading rules did not yield significant returns for most of the currencies and for most of the time, this comparison might be considered somewhat futile anyway.

A further limitation with the study by Neely and Weller (1998) is that their results are

[2] Only when portfolio rules are constructed on the basis of the 100 rules generated do they find that returns are significantly greater than zero (Neely and Weller 1998, p. 11f).

based on a validation period beginning in 1986. However, it is of particular interest to analyse the profitability of technical trading rules for the period between 1979 and 1986 because re-alignments of ERM exchange rate parities were frequent in this period. Another question which Neely and Weller (1998) leave open is whether there exist significant differences in trading rule profitability during periods when an exchange rate was a member of the ERM and when it was not. These two issues will be addressed in the subsequent sections. Our first result is that technical trading rules were not even profitable when ERM parities were non-credible. Moreover, we find that this is to some extent due to interest differentials compensating for subsequent exchange rate movements. As regards the question whether there exist regime-specific differences in trading rule profitability, we examine the cases of the ITL and the GBP and find strong evidence that participation in the ERM significantly reduces technical trading rule profitability.

The structure of this chapter is as follows: After discussing the data in Section 5.2, we address the relationship between credibility and trading rule profitability in Section 5.3. Section 5.4 analyses the effect of leaving/joining the ERM on trading rule returns. The results are discussed in Section 5.5.

5.2 Data Summary

The analysis uses daily exchange rate and interest rate data running from January 2, 1979, to November 28, 1997. We consider the Deutsche Mark (DEM) exchange rates of the Dutch Guilder (NLG), the French Franc (FRF), the Italian Lira (ITL) and the British Pound (GBP). Interest rates are daily overnight eurorates for each currency.[3] Table 5.1 gives the summary statistics of log first differences of daily exchange rates and of the daily interest differential versus the German mark. Both for the FRF and for the ITL the mean daily change in the value of the exchange rate versus the DEM is significantly smaller than zero. All exchange rate changes are skewed to the left (i.e. towards depreciation against the DEM), there is strong evidence of excess kurtosis (fat tails) and thus, not surprisingly, the Null-hypothesis of Normality can be rejected in a Jarque-Bera test for all exchange rates. The interest differentials versus Germany

[3] Exchange rates are ECU exchange rates at 2:15 Brussels time (prior to September 1988: 2:30 pm) as communicated by the Commission of the European Communities. Euromarket rates are Bid rates around 10 am Swiss time. Source: BIS.

Table 5.1: Exchange Rate and Interest Differential Summary Statistics

	NLG		FRF		ITL		GBP	
	Δs_t	$\frac{1}{260}\Delta i_t$	Δs_t	$\frac{1}{260}\Delta i_t$	Δs_t	$\frac{1}{260}\Delta i_t$	Δs_t	$\frac{1}{260}\Delta i_t$
Mean^a	-0.067	-0.152*	-0.408*	-1.2*	-1.58*	-2.10*	-0.522	-1.58*
Std. Dev.^{a}	6.41	0.423	15.23	1.38	35.20	1.52	50.32	1.14
Skewness	-0.665*	-1.163*	-0.720*	-5.903*	-2.59*	-6.901*	-0.66*	0.309*
Kurtosis	17.40*	14.98*	26.47*	74.56*	48.45*	119.04*	8.336*	2.484*
Jarque-Bera	40866*	29806*	108032*	105314*	108925*	273322*	5905.27*	129.57*

^{a} $(*10^4)$. * indicates significance at the 1% level.

Table 5.2: Phillips-Perron Unit Root Test for Daily Changes in EMS Exchange Rates

Equation	$\Delta s_t - \Delta s_{t-1} = \alpha + \beta \Delta s_{t-1} + u_t$			
	DEM/NLG	DEM/FRF	DEM/ITL	DEM/GBP
$\widehat{\alpha}$	-8.94E-06	-6.25E-05	-0.00151	-7.10E-05
S. E.	9.32E-06	2.75E-05	5.22E-05	7.43E-05
t-value	-0.959	-2.270	-2.896	-0.9566
$\widehat{\beta}$	-1.076	-0.9955	-0.9854	-0.9850
S. E.	0.0145	0.0148	0.0147	0.0147
t-value	-74.17	-67.10	-66.88	-67.08
PP Test Statistic	-73.80**	-67.12**	-66.91**	-67.14**

* (**) indicates significance at the 5% (1%) level. Truncation lag equals 9.

are significantly negative on average for all currencies and again there is strong evidence of skewness and excess kurtosis.

All four exchange rate time series were also checked for non-stationarity. Table 5.2 contains the results of a Phillips-Perron test. The Null-hypothesis of a unit root is strongly rejected for all exchange rates. As for the case of flexible exchange rates we carried out Engle's (1982) ARCH-test of conditional heteroskedasticity. Table 5.3 contains the results. The Null-hypothesis of homoskedasticity is strongly rejected for all exchange rates except the FRF. An analysis of the autocorrelation function of exchange rate changes revealed evidence of negative autocorrelation for the NLG and FRF (of first and second order respectively) and of positive autocorrelation for the ITL and the GBP (of fourth and tenth order respectively).[4] This is a first indication of a difference in the time series characteristics of fixed compared to flexible exchange rates and foreshadows to some extent the results of section 5.4.

[4]The autocorrelation functions are to be found in Appendix B. See Tables B.6 to B.9.

Table 5.3: ARCH-Test for Daily Changes in EMS Exchange Rates

Equation	$(s_t - \overline{s})^2 = \alpha + \beta(s_{t-1} - \overline{s})^2 + u_t$			
	DEM/NLG	DEM/FRF	DEM/ITL	DEM/GBP
$\widehat{\alpha}$	2.88E-07	3.42E-06	1.07E-05	2.05E-05
S. E.	2.34E-08	8.58E-07	1.28E-06	1.06E-06
t-value	12.31	3.989	8.375	19.31
$\widehat{\beta}$	0.2725	0.0146	0.1389	0.1868
S. E.	0.0135	0.0148	0.0146	0.0144
t-value	20.13	0.9911	9.499	12.964
F-statistic	406.06	0.9824	90.23	168.07
p-value	0	0.3217	0	0
$Obs * R^2$	372.35	0.9826	88.53	162.19
p-value	0	0.3216	0	0
Durbin-Watson	2.0449	2.0454	2.0589	2.0748

5.3 Profitability and Credibility

In an influential paper Paul Krugman (1991) showed how the introduction of a perfectly credible exchange rate target zone would lead to a reduction in exchange rate volatility due to rational stabilising speculation (the so called 'honeymoon'-effect). Krugman's target zone model has been subjected to a number of empirical tests and has, on the whole, performed quite poorly.[5] In an extension of the target zone model, Bertola and Svensson (1993) allow for imperfect and time-varying credibility and introduce a method for measuring the credibility of fixed exchange rate parities. The general consensus in the literature that tries to quantify the credibility of EMS exchange rate parities is that they have not been credible before 1987 but credible between 1987 and 1992.[6] Since the ERM fluctuation bands have been widened in August 1993 to ±15%, the analysis of the credibility of the central parities has become obsolete.

When examining the relationship between the profitability of technical trading rules for the EMS exchange rates and the credibility of their underlying central parities, we will therefore distinguish three regimes:[7] Firstly non-credible narrow fluctuation bands, secondly credible narrow exchange rate parities and thirdly wide/no fluctuation bands. Consequently, we will

[5] See, e.g., Bertola and Caballero (1992).

[6] See, e.g., Rose and Svensson (1994). An exception is the Dutch guilder, whose central parity is said to have been credible from 1983 onwards.

[7] N.B. To avoid repetition, in this section we consider explicitly only cases of the NLG and the FRF. The ITL and the GBP will be dealt with separately in the next section.

Table 5.4: Profitability of Peak-and-Trough-Progression Rule for the DEM/NLG and the DEM/FRF Exchange Rate

		Full Sample	1st Period	2nd Period	3rd Period
NLG	Mean	-0.0166	-0.0127	-0.0192	-0.0139
	S.E.	0.0028	0.0105	0.0027	0.0029
	t-val	-6.0265	-1.2103	-7.0768	-4.8167
	Gross Return	-0.0001	0.0036	-0.0020	0.0012
	Interest Eff.	0.0010	0.0018	0.0009	0.0003
	Transaction Cost	-0.0168	-0.0170	-0.0175	-0.0148
FRF	Mean	-0.0142	-0.0177	-0.0170	-0.0035
	S.E.	0.0046	0.0060	0.0024	0.0160
	t-val	-3.1150	-2.9403	-7.1131	-0.2155
	Gross Return	0.0021	-0.0007	-0.0001	0.0105
	Interest Eff.	0.0011	0.0013	0.0010	0.0007
	Transaction Cost	-0.0168	-0.0174	-0.0173	-0.0148

split up the sample slightly differently for each of the exchange rates considered: For the NLG, the subsamples are: 2/1/1979-31/12/1982, 2/1/1983-1/8/1993 and 2/8/1993-28/11/1997. For the FRF they are: 2/1/1979-31/12/1986, 2/1/1987-1/8/1993 and 2/8/1993-28/11/1997.[8] Table 5.4 shows the results of using the peak-and-trough-progression rule for both exchange rates. The first point to note is that trading rules do not appear to be profitable even during periods of low credibility of the underlying exchange rate parities. The table also shows the results of splitting up trading rule returns into the components gross returns, interest effects and transaction costs. In general, trading rule returns tend to be dominated by transaction costs. In order to see whether these results are robust, we carried out the same calculations for the moving average trading rule. Figures 5-1 and 5-2 show annual rates of return from following moving average trading rules for both exchange rates during the three periods.

For all exchange rates, for all lengths of the moving average and for all periods, returns are smaller than 1%. It is interesting to examine more closely why the trading rules were not profitable during the period when the exchange rate parities were not credible. For this purpose net annual returns for the DEM/FRF exchange rate are again split up into gross annual rates of return, average annual transaction costs and annual interest effects. Figure 5-3 shows that

[8] It is debatable whether one should not rather choose September 1992 as the end of the second subperiod for the NLG and FRF. However, since neither currency was devalued between September 1992 and August 1993 the difference is likely to be small.

Figure 5-1: Annual Rates of Return from Following MA Trading Rule for the DEM/NLG Exchange Rate

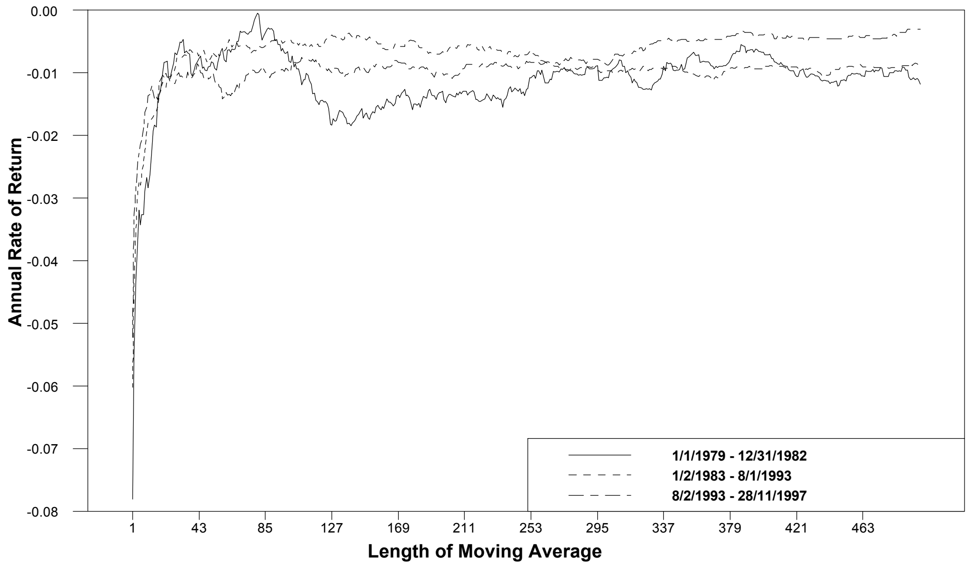

Figure 5-2: Annual Rates of Returns from Following MA Trading Rule for the DEM/FRF Exchange Rate

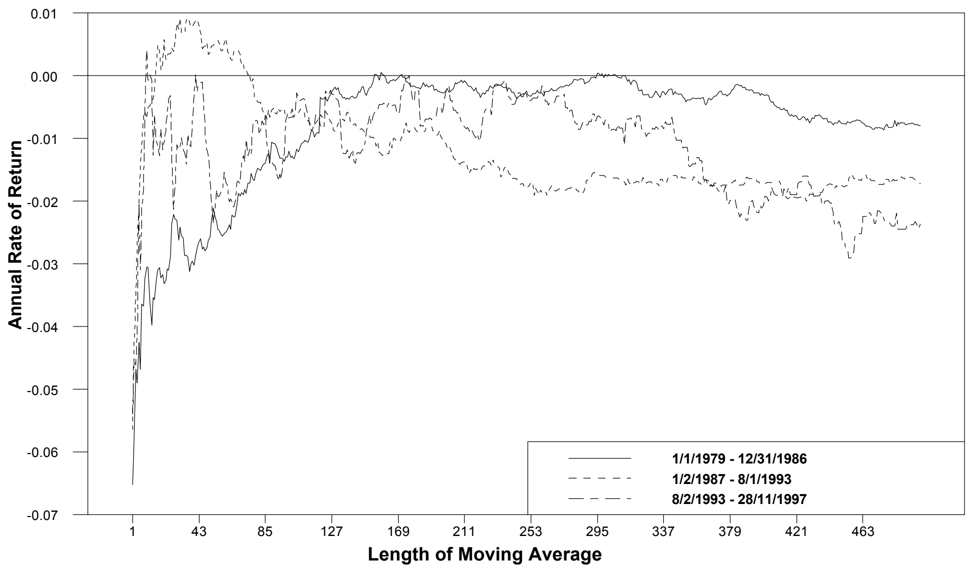

Figure 5-3: Components of Moving Average Trading Rule Returns for the DEM/FRF Exchange Rates

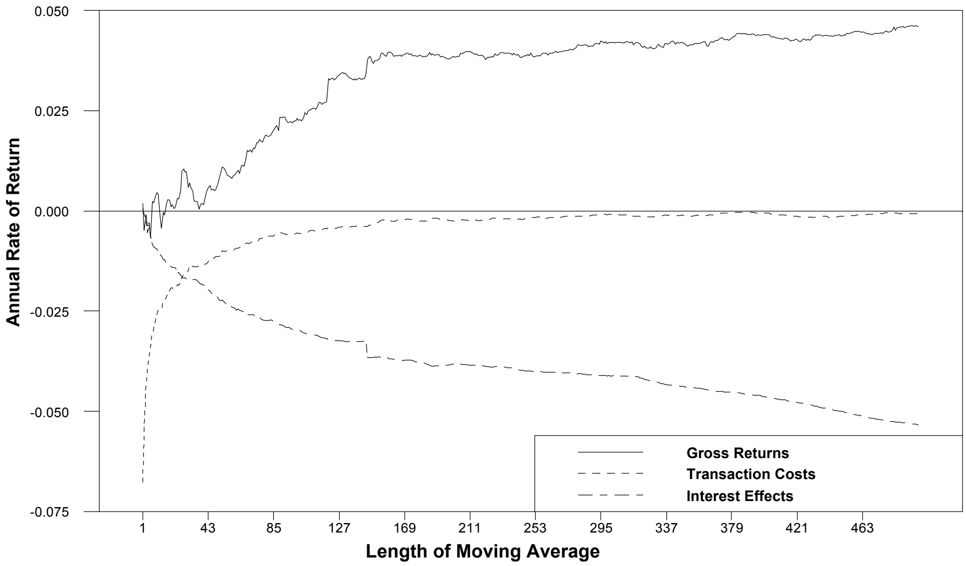

gross returns for the FRF almost reach 5% in this period for long lengths of the moving average. The reason for this is that for long laglengths moving average trading rules pick up long run trends in exchange rates and that the FRF depreciated considerably against the DM during the period between 1979 and 1987. These returns were, however, overcompensated by the interest differential.[9] This suggests that depreciations against the DEM were taken ahead in the interest differentials and can be interpreted as casual evidence in favour of uncovered interest parity.

5.4 Effects of ERM-Membership on Trading Rule Profitability

The results in the previous section suggest that technical trading rules are unprofitable when applied to exchange rates that belong to a fixed exchange rate regime, whilst earlier we found considerable evidence in favour of trading rule profitability when they were applied to floating exchange rates. This might be interpreted as evidence that the introduction of fixed exchange rates with fluctuation bands was a way of making technically motivated trading unprofitable.

[9] N.B. The same effect is observable for the DEM/ITL exchange rate. The corresponding figure is contained in Appendix A.

Table 5.5: Profitability of Peak-and-Trough-Progression Rule for the DEM/ITL and the DEM/GBP Exchange Rate

		Full Sample	1st Period	2nd Period	3rd Period
ITL	Mean	0.0047	-0.0136	-0.0163	0.0555
	S.E.	0.0136	0.0131	0.0127	0.0429
	t-val	0.3437	-1.0349	-1.2841	1.2928
	Gross Return	0.0262	0.0182	-0.0015	0.0687
	Interest Eff.	-0.0043	-0.0129	0.0027	0.0011
	Transaction Cost	-0.0166	-0.0171	-0.0169	-0.0153
GBP	Mean	0.0514	0.0636	0.0180	0.0365
	S.E.	0.0193	0.0249	0.0353	0.0400
	t-val	2.6593	2.5565	0.5104	0.9133
	Gross Return	0.0628	0.0729	0.0339	0.0509
	Interest Eff.	0.0044	0.0059	0.0006	0.0026
	Transaction Cost	-0.0156	-0.0157	-0.0156	-0.0154

The problem with this interpretation is that it is unclear whether it is the fact of belonging to a fixed exchange rate regime that makes trading rule unprofitable or whether this lack of profitability is not rather due to aspects specific to the exchange rates considered. However, this problem can be dealt with by examining the cases of the ITL and the GBP, both of which were members of the ERM during part of our sample period and did not participate at other times: Whereas the Lira was a member of the ERM from the outset, the Pound only joined the ERM in October 1990. Both currencies left the ERM on September 16th 1992.

Correspondingly, we will examine the profitability of the peak-and-trough-progression rule and of moving average rules for both currencies during different subsamples. For the ITL these subsample are 2/1/1979-31/12/1986, 2/1/1987-16/9/1992 and 17/1993-11/1997, whilst for the GBP they are 1/2/1979-5/10/1990, 8/10/1990-9/16/1992 and 9/17/1992-11/28/1997. Table 5.5 shows the results for the peak-and-trough-progression rule. Whilst there is no evidence of profitability of the trading rule for the ITL during its membership in the ERM (again no evidence even during the period of low credibility before 1987), we find a pronounced increase in the profitability of the trading rule after the ITL left the ERM. This confirms Neely and Weller's (1998) result. As regards the GBP, however, we do not find any signs of profitability of the trading rule for the period when it belonged to the ERM. Like Neely and Weller (1998) we find that trading rule returns for the GBP during the third period are quite small.

Figure 5-4: Annual Rates of Return from Follwing MA Trading Rules for the DEM/ITL Exchange Rate

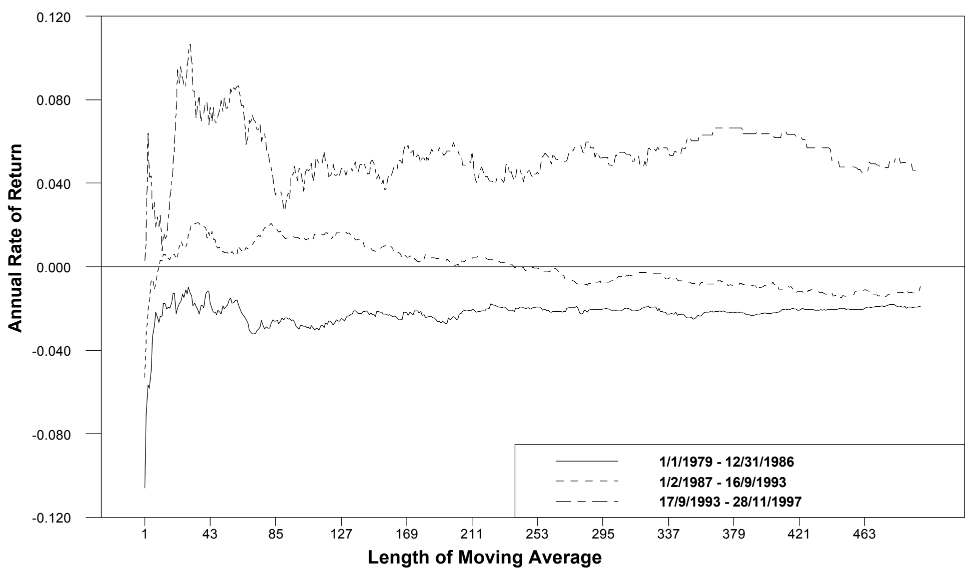

In order to check the reliability of the results we carried out analogous calculations for the moving average trading rule. The results are contained in Figures 5-4 and 5-5. Again, technical trading rule profitability during membership in the ERM appears to behave quite differently than during non-membership. Moreover, for GBP there is no evidence of a reduction in profitability of the moving average trading rule for the last subperiod, which suggest that Neely and Weller's finding of a lack of trading rule profitability is not robust.

It suggests itself to test statistically for a difference between membership and non-membership periods. As a starting point, we test the hypothesis that the difference in mean returns between exchange rate regimes is zero against the hypothesis that it is positive by considering the following test statistic for each individual trading rule:

$$t = \frac{(\bar{x}_{out} - \bar{x}_{in})}{\sqrt{\frac{s^2_{out}}{N_{out}} + \frac{s^2_{in}}{N_{in}}}} \tag{5.1}$$

where $\bar{x}_{(.)}$, $s^2_{(.)}$ and $N_{(.)}$ denote the mean, variance and size of the sample of trading rule returns for each regime. For the DEM/ITL exchange rate we test for a difference between 1/2/1979-

Figure 5-5: Annual Rates of Return form Following MA Trading Rules for the DEM/GBP Exchange Rates

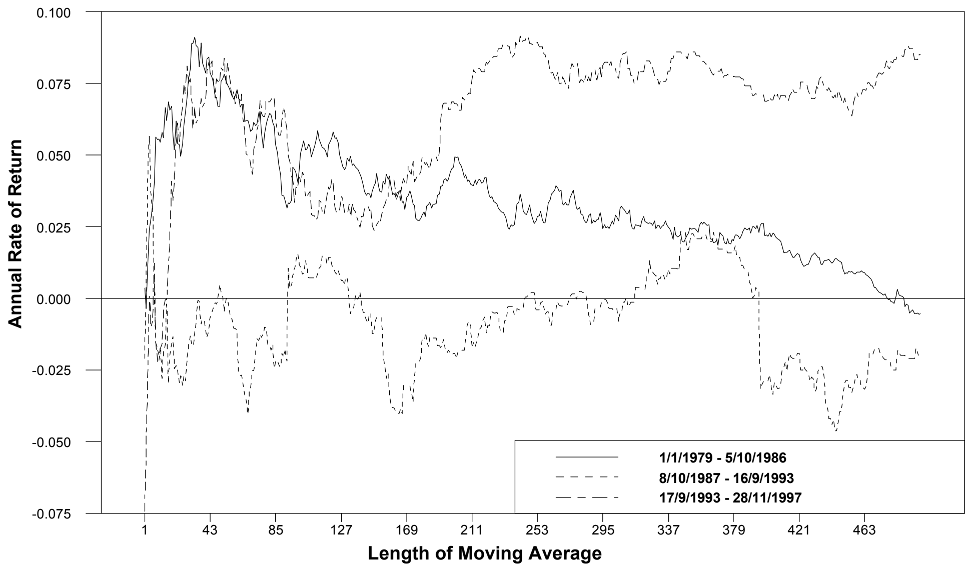

9/16/1992 and 9/17/1992-11/28/1997, whilst for the DEM/GBP exchange rate, we test for a difference between 1/2/1979-5/10/1990+9/17/1992-11/28/1997 and 8/10/1990-16/9/1992. Table 5.6 contains the results for the peak-and-trough-progression rule. Evidence of a difference in profitability between the two regimes is at best marginal. Figure 5-6 shows the values of the test-statistic for varying lengths of the moving average for both exchange rates. Whilst for a wide range of parameter values there are now some indications of a difference in mean returns, the evidence still cannot be considered as conclusive. However, so far we have only tested for a difference in means for each trading rule separately. The next step is to carry out a joint test. This is not straightforward because trading rule returns are correlated, which makes it difficult to derive the joint distribution of the test-statistic analytically.

To deal with this problem, we again use bootstrapping simulations. For the purpose of explaining how this can be done, suppose for an instance that we are dealing with the simpler problem of using bootstrapping simulations to test for a difference in mean returns of only one trading rule between two subsamples. Suppose further that the two subsamples have sizes N_1 and N_2. In this case we must draw random samples of size $(N_1 + N_2)$ with replacement from

Table 5.6: Test of a Difference in Mean Peak-and-Trough-Progression Rule Returns due to Joining/Leaving the EMS

	ITL	GBP
Membership	-0.0147	0.0180
Non-Membership	0.0555	0.0552
Difference	0.0702	0.0372
S.E. of Difference	0.0439	0.0412
t-value	1.5989	0.9033
Bootstrap p-value	0.0375	0.1155

Figure 5-6: Test of a Difference in Mean MA Trading Rule Returns due to Joining/Leaving the EMS

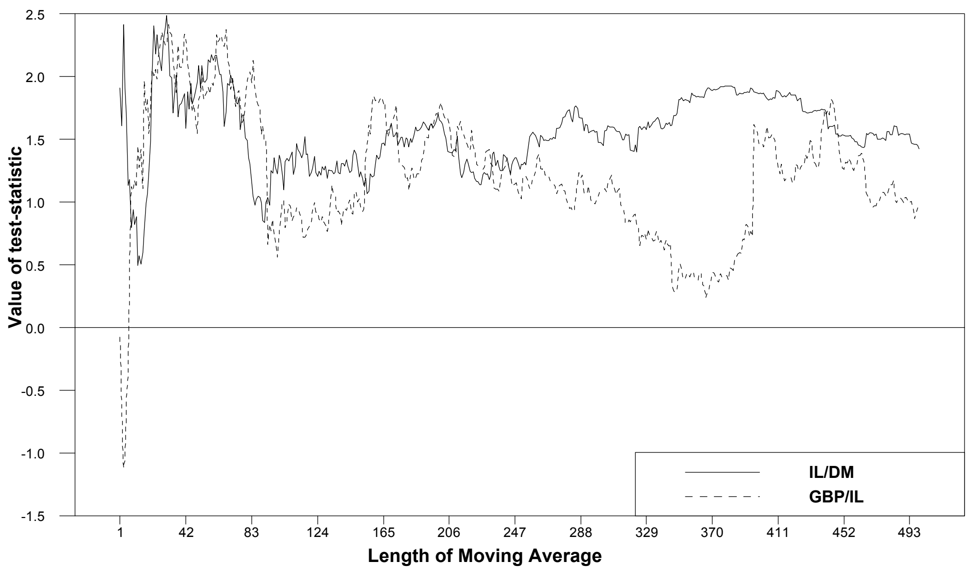

the combined subsamples. Each such random sample is split into two subsamples of size N_1 and N_2 and the value of the test-statistic (5.1) is calculated. Repeating this procedure 2000 times then enables us to derive an empirical distribution of the test-statistic.

Since we want to carry out a joint-test for several trading rules, we need to proceed slightly differently. In order to preserve the autocorrelation in trading rule returns, we need to draw samples of trading rule returns in a 'parallel fashion', which is to say that returns of different trading rules must be sampled on the same dates. For each rule the value of the test-statistic is then determined. Repeating this procedure, we can determine the proportion of simulations in which the average of the simulated test-statistics is greater than or equal to the corresponding average for the original series. If this proportion is small, we can reject the hypothesis of no difference in trading rule returns during subperiods.

The problem with this approach is that if we were to apply all trading rules we considered before, we would need to evaluate 2000 times the value of (5.1) of 500 trading rules for a 'time series' containing almost 5000 data points. The computing time needed for this is (at least at present) prohibitive. We therefore do not consider all 500 trading rules but only 11, namely moving average rules with lengths 50, 100, ..., 500 as well as the peak-and-trough-progression rule.[10] Table 5.6 also contains the bootstrapping p-value of this joint test. For the ITL the hypothesis of no difference can be rejected at the 5% significance level. For the GBP, however, the difference between ERM-participation and non-participation is only marginally significant. As regards the evidence for the GBP it should be born in mind the GBP participated only for less than two years in the ERM, whilch implies that the sample is not very large and as a result the test is not very powerful. Moreover, one might argue that another reason why the results are not stronger is that it takes some time before the full effects of joining a fixed exchange rate regime materialise. In the light of these considerations our results represent strong evidence that joining/leaving a fixed exchange rate regime has a significant impact on the profitability of technical trading rules.

[10] Note that testing all trading rules jointly instead of just 11 would make it more likely that one finds evidence of a difference. The test carried out here is therefore conservative.

5.5 Discussion

The analysis of the profitability of technical trading rules for EMS currencies showed that whenever an exchange rate participated in the ERM both the peak-and-trough-progression rule and moving average trading rules were unprofitable. An obvious limitation of the validity of this result is that only two trading rule classes were considered. Whilst it is probably fair to assume that the results will remain valid for other trend-following trading rules like momentum or filter rules, the same cannot be said for technical trading rules in general. Thus, it remains to be seen whether there exist technical trading rules which can exploit time dependencies in ERM exchange rates.[11]

For moving average trading rules we showed that positive gross returns from following the trading rules were to a large extent compensated by interest differentials. In contrast, when examining the profitability of moving average trading rules for flexible exchange rates, we found that interest differentials tended to augment gross trading profits. Given that trading profits on foreign exchange markets are always due to systematic deviations from UIP being exploited, these results square nicely with the fact that UIP works much better for fixed than for flexible exchange rates, as was shown by Flood and Rose (1996).

There also exists a connection between our results concerning the relative profitability of technical trading rules for fixed versus flexible exchange rates and those of the study by Cavaglia et al. (1993), which concerns differences in the process of expectation formation between fixed and flexible exchange rates. Using monthly data on exchange rate expectations for both EMS and non-EMS exchange rates, Cavaglia et al. (1993) find evidence of extrapolative expectation formation only for flexible exchange rates. This leaves us with the interesting question whether EMS exchange rate expectations are not extrapolative because technical trading rules are not profitable, or whether technical trading rules are not profitable because expectations are not formed extrapolatively. Of course, one possibility is that both propositions are true, which is equivalent to saying that there a exists a self-reinforcing relationship between the use and the success of extrapolative expectation formation or technical analysis. If this is the case, then

[11] N.B. This shortcoming is to some extent alleviated by Neely and Weller's (1998) result that genetic programming based trading rules (which are generated on the basis of a very general set of rules) did not perform well during periods when exchange rates belonged to the ERM.

our results concerning the influence of ERM-participation on trading rule profitability for the ITL/DEM and GBP/DEM exchange rates indicate that participation in a fixed exchange rate regime may be one way to break this self-reinforcing chain.

Chapter 6

Conclusions

This investigation dealt with the profitability of technical trading rules and its relationship to central bank interventions and exchange rate regimes. Throughout we analysed the profitability of two classes of trading rules which are known to have been in wide use during the period over which our samples of exchange rate and interest rate data ran. For flexible exchange rates we showed that for both trading rule classes it is extremely unlikely that the observed trading rule returns should have come about by chance and thereby added to the growing body of evidence suggesting that simple technical trading rules have predictive ability in foreign exchange markets. A second important result we confirmed was that the observed trading rule returns do not appear to be a compensation for bearing high levels of risk. In contrast, we have found strong indications that following trading rules is less risky than buying and holding currencies or stock market indices. Moreover, we found no evidence at all of decreasing or even vanishing trading rule returns over time. It thus seems as if over a period of 14 years technical analysis had been a simple way to earn excess returns without having to pay for these with the incurrence of additional risk.

Whilst many economists feel uneasy about embracing this conclusion, some attempts have been made to provide theoretical explanations for the existence and persistence of trading which is not based on economic fundamentals. The model by DeLong et al. (1990), which was already mentioned, tries to explain the success of extrapolative trading by allowing for noise traders whose erroneous beliefs cause enough additional volatility in prices, that rational arbitrageurs are deterred from driving both prices to their fundamental values and noise traders out of the market. The main contribution of DeLong et al.'s (1990) paper is that it demonstrates that it is possible that noise traders themselves create the conditions which are necessary for their survival. Whilst this self-fulfilling component may be an intuitively appealing explanation for the existence of technical traders, in other respects the model of DeLong et al. (1990) is not in agreement with the empirical regularities we found. One implication of the model is that whilst noise traders have higher expected returns than rational arbitrageurs, the volatility of these returns is overproportionally high. Our finding that trading rule returns are less risky than returns from a Buy-and-Hold strategy suggests that a simple identification of noise traders with technical analysts would be inaccurate. Another respect in which the noise trading model is at odds with evidence on technical analysis is that in the model noise traders are the less well informed market participants, whereas in reality technical analysis is particularly popular with foreign exchange professionals.[1] Thus there exist important respects in which the noise trading model needs to be modified in order to capture the nature of technical trading in asset markets more adequately.

As regards the question of the role of central bank interventions for the profitability of technical trading rules, we started off by showing that LeBaron's (1996) finding that technical trading rules are remarkably useful on days on which central banks intervene is robust to considering a wider range of technical trading rules and using not only Fed but also Bundesbank intervention data. This finding strongly suggests that interventions and technical trading are related. The nature of this relationship, however, is not as easy to determine. We argued that the explanation that central banks introduce exploitable time dependencies into the paths of exchange rates and thereby transfer funds to technical traders does not square with our finding of very substantial intervention profits for both the Fed and the Bundesbank. As an alternative explanation we suggested and provided evidence that interventions are related to technical analysis in the sense that interventions are intended to counter the influence of technically motivated trading. If this is the case then the reason for the remarkably high profitability of technical trading rules on intervention days is simply that interventions are more reliable signs for the existence of technical trends than trading rule signals by themselves. However,

[1] See Menkhoff (1997) for an analysis of the characteristics of foreign exchange market participants engaging in technical analysis.

our results also indicate that in the short run technical trends continue in spite of central bank interventions and that exchange rates move in a direction consistent with central banks' intentions only in the long run. Moreover, it emerges that this empirical regularity can even be exploited using information contained in daily newspapers.

We have argued that the peculiar discrepancy between the short and long term behaviour of exchange rates after interventions is not explained by any existing theoretical model of the effects of central bank interventions. It is worth noting, however, that there exists a corresponding result which Frankel and Froot (1990) found analysing market participants' process of expectation formation. Examining survey data on exchange rate expectations, Frankel and Froot (1990) observed that whilst short term (1-week to 3-month) expectations exhibit bandwagon tendencies, in the long term market participants tend to forecast a return to a long-run equilibrium such as Purchasing Power Parity.[2] If one assumes that central bank interventions are intended to counter extrapolative technical trading which drives exchange rates from their fundamental values, then our results indicate that this twist in the expectations of market participants is nothing but an indication of profit-maximising behaviour. Moreover, this interpretation also squares with the results of Taylor and Allen's (1992) study of the use of technical analysis in foreign exchange markets. About 90% of the respondents of their questionnaire survey use at least some chartist input at short horizons (intraday to one week), whilst at long horizons (one year or longer) 85% of respondents view fundamental analysis as more important than chart analysis. Given the results of these studies and the findings made in this investigation, it seems fair to count the observed divergence between short and long term dynamics of exchange rates as an important stylised fact in need of explanation. It remains one of the main challenges of theoretical research in international economics to develop a model that can rationalise it.

Our results concerning the profitability of technical trading rules for EMS exchange rates indicate that for these there do not exist time dependencies which can be exploited by means of technical analysis. Moreover, we have provided strong evidence that the GBP's participation and the ITL's departure from the ERM significantly affected technical trading rule profits.

[2] Frankel and Froot (1990), p. 96ff. See also the evidence in Froot and Ito (1989) and Ito (1990) on the inconsistency of short- and long term exchange rate expectations.

These results are in perfect agreement with the above mentioned paper by Jeanne and Rose (1999). In this paper, a noise trading framework is employed to explain the difference in volatility between fixed and flexible exchange rate regimes. In this model the presence of noise traders can lead to multiple equilibria in flexible exchange rates: One equilibrium with a large number of noise traders and high volatility and one with few noise traders and low volatility. Joining a fixed exchange rate regime like the EMS can in this context be seen as a way to pin down the economy in a low volatility equilibrium. The evidence we provided concerning the influence of joining/leaving the ERM on trading rule profitability can be considered as the closest thing to an empirical verification of the model.

We also found that moving average technical trading rules were unprofitable in fixed exchange rate regimes because gross returns were compensated by interest differentials and we noted that this was at least partly a reflection of the fact that uncovered interest parity (UIP) works much better in fixed than in flexible exchange rate regimes. Again it is worth noting that this finding is in line with Mark and Wu's (1998) recent attempt at explaining the failure of UIP within a noise trading model. If noise trading is responsible for deviations from UIP for flexible exchange rates, then the fact that fixing exchange rates is one way to drive noise traders out of the market might well explain why UIP works much better for fixed exchange rates.

In this context it is also worth noting that we found that technical trading rules were not even profitable during periods of low credibility of the EMS exchange rate parities. This indicates that the rules are of little use as regards predicting the timing of realignments of fixed exchange rate parities and suggests that technical trading and speculative attacks are distinct phenomena, associated with different exchange rate regimes. Whilst technical trading goes on continually in flexible exchange rate markets and may be the root of excess volatility, in a system of fixed exchange rates speculative attacks occur from time to time and cause sporadic, though potentially very costly turbulence. Seen in this way, the choice between exchange rate regimes comes down to opting for the lesser evil.[3]

Two other policy issues should also be addressed: The first is the obvious question what

[3]Arguably, speculative attacks need not occur in fixed exchange rate markets as long as the authorities in charge of fixing the central parities do so only taking into account economic factors (i.e. not political ones, or national pride). Experience shows, however, that this rarely happens (e.g. the UK in 1992).

our findings tell us concerning the desirability of conducting central bank interventions. At first sight it might seem hard to make out a case against interventions, given that they are highly profitable and appear to be in accordance with economic fundamentals. One potential danger might be, however, that interventions come to be seen as confirmations of the existence of technical trends and thereby attract additional technical trading. In this case, they would tend to increase exchange rate volatility and would therefore be undesirable. Given the recent results by Dominguez (1998) and Chang and Taylor (1998), which indicate that interventions tend to increase volatility, and our finding that newspaper reports of interventions are exploitable, this danger should be considered real.

The second policy issue is the more general one of whether a case can be made for the introduction of a Tobin tax, i.e. a small tax on foreign exchange transactions intended to reduce volume and volatility on foreign exchange markets.[4] The model by Jeanne and Rose (1999) suggests that volatility in flexible exchange rate markets is unnecessarily high due to the influence of noise traders. The question is whether introducing a Tobin tax can be an alternative to joining a fixed exchange rate regime when the aim is to drive noise traders out of the market. Experiments with our observed technical trading rule returns show that a tax (modelled as an increase in transaction costs) of between 0.35% and 0.8% would be necessary to render trading rule returns insignificant.[5] Given that technical trading rules yield significant profits in a number of other asset markets (e.g. stock and commodity markets), it might be expected that the introduction of such a tax would indeed drive technical traders out of the market. However, this ignores another possible reaction to the introduction of a Tobin tax, namely that foreign exchange markets move to a place where no such tax is levied. Whilst it might be desirable to forge an international agreement to introduce such a tax, the gains for countries who do not participate are too large for it to seem likely that such an agreement could be reached. Thus, although our results suggest that flexible exchange rates do indeed display the symptoms of a disease that is curable with a Tobin tax, without a sensible plan to make such a tax enforceable, its introduction does not belong onto the political agenda.

[4] Tobin (1978); see also Frankel (1996) for a recent discussion of the Tobin tax.

[5] N.B. These figures are quite close to the 0.5% suggested by Tobin.

Bibliography

[1] Alexander, S.S. (1961) Price movements in speculative markets: Trends or random walks. Industrial Management Review 2, 7-26.

[2] Alexander, S.S. (1964) Price movements in speculative markets: Trends or random walks, No. 2. In The Random Character of Stock Market Prices, ed. Paul H. Cootner. MIT Press, Cambridge, MA.

[3] Allen, H. and Taylor, M.P. (1990) Charts, noise and fundamentals in the London foreign exchange market. Economic Journal 100, 49-59.

[4] Almekinders, G.J. (1995) Foreign Exchange Interventions: Theory and Evidence. Edward Elgar, Aldershot, UK.

[5] Avery, Ch. and Zemsky, P. (1998) Multidimensional uncertainty and herd behavior in financial markets. American Economic Review 88, 724-748.

[6] Baillie, R.T. and Osterberg, W.P. (1997) Why do central banks intervene? Journal of International Money and Finance 16, 909-919.

[7] Bank of England (1983) Intervention, stabilisation and profits. Quarterly Bulletin 23, 384-391.

[8] Baumol, W. (1957) Speculation, profitability, and stability. Review of Economics and Statistics 39, 263-271.

[9] Baxter, M. and Stockman, A.C. (1989) Business cycles and the exchange rate regime: Some international evidence. Journal of Monetary Economics 23, 377-400.

[10] Berndt, E., Hall, B. Hall, R. and Hausman, J. (1974) Estimation and inference in nonlinear structural models. Annals of Economic and Social Measurement 3, 653-665.

[11] Bertola, G. and Caballero, R.J. (1992) Target zones and realignments. American Economic Review 82, 520-536.

[12] Bertola, G. and Svensson, L.E.O. (1993) Stochastic devaluation risk and the empirical fit of target zone models. Review of Economic Studies 60, 689-712.

[13] Bhattacharya, U. and Weller, P. (1997) The advantage of hiding one's hand: Speculation and central bank intervention in the foreign exchange market. Journal of Monetary Economics 39, 251-277.

[14] BIS - Bank for International Settlements (1996), Central Bank Survey of Foreign Exchange and Derivatives Market Activity 1995, Basle.

[15] Black, F. (1986) Noise. Journal of Finance XLI, 529-543.

[16] Blake, D., Beenstock, M. & Brasse, V. (1986) The performance of UK exchange rate forecasters. Economic Journal 96, 986-999.

[17] Bollerslev, T (1986) Generalized autoregressive conditional heteroskedasticity. Journal of Econometrics 31, 307-327.

[18] Bollerslev, T. and Wooldridge, J. (1992) Quasi-maximum likelihood estimation and inference in dynamic models with time varying covariances. Econometric Reviews 11, 143-172.

[19] Bonser-Neal, C. and Tanner, G. (1996) Central bank intervention and the volatility of foreign exchange rates: Evidence from the options market. Journal of International Money and Finance 15, 853-878.

[20] Brock, W., Lakonishok, J. & LeBaron, B. (1992) Simple technical trading rules and the stochastic properties of stock returns. Journal of Finance 47, 1731-64.

[21] Bundesbank (1998) 50 Jahre Deutsche Mark: Monetäre Statistiken 1948-1997 auf CD-ROM, Vahlen, Munich.

[22] Catte, P., Galli, G. and Rebecchini, S. (1994) Concertated interventions and the dollar: An analysis of daily data. In The International Monetary System, eds P. Kenen, F. Papadia and S. Saccomanni, 201-239, Cambridge University Press, Cambridge, UK.

[23] Cavaglia, S., Verschoor, W.F.C. & Wolff, C.C.P. (1993) Further evidence on exchange rate expectations. Journal of International Money and Finance 12, 78-98.

[24] Chang, P.H.K. and Osler, C.L. (1998) Is more always better? Head-and-shoulders and filter rules in foreign exchange markets. In Advanced Trading Rules, eds E. Acar and S. Satchell. Butterworth Heinemann.

[25] Chang, Y. and Taylor, S.J. (1998) Intraday effects of foreign exchange intervention by the Bank of Japan. Journal of International Money and Finance 17, 191-210.

[26] Cornell, W. B. and Dietrich, J. K. (1978) The efficiency of the market for foreign exchange under floating exchange rates. Review of Economics and Statistics 60, 111-120.

[27] Cornett, M.M., Schwarz, T.V. & Szakmary, A.C. (1995) Seasonalities and intraday return patterns in the foreign exchange market. Journal of Banking and Finance 19, 843-869.

[28] Corrado, C.J. and Taylor, D. (1986) The cost of a central bank leaning against a random walk. Journal of International Money and Finance 5, 303-314.

[29] Dacorogna, M.M., Müller, U.A. and Pictet, O.V. (1991) A measure of trading model performance with a risk component. Olsen and Associates Research Group Discussion Paper October 1991.

[30] De Grauwe, P. and Dewachter, H. (1992) Chaos in the Dornbusch model of the exchange rate. Kredit und Kapital, 27-54.

[31] De Grauwe, P. Dewachter, H. & Embrechts, M. (1993) Exchange Rate Theory: Chaotic Models of Foreign Exchange Markets. Basil Blackwell, Oxford, UK.

[32] DeLong, J.B., Shleifer, A., Summers, L.H. & Waldman, R.J. (1990) Noise trader risk in financial markets. Journal of Political Economy 98, 703-738.

[33] Dewachter, H. (1997) Can Markov-switching models replicate chartist profits in the foreign exchange market? Center for Economic Studies Discussion Paper Series DPS 97.27.

[34] Dominguez, K.M. (1990) Market responses to coordinated central bank intervention. Carnegie Rochester Series on Public Policy 32, 121-164.

[35] Dominguez, K.M. and Frankel, F.A. (1993) Does Foreign Exchange Intervention Work? Institute for International Economics. Washington D.C.

[36] Dominguez, K.M. (1998) Central bank intervention and exchange rate volatility. Journal of International Money and Finance 17, 161-190.

[37] Dornbusch, R. (1976) Expectations and exchange rate dynamics. Journal of Political Economy 84, 1161-1176.

[38] Dornbusch, R. and Frankel, J.A. (1987) The flexible exchange rate system: Experience and alternatives. NBER Working Paper # 2464.

[39] Edwards, R.D. and Magee, J. (1997) Technical analysis of stock trends. 7th Ed.. John Magee Inc. Chicago, Illinois.

[40] Efron, B. and Tibshirani, R.J. (1993) An Introduction to the Bootstrap. Chapman & Hall, New York, NY.

[41] Eijffinger, S.C.W. and Verhagen, W.H. (1997) The advantage of hiding both hands: Foreign exchange intervention, ambiguity and private information. Center for Economic Research Discussion Paper # 9730, Tilburg University.

[42] Engel, Ch. and Hamilton, J.D. (1990) Long swings in the dollar: Are they in the data and do the markets know it? American Economic Review 80, 689-713.

[43] Engel, Ch. (1994) Can Markov switching models forecast exchange rates? Journal of International Economics 36, 151-165.

[44] Engle, R.F. (1982) Autoregressive conditional heteroskedasticity with estimates of the variance of United Kingdom inflation. Econometrica 50, 987-1007.

[45] Fase, M.M.G. and Huijser, A.P. (1991) Foreign exchange rate stabilisation and the profitability of official market intervention: A case study for the Netherlands 1974-1989. Kredit und Kapital 24, 212-234.

[46] Flood, R.A. and Rose, A.K. (1995) Fixing exchange rates: A virtual quest for fundamentals. Journal of Monetary Economics 36, 3-37.

[47] Flood, R.P. and Rose, A.K. (1996) Fixes: Of the forward discount puzzle. Review of Economics and Statistics , 748-752.

[48] Flood, R.P. and Rose, A.K. (1999) Understanding exchange rate volatility without the contrivance of macroeconomics. CEPR Working Paper # 1944 (forthcoming in the Economic Journal).

[49] Frankel, J.A. (1990) The making of exchange rate policy in the 1980s. NBER Working Paper # 3539.

[50] Frankel, J.A. (1996) How well do foreign exchange markets work: Might a Tobin tax help? NBER Working Paper # 5422.

[51] Frankel, J.A. and Froot, K.A. (1986) Understanding the US dollar in the Eighties: The expectations of chartists and fundamentalists. Economic Record 62, 24-38.

[52] Frankel, J.A. and Froot, K (1987) Using survey data to test standard propositions regarding exchange rate expectations. American Economic Review 77, 133-153.

[53] Frankel, J.A. and Froot, K (1990) Chartists, fundamentalists, and the demand for dollars. In Policy Issues for Interdependent Economies, eds T. Courakis. and M. Taylor, 73-126, Macmillan, London.

[54] Frankel, J.A. and Rose, A.K. (1995) Empirical research on nominal exchange rates. In Handbook of International Economics,Vol. III, G. M. Grossman and K. Rogoff (Eds.), 1689-1730. Elsevier, Amsterdam.

[55] Frenkel, J.A. (1976) A monetary approach to the exchange rate: Doctrinal aspects and empirical evidence. Scandinavian Journal of Economics 78, 200-224.

[56] Frenkel, J.A. and Mussa, M.L. (1980) The efficiency of foreign exchange markets and measures of turbulence. American Economic Association Papers and Proceedings 70, 374-381.

[57] Friedman, M. (1953) The case for flexible exchange rates. In Essays in Positive Economics. University of Chicago Press, Chicago.

[58] Froot, K.A. and Ito, T. (1989) On the consistency of short-run and long-run exchange rate expectations. Journal of International Money and Finance 8, 487-510.

[59] Funabashi, Y. (1988) Managing the Dollar: From the Plaza to the Louvre, Institute for International Economics, Washington DC.

[60] Gençay, R. (1999) Linear, non-linear and essential foreign exchange prediction with simple technical trading rules. Journal of International Economics 47, 91-109.

[61] Ghosh (1992) Is it signalling? Exchange intervention and the dollar-Deutschemark rate. Journal of International Economics 32, 201-220.

[62] Gleske, L. (1982) Die Devisenpolitik der Deutschen Bundesbank: Interventionen am DM-$-Markt und im Europäischen Währungssystem sowie geldmarktorientierte Devisentransaktionen. Kredit und Kapital, 15, 259-274.

[63] Goodhart, C.A.E. and Hesse, T. (1993) Central bank forex interventions assessed in continuous time. Journal of International Money and Finance 12, 368-389.

[64] Group of Thirty (1985) The Foreign Exchange Market in the 80's, New York.

[65] Hamilton, J.D. (1989) A new approach to the economic analysis of non-stationary time series and the business cycle. Econometrica 57, 357-384.

[66] Hamilton, J.D. (1994) Time Series Analysis. Princeton University Press.

[67] Hart, O.D. and Kreps, D.M. (1986) Price destabilizing speculation. Journal of Political Economy 94, 927-954.

[68] Hung, J.H. (1997) Intervention strategies and exchange rate volatility: A noise trading perspective. Journal of International Money and Finance 16, 779-793.

[69] Ito, T. (1990) Foreign exchange rate expectations: Micro survey data. American Economic Review 80, 434-449.

[70] Jacobson, L.R. (1983) Calculations of profitability for US dollar-deutschemark intervention. Staff Series # 131, Board of Governors of the Federal Reserve System. Washington, DC.

[71] Jeanne, O. and Rose, A.K. (1999) Noise trading and exchange rate regimes. NBER Working Paper # 7104.

[72] Kaufman, P.J. (1978) Commodity trading systems and methods. John Wiley, New York.

[73] Kaminsky, G.L. and Lewis, K.K. (1996) Does foreign exchange intervention signal future monetary policy? Journal of Monetary Economics 37, 285-312.

[74] Keim, D.B. and Stambaugh, R. (1984) A further investigation of the weekend effect in stock returns, Journal of Finance 39, 819-840.

[75] Kho, B.-Ch. (1996) Time-varying risk premia, volatility, and technical trading profits: Evidence from foreign currency futures markets. Journal of Financial Economics 41, 249-290.

[76] Kirman, A. (1993) Ants, rationality and recruitment. Quarterly Journal of Economics 108, 137-156.

[77] Klein, M.W. (1993) The accuracy of reports of foreign exchange intervention. Journal of International Money and Finance 12, 644-653.

[78] Krugman, P. (1991) Target zones and exchange rate dynamics. Quarterly Journal of Economics 106, 669-682.

[79] Krugman, P. and Miller, M.H. (1992) Why have a target zone? CEPR Discussion Paper #718.

[80] Leahy, M.P. (1995) The profitability of US intervention in the foreign exchange markets. Journal of International Money and Finance 14, 823-44.

[81] LeBaron, B. (1991) Technical trading rules and regime shifts in foreign exchange markets. Working Paper, University of Wisconsin - Madison; republished in: Advanced Trading Rules, eds E. Acar and S. Satchell. Butterworth Heinemann, 1998.

[82] LeBaron, B. (1996) Technical trading rule profitability and foreign exchange intervention. NBER Working Paper 5505; forthcoming in Journal of International Economics.

[83] Lee, Ch.I. and Mathur, I. (1996) Trading rule profits in European currency spot cross-rates. Journal of Banking and Finance 20, 949-962.

[84] Levich, R. and Thomas, L. (1993) The significance of technical trading rule profitability: A bootstrap approach. Journal of International Money and Finance 12, 451-474.

[85] Lewis, K.K. (1995a) Are foreign exchange intervention and monetary policy related, and does it really matter? Journal of Business 86, 185-214.

[86] Lewis, K.K. (1995b) Puzzles in International Financial Markets. In Handbook of International Economics,Vol. III, G. M. Grossman and K. Rogoff (Eds.), 1913-1971. Elsevier, Amsterdam.

[87] Luca, C. (1997) Technical analysis applications in the global currency markets. Prentice Hall. New York Institute of Finance. New York.

[88] Lui, Y.-H. and Mole, D. (1998) The use of fundamental and technical analysis by foreign exchange dealers: Hong Kong evidence. Journal of International Money and Finance 17, 535-545.

[89] Lukac, L.P. Brorsen, B.W. and Irwin, S.H. (1988) A test of futures market disequilibrium using twelve different technical trading systems. Applied Economics 20, 623-639.

[90] Lukas, R.E. (1987) Models of Business Cycles. Basil Blackwell, Oxford.

[91] Lumsdaine, R. (1995) Finite-sample properties of the maximum likelihood estimator in GARCH(1,1) and IGARCH(1,1) models: A Monte Carlo investigation. Journal of Business and Economic Statistics 13, 1-10.

[92] Lux, Th. (1995) Herd behaviour, bubbles and crashes. Economic Journal 105, 881-896.

[93] Lux, Th. (1998) The socio-economic dynamics of speculative markets: Interacting agents, chaos, and the fat tails of return distributions. Journal of Economic Behavior and Organization 33, 143-165.

[94] Lyons, R.K. (1995) Foreign exchange volume: Sound and fury signifying nothing? NBER Working Paper # 4984.

[95] Lyons, R.K. (1997) A simultaneous trade model of the foreign exchange hot potato. Journal of International Economics 42, 275-298.

[96] Malkiel, B. (1990). A Random Walk down Wallstreet. 5th ed..Norton, New York.

[97] Mark, N.C. (1995) Exchange rates and fundamentals: Evidence on long-horizon predictability. American Economic Review 85, 201-218.

[98] Mark, N.C. and Wu, Y. (1998) Rethinking deviations from uncovered interest parity: The role of covariance risk and noise. Economic Journal 108, 1686-1706.

[99] Meese, R.A. and Rogoff, K (1983) Empirical exchange rate models of the seventies: Do they fit out of sample? Journal of International Economics 14, 3-24.

[100] Menkhoff, L. and Schlumberger, M. (1995) Persistent profitability of technical analysis on foreign exchange markets?. Banca Nazionale Di Lavoro Quarterly Review 193, 189-216.

[101] Menkhoff, L. (1997) Examining the use of technical currency analysis. International Journal of Finance and Economics 2, 307-318.

[102] Menkhoff, L. (1998) The noise trader approach - questionnaire evidence from foreign exchange. Journal of International Money and Finance 17, 547-564.

[103] Murphy, J.J. (1986) Technical analysis of the futures markets. Prentice Hall. New York Institute of Finance. New York.

[104] Mussa, M.M. (1986) Nominal exchange rate regimes and the behavior of the real exchange rate. Carnegie Rochester Series on Public Policy 11, 9-58.

[105] Neely, Ch., Dittmar, R. and Weller, P. (1996) Is technical analysis in the foreign exchange market profitable? A genetic programming approach. CEPR Discussion Paper 1480.

[106] Neely, Ch. and Weller, P. (1997) Technical analysis and central bank intervention. Federal Reserve Bank of St. Louis Working Paper 97-002A.

[107] Neely, Ch. and Weller, P. (1998) Technical trading rules in the European Monetary System. Federal Reserve Bank of St.Louis Working Paper.

[108] Newey, W.K. and West , K.D. (1987) A simple positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix. Econometrica 55, 703-708.

[109] Osler C.L. and Chang P.H.K. (1995) Head-and-shoulders: Not just a flaky pattern. Federal Reserve Bank of New York Staff Reports 4, August 1995.

[110] Pring, M.J. (1991) Technical analysis explained. 3rd. Ed.. McGraw-Hill. New York.

[111] Rose, A.K. and Svensson, L.E.O. (1994) European exchange rate credibility before the fall. European Economic Review 38, 1185-1216.

[112] Scharfstein, D.S. and Stein, J.C. (1990) Herd behaviour and investment. American Economic Review 80, 465-479.

[113] Schulmeister, S. (1987) An essay on exchange rate dynamics. Discussion paper IIM/LMP 87-8. Wissenschaftszentrum Berlin für Sozialforschung.

[114] Schwarz, G. (1978) Estimating the dimension of a model. Annals of Statistic 6, 461-464.

[115] Shleifer, A. and Summers, L.H. (1990) The noise trader approach to finance. Journal of Economic Perspectives 4, 19-33.

[116] Sjöö, B. and Sweeney, R.J. (1999) The foreign exchange cost of central bank intervention: Evidence from Sweden. Forthcoming in: Journal of International Money and Finance.

[117] Spencer, P.D. Official intervention in the foreign exchange market. (1985) Journal of Political Economy 93, 1019-1024.

[118] Spencer, P.D. How to make the central bank look good: A reply. (1989) Journal of Political Economy 97, 233-235.

[119] Surajaras, P. and Sweeney, R.J. (1992) Profit-making speculation in foreign exchange markets. Westview Press, Boulder, Co.

[120] Sweeney, R. J. (1986). Beating the Foreign Exchange Market. Journal of Finance 41, 163-182.

[121] Sweeney, R. J. (1997) Do central banks lose on foreign-exchange intervention? A review article. Journal of Banking and Finance 21, 1667-1684.

[122] Sweeney, R.J. (1999) Does the Fed beat the foreign exchange market? Forthcoming in: Journal of Banking and Finance.

[123] 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, 513-535.

[124] Szpiro, G.G. (1994) Exchange rate speculation and chaos inducing speculation. Journal of Economic Behaviour and Organization 24, 363-368.

[125] Taylor, D. (1982) Official intervention in the foreign exchange market, or bet against the central bank. Journal of Political Economy 90, 356-368.

[126] Taylor, D. (1989) How to make central banks look good. Journal of Political Economy 97, 226-232.

[127] Taylor, M. P. and Allen, H. (1992) The use of technical analysis in the foreign exchange market. Journal of International Money and Finance 11, 304-314.

[128] Taylor, S.J. (1992) Rewards available to currency futures speculators: Compensation for risk of evidence of inefficient pricing? Economic Record 68, 105-116.

[129] Tobin, J. (1978) A proposal for international monetary reform. Eastern Economic Journal 4, 153-159.

[130] Truman, E. (1992) Comments on 'Concerted Interventions and the Dollar: An Analysis of Daily Data'. Ossola Memorial Conference, Banca d'Italia, Perugia, Italy.

[131] Vigfusson, R. (1997) Switching between chartists and fundamentalists: A Markov regime-switching approach. International Journal of Finance and Economics 2, 291-305.

[132] Weber, A.A. (1994) Foreign exchange intervention and international policy coordination: Comparing the G-3 and EMS experience. CEPR Discussion Paper # 1038.

[133] White , H. (1980). A heteroskedasticity-consistent covariance matrix estimator and a direct test for heteroskedasticity. Econometrica, 48, 817-838.

Appendix A

Omitted Graphs

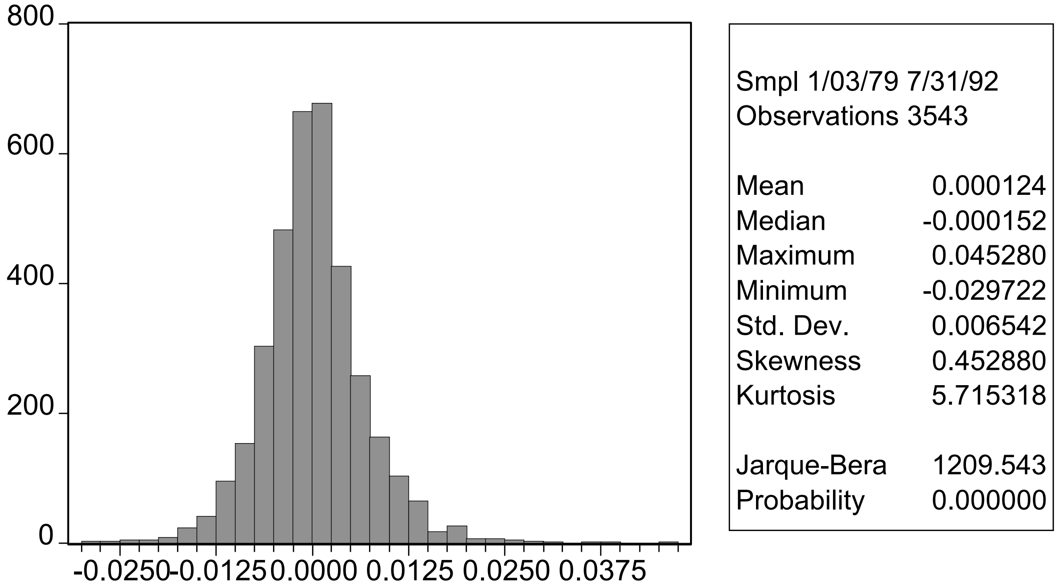

Frequency Distribution and Summary Statistics of Daily JPY/USD Exchange Rate Changes

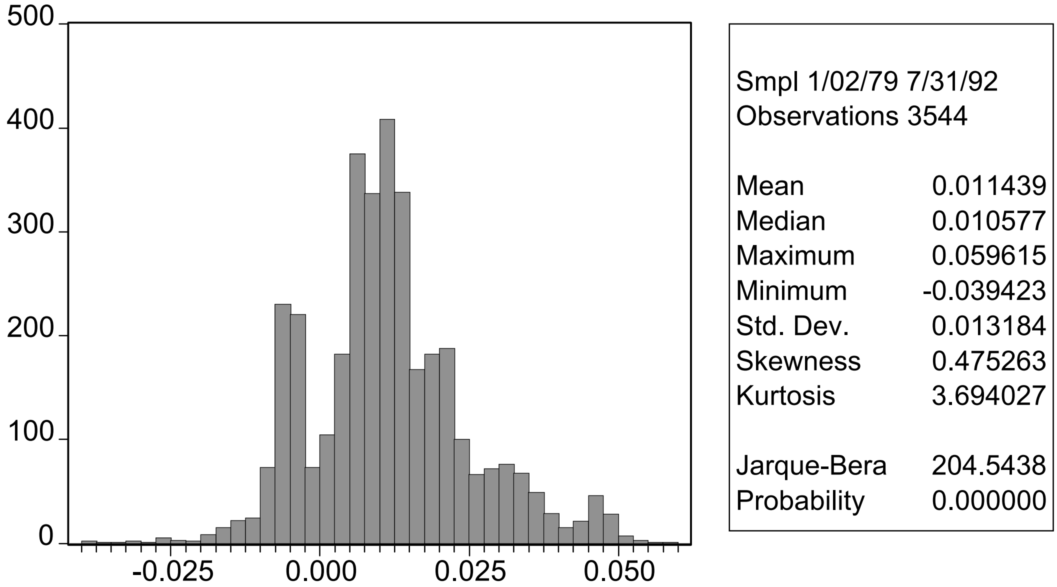

Frequency Distribution and Summary Statistics of Daily USD-JPY Interest Differentials

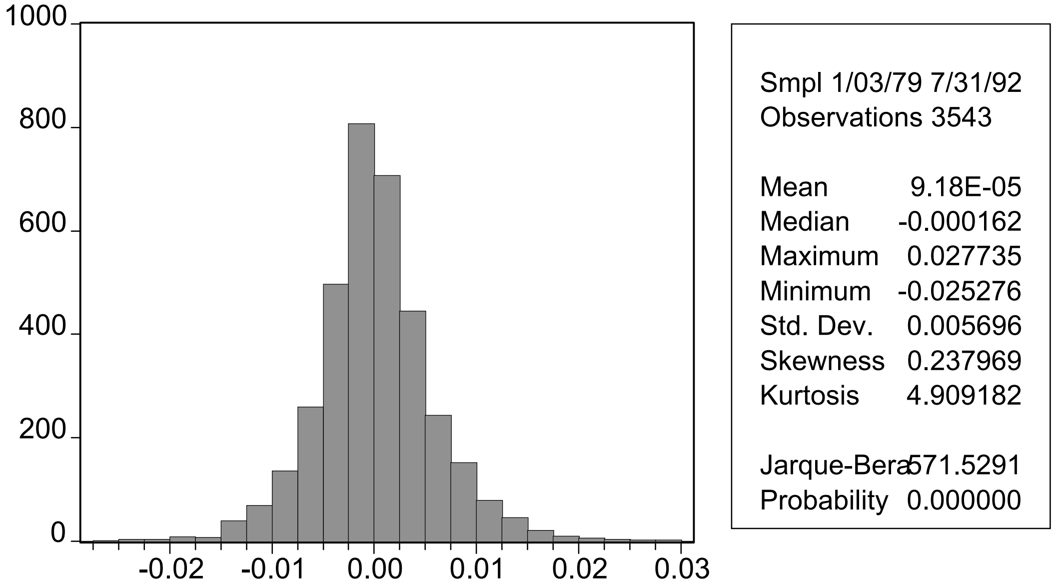

Frequency Distribution and Summary Statistics of Daily JPY/DEM Exchange Rate Changes

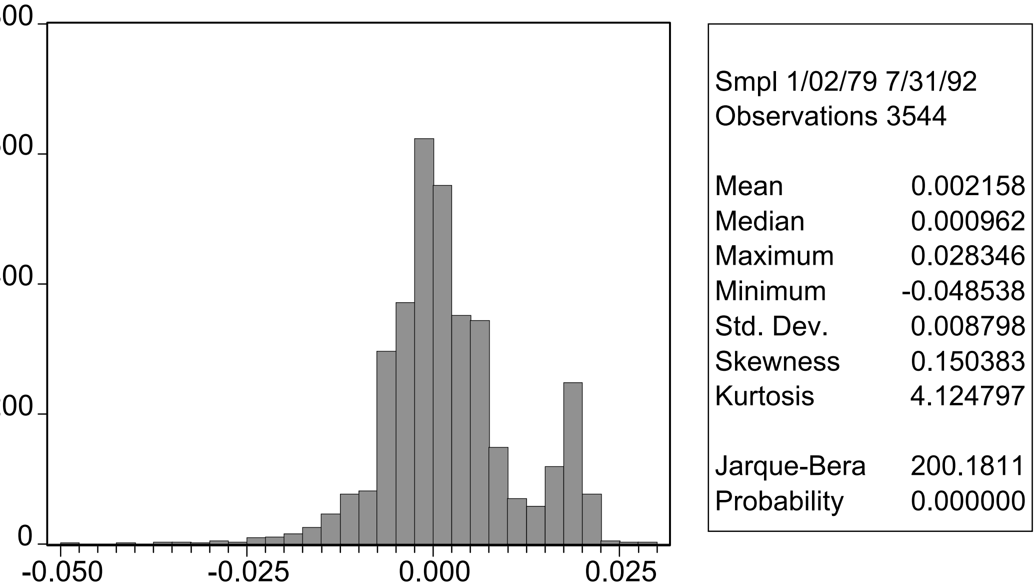

Frequency Distribution and Summary Statistics of Daily DEM-JPY Interest Differentials

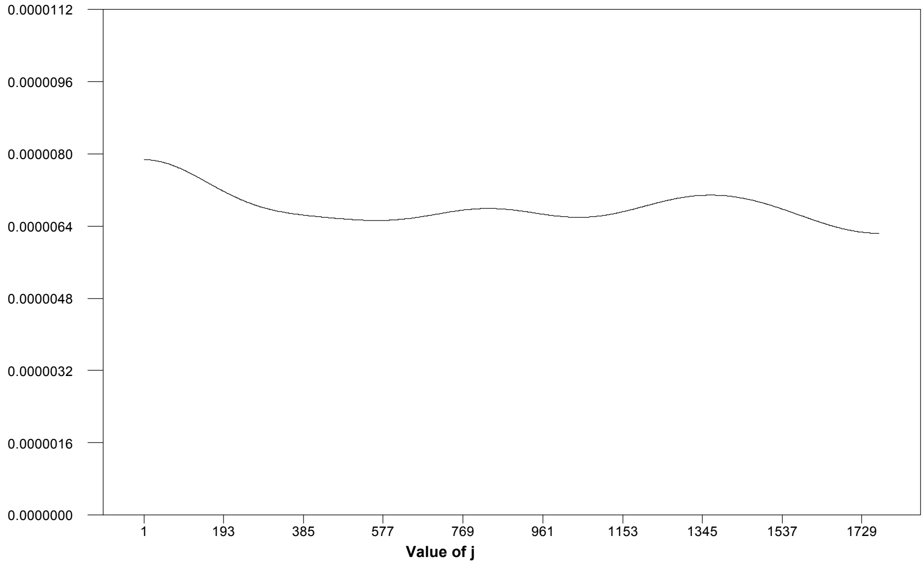

Spectrum of Daily JPY/USD Exchange Rate Changes

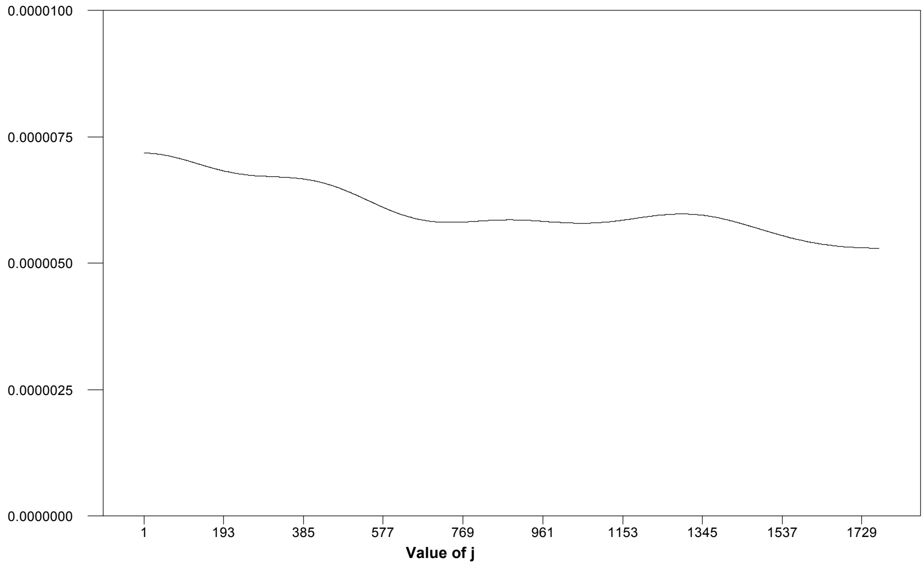

Spectrum of Daily JPY/DEM Exchange Rate Changes

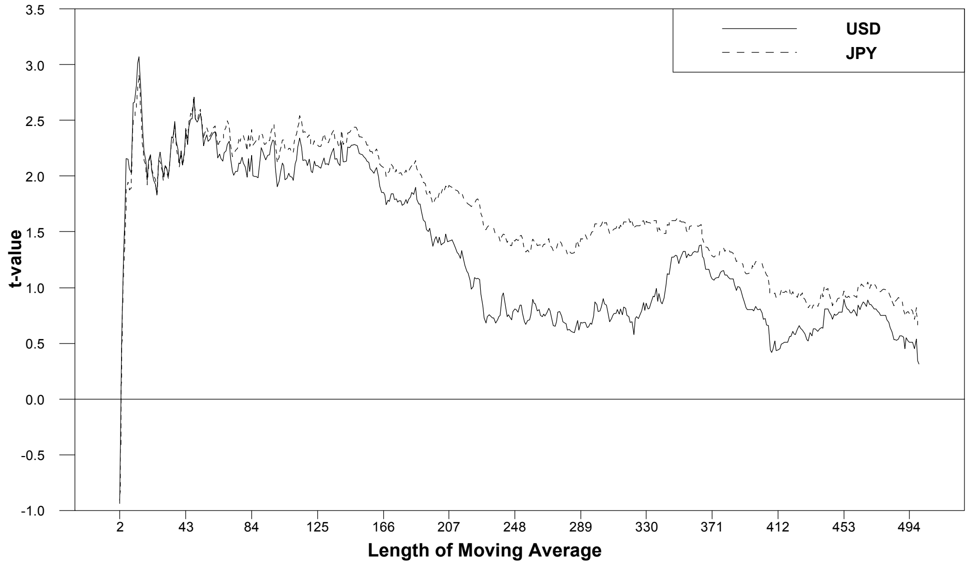

Results of Sweeney's (1986) X-test of the Significance of Risk-adjusted MA Trading Rule Returns for the JPY/USD Exchange Rate

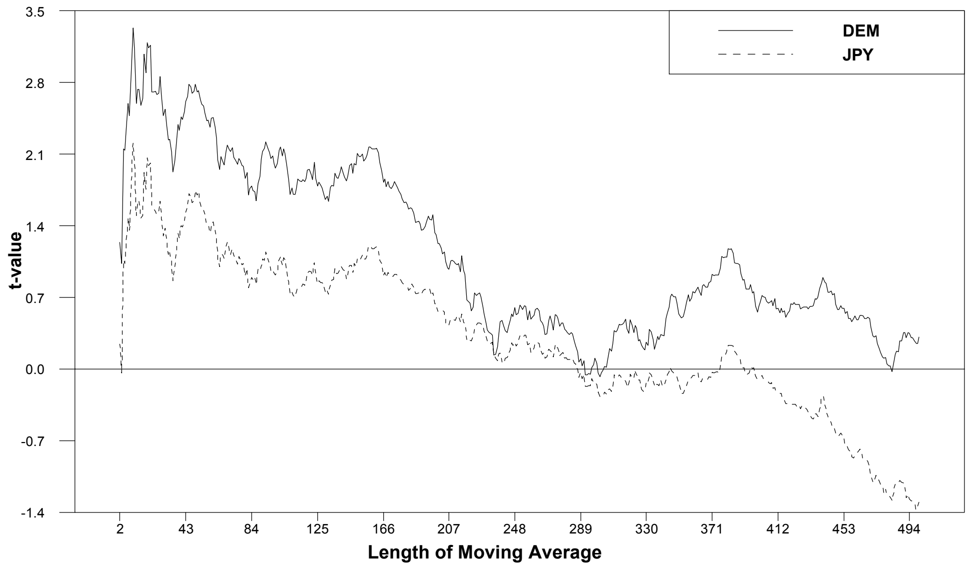

Results of Sweeney's (1986) X-test of the Significance of Risk-adjusted MA Trading Rule Returns for the JPY/DEM Exchange Rate

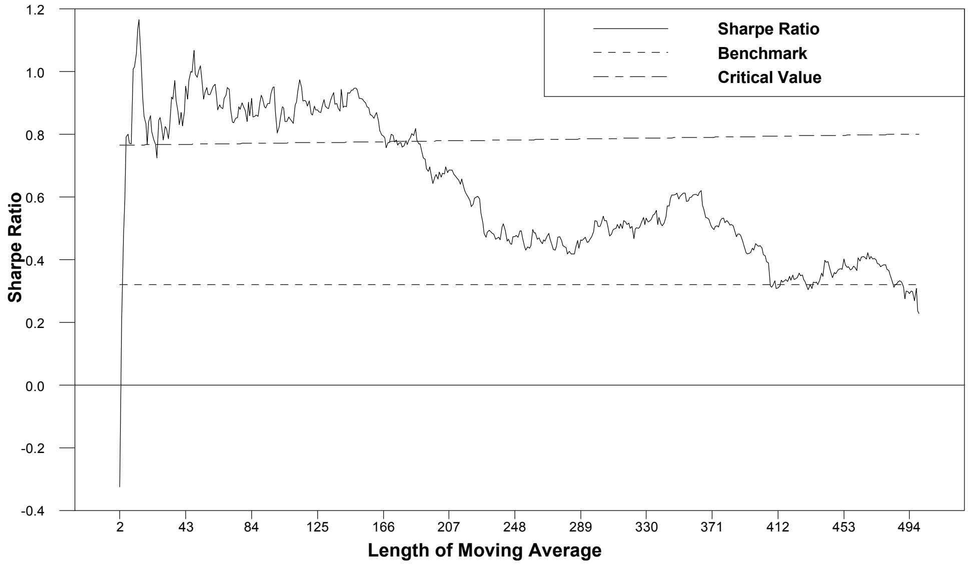

Sharpe Ratio of Moving Average Trading Rule Returns for the JPY/USD Exchange Rate

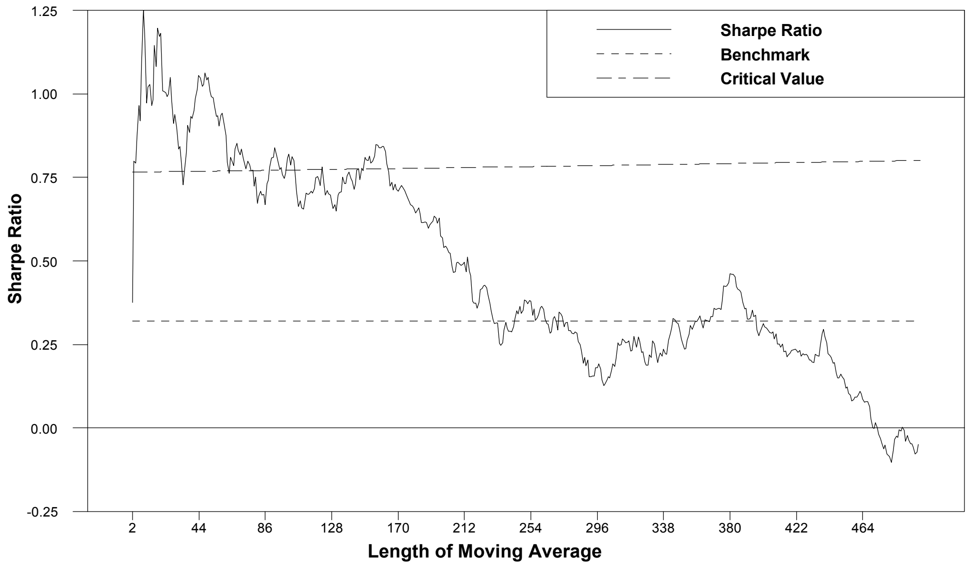

Sharpe Ratio of Moving Average Trading Rule Returns for the JPY/DEM Exchange Rate

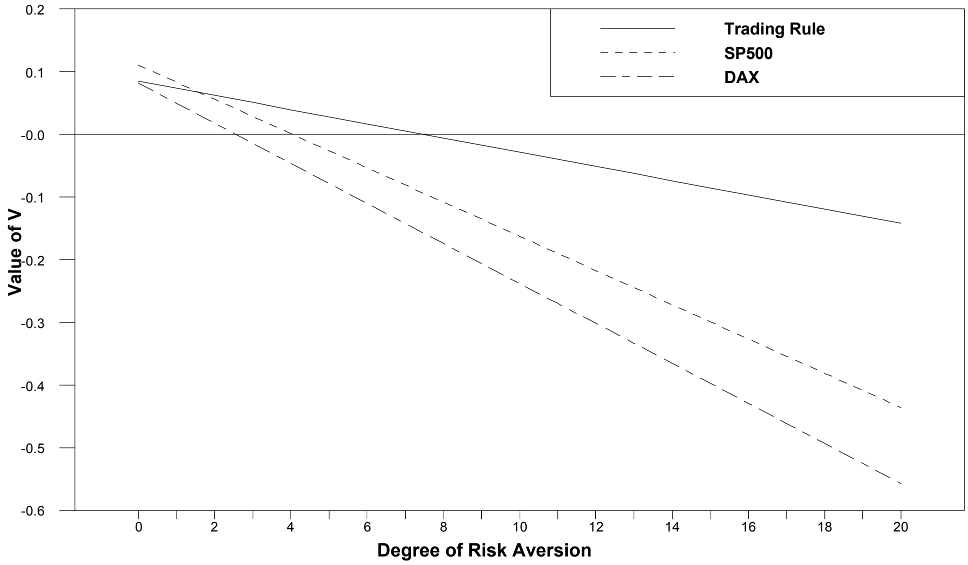

Expected Utility Comparison between the Peak-and-Trough-Progression Rule and the SP500 and DAX indices for the JPY/USD Exchange Rate

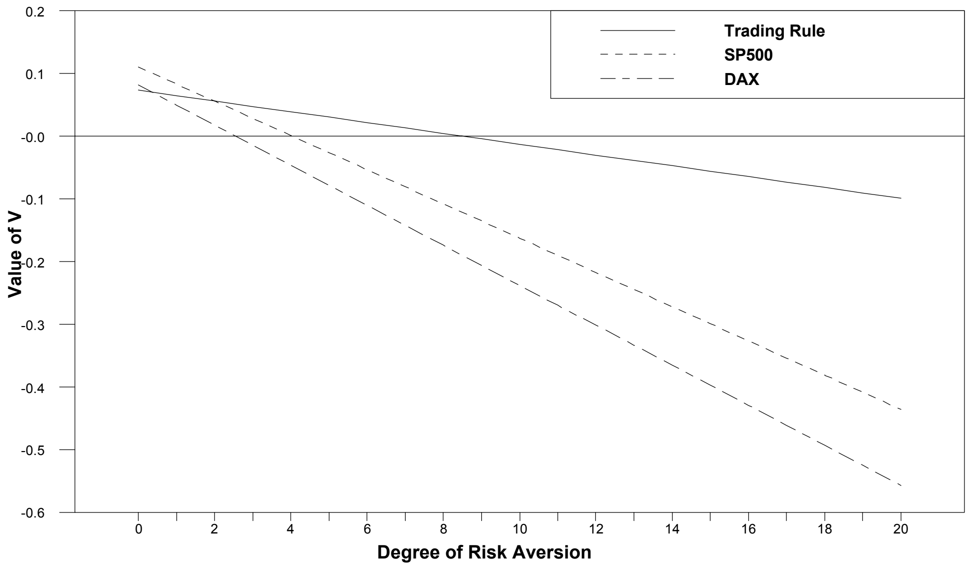

Expected Utility Comparison between the Peak-and-Trough-Progression Rule and the SP500 and DAX indices for the JPY/DEM Exchange Rate

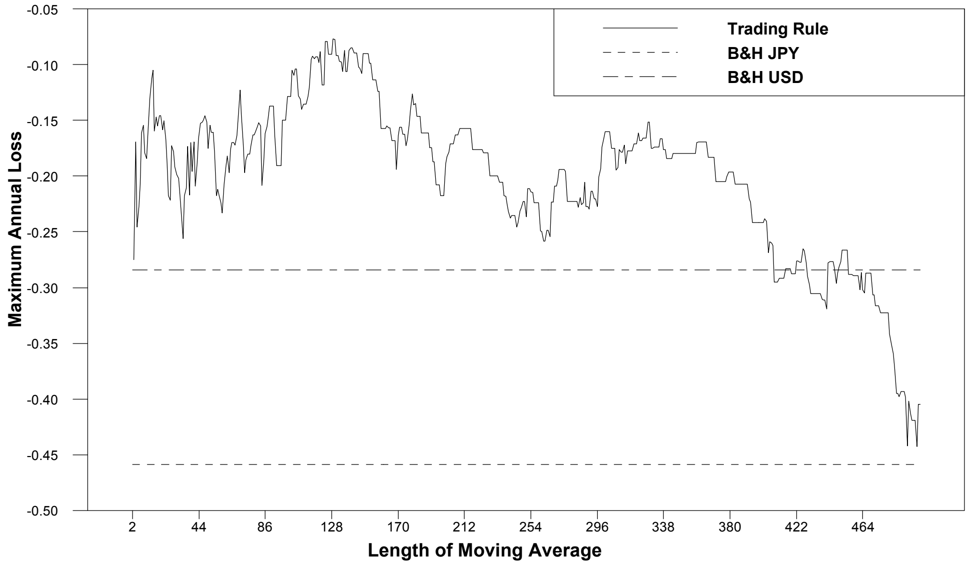

Maximal Annual Loss of Moving Average Trading Rules for the JPY/USD Exchange Rate

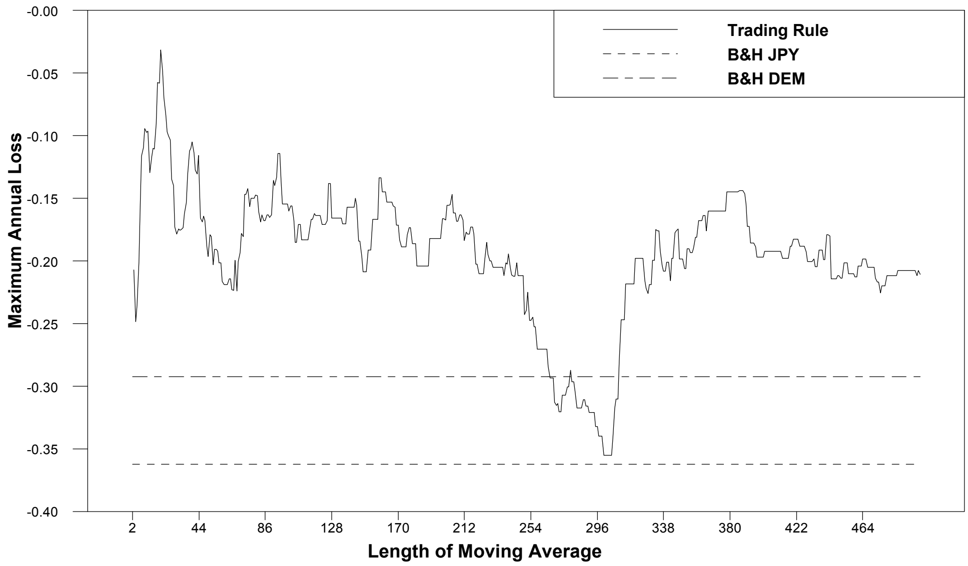

Maximal Annual Loss of Moving Average Trading Rules for the JPY/DEM Exchange Rate

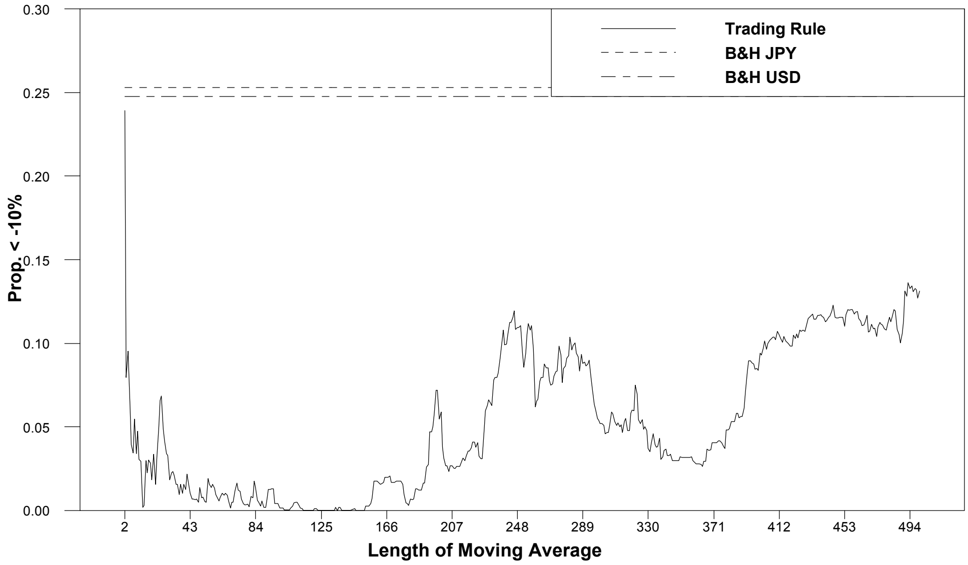

Proportion of One Year Periods over which Moving Average Rule Returns are Smaller than -10% for the JPY/USD Exchange Rate

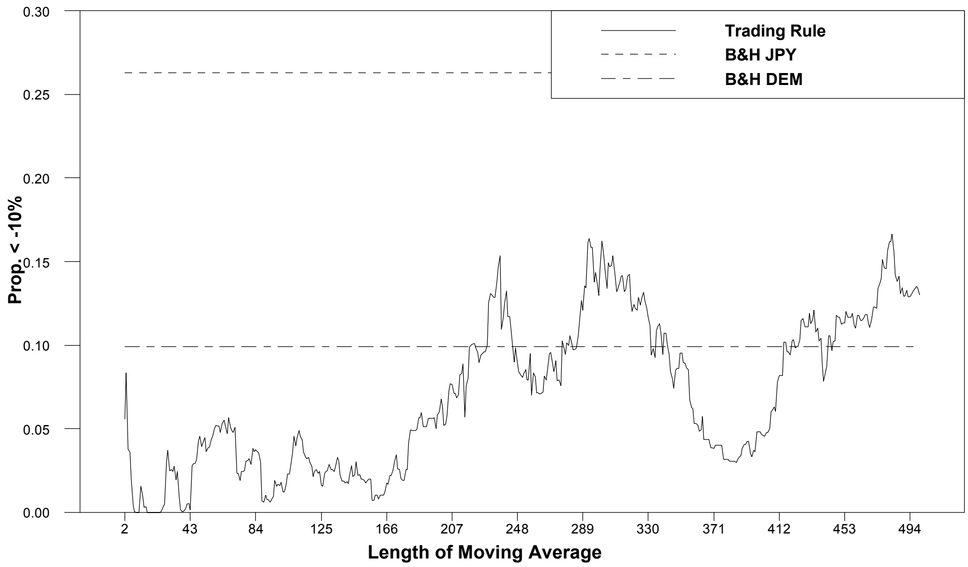

Proportion of One Year Periods over which Moving Average Rule Returns are Smaller than -10% for the JPY/DEM Exchange Rate

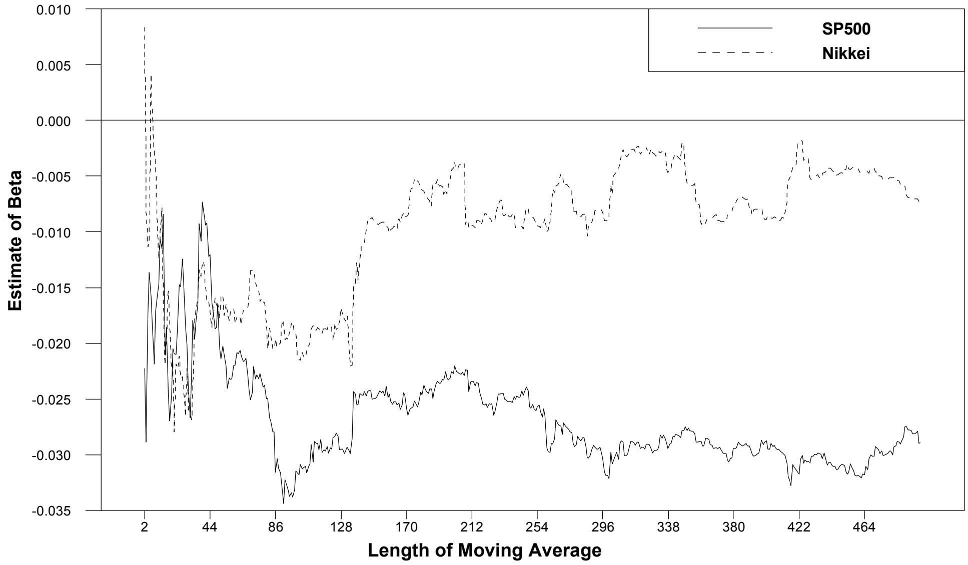

Estimates of CAPM Betas for Moving Average Trading Rule Returns for the JPY/USD Exchange Rate

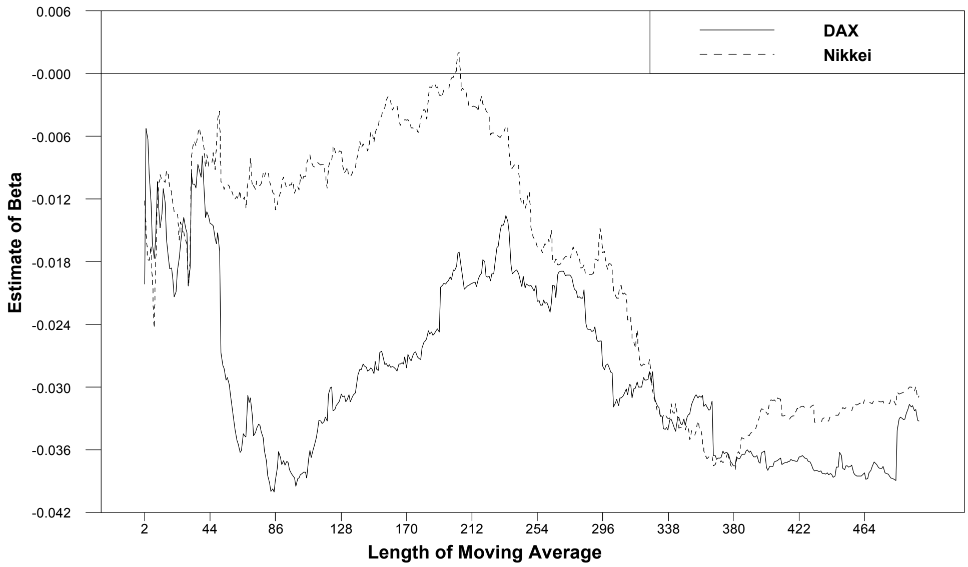

Estimates of CAPM Betas for Moving Average Trading Rule Returns for the JPY/DEM Exchange Rate

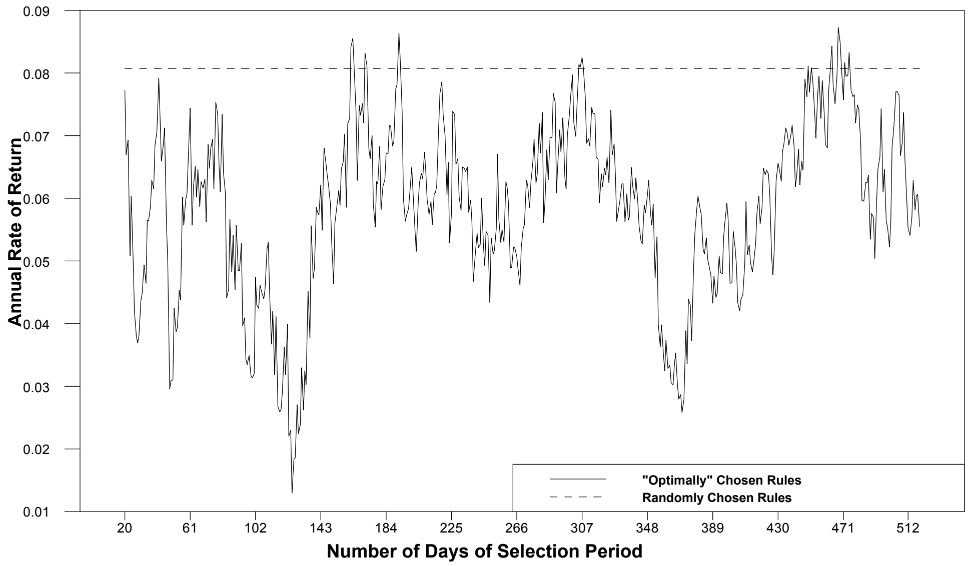

Profitability of Always Choosing the Best Rule for the JPY/USD Exchange Rate

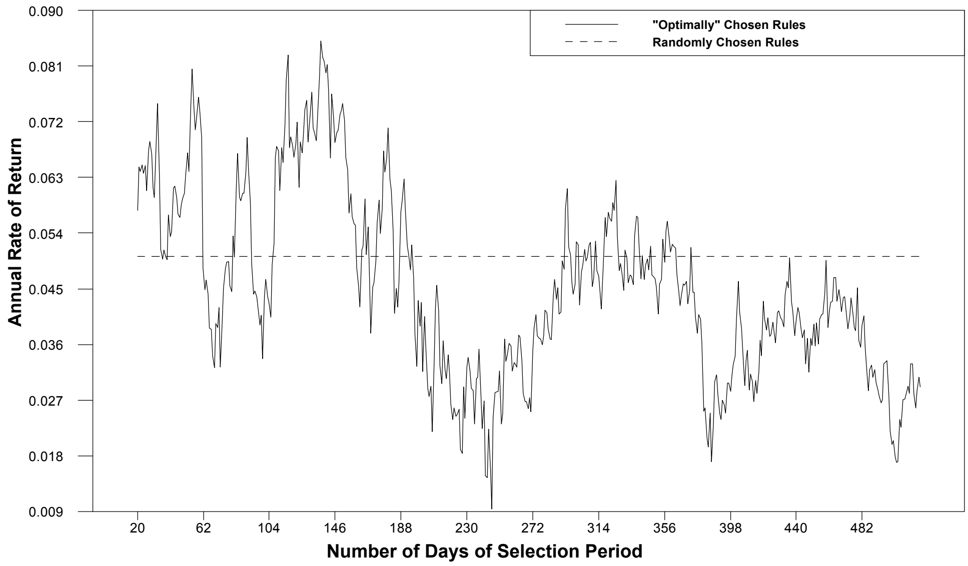

Profitability of Always Choosing the Best Rule for the JPY/DEM Exchange Rate

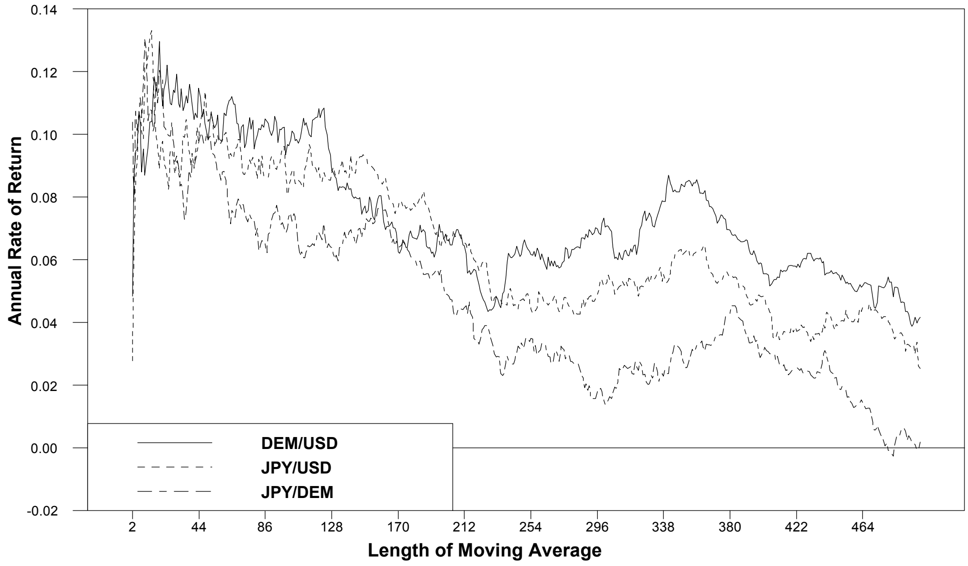

Gross Rates of Return from Following MA Trading Rules for Flexible Exchange Rates

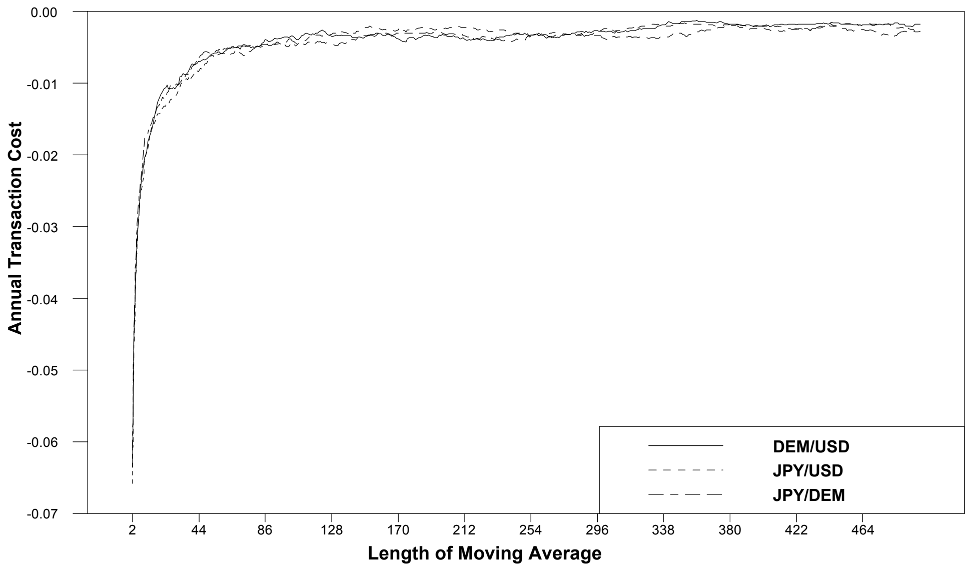

Average Annual Transaction Costs Incurred when Following MA Trading Rules for Flexible Exchange Rates

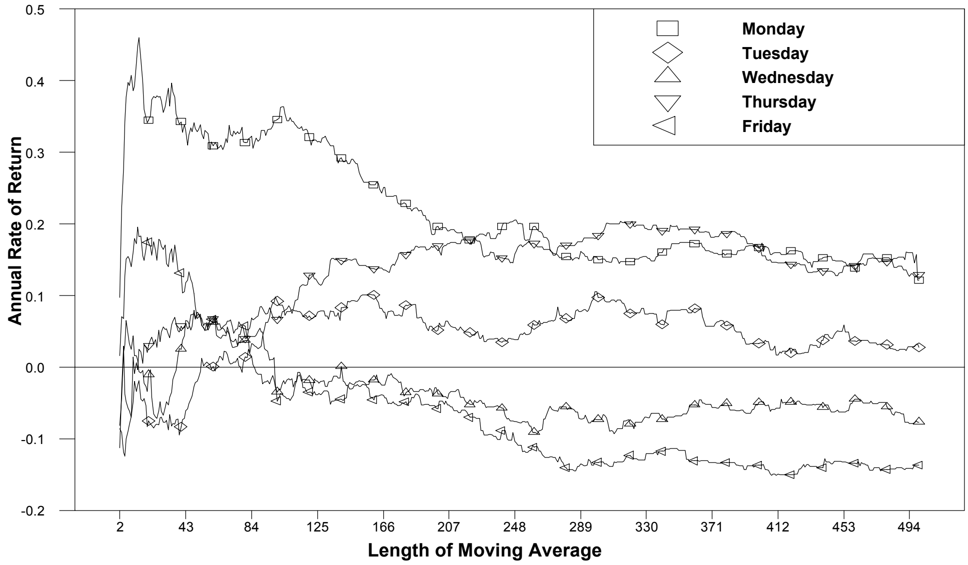

Moving Average Trading Rule Returns for Different Days of the Week for the JPY/USD Exchange Rate

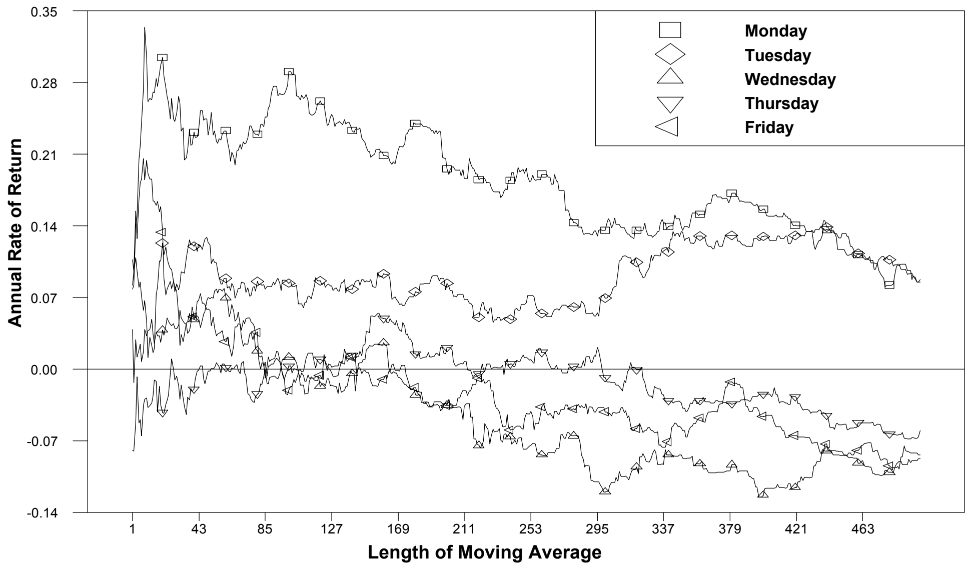

Moving Average Trading Rule Returns for Different Days of the Week for the JPY/DEM Exchange Rate

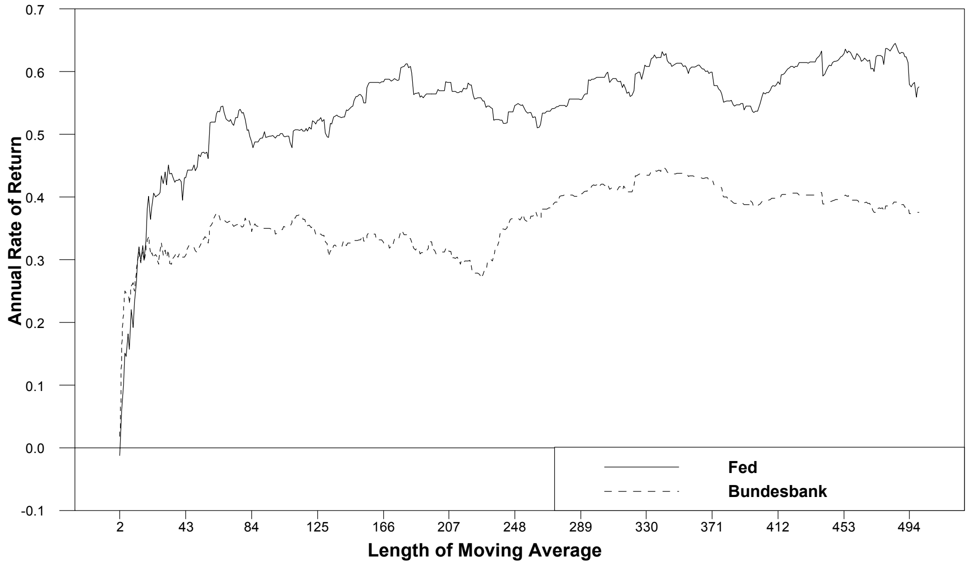

Annual Rates of Return of MA Trading Rules on Days when the Fed / the Bundesbank Intervened

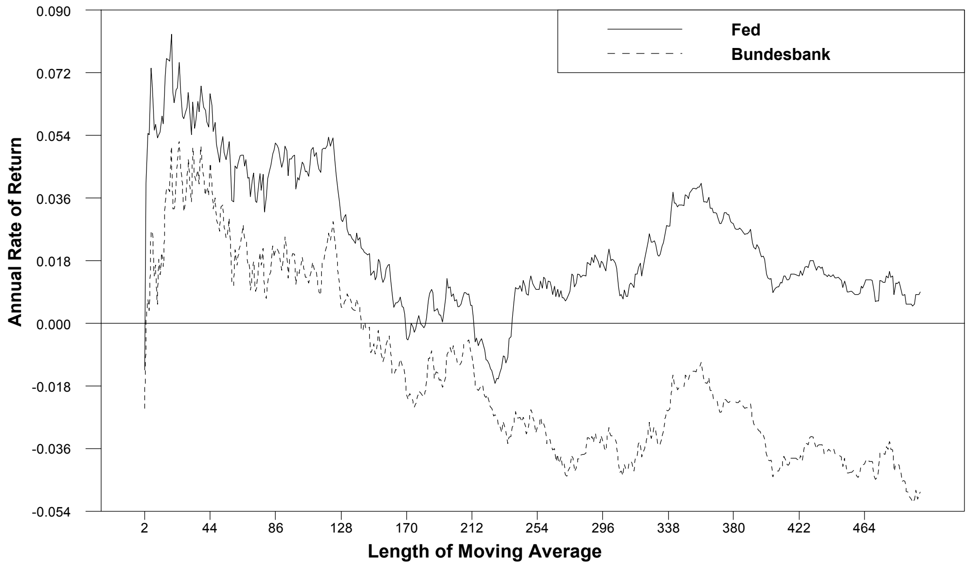

Annual Rate of Return of MA Trading Rules on Days when the Fed / the Bundesbank did not Intervene

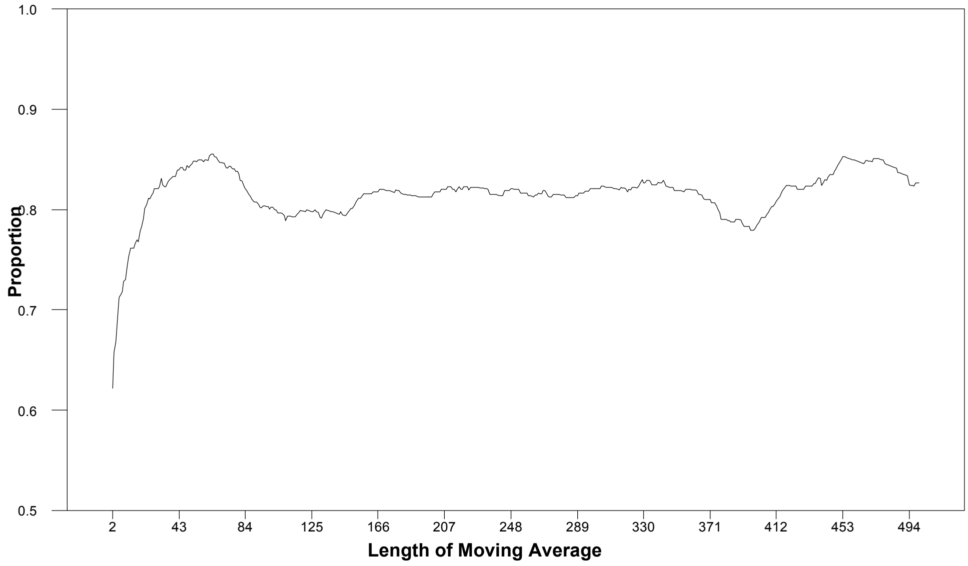

Proportion of Trades of MA Trading Rules against the Fed

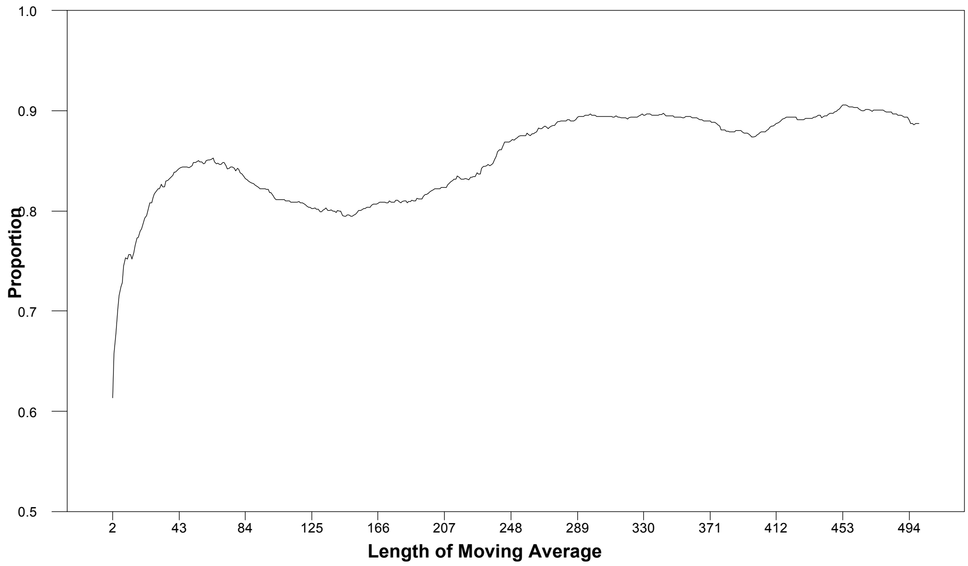

Proportion of Trades of MA Trading Rules against the Bundesbank

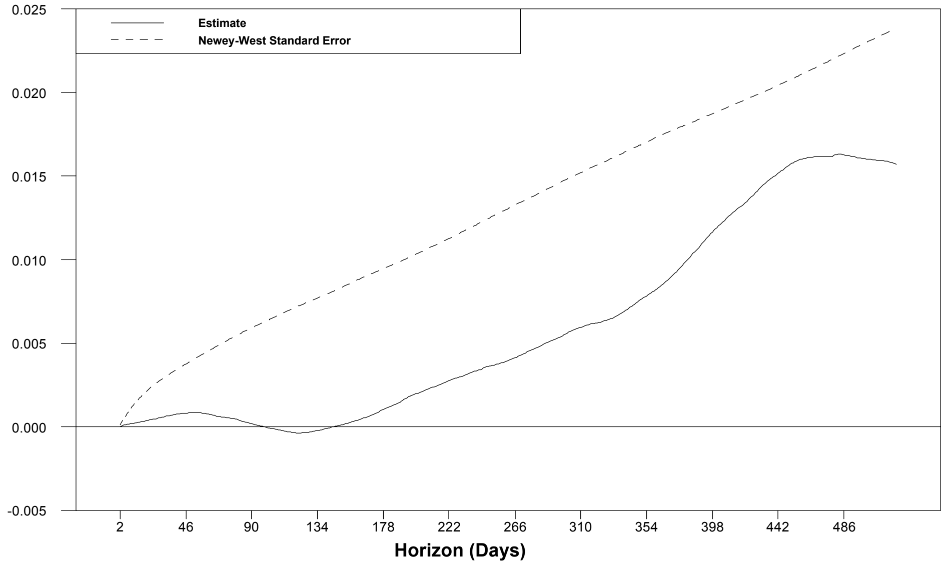

Regressing Fed Interventions on Subsequent Deviations from UIP: Estimates of Alpha

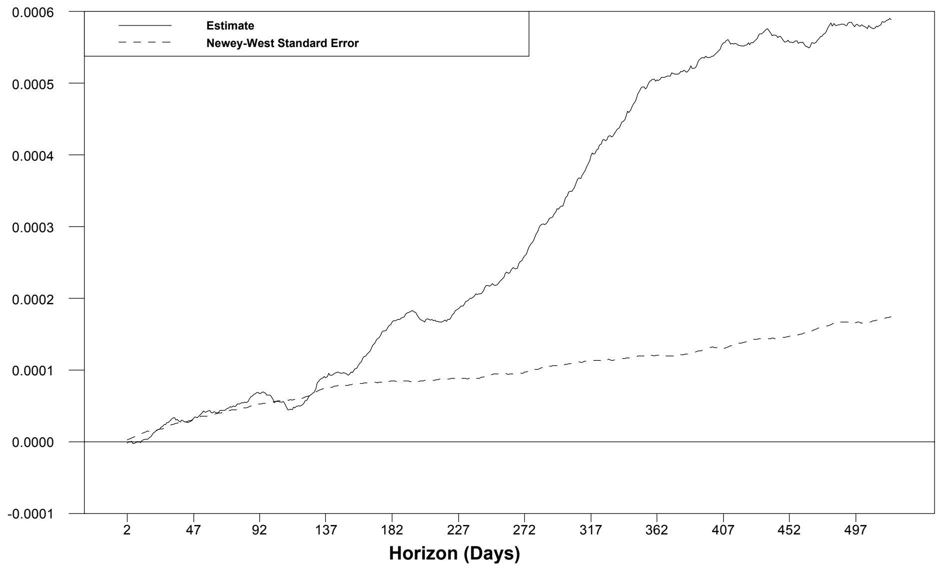

Regressing Fed Interventions on Subsequent Deviations from UIP: Estimates of Beta

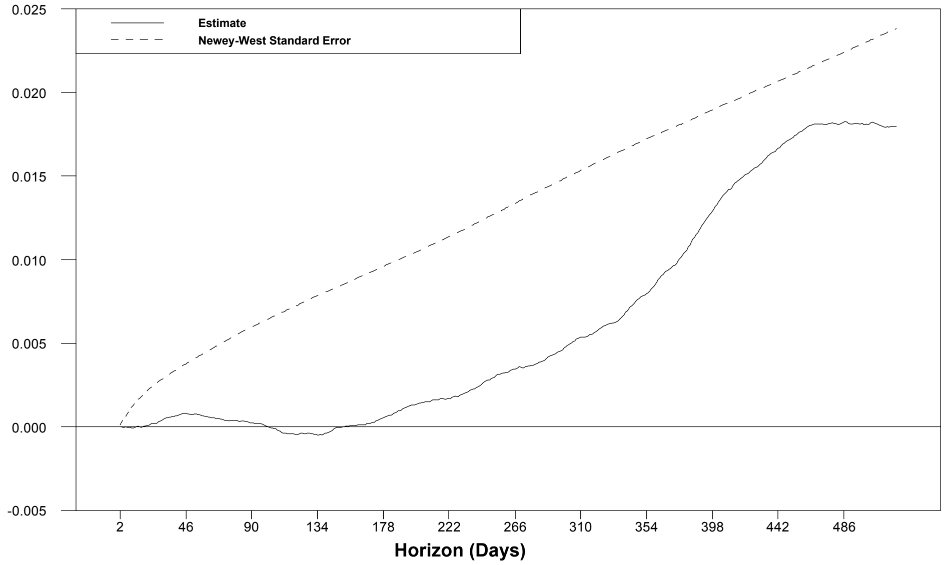

Regressing Bundesbank Interventions on Subsequent Deviations from UIP: Estimates of Alpha

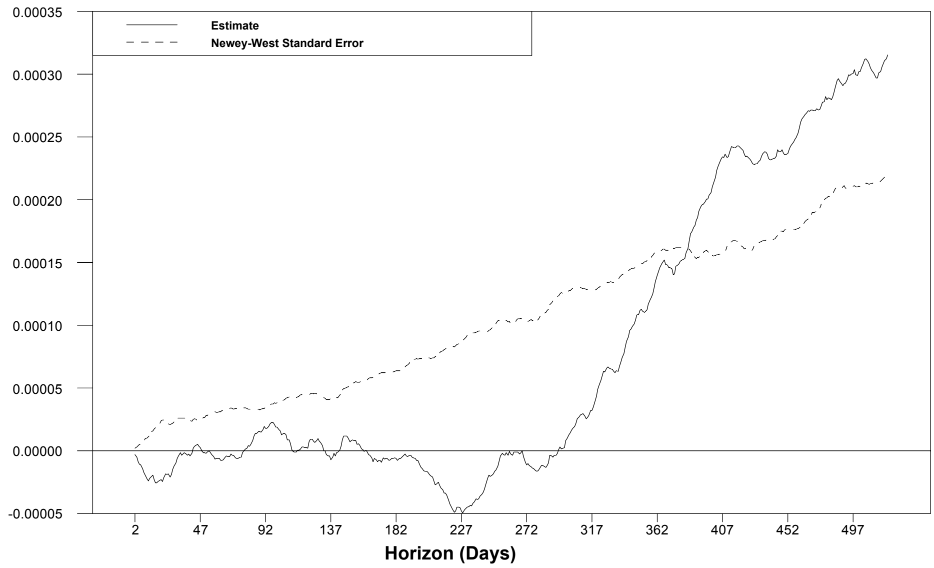

Regressing Bundesbank Interventions on Subsequent Deviations from UIP: Estimates of Beta

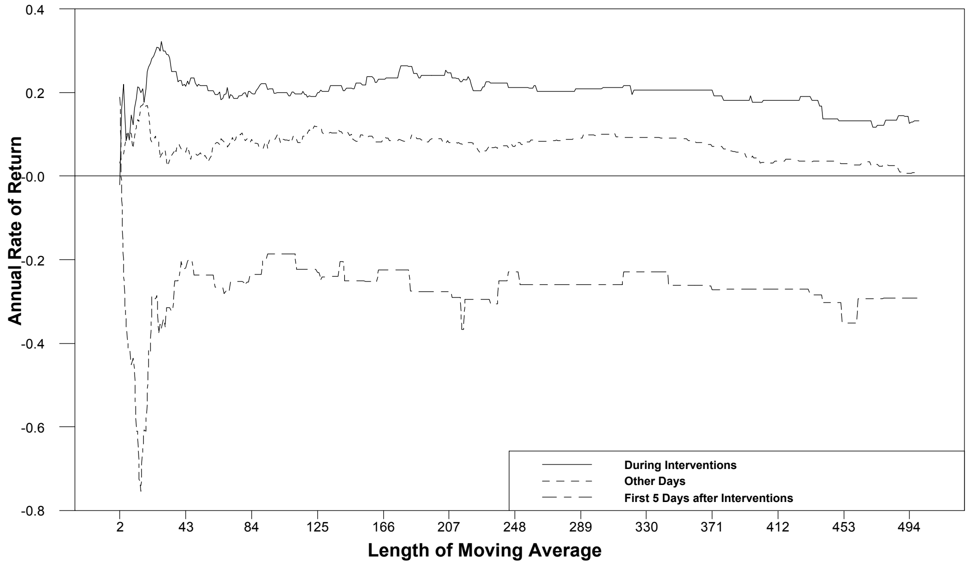

MA Trading Rule Returns During and After Intervention Periods

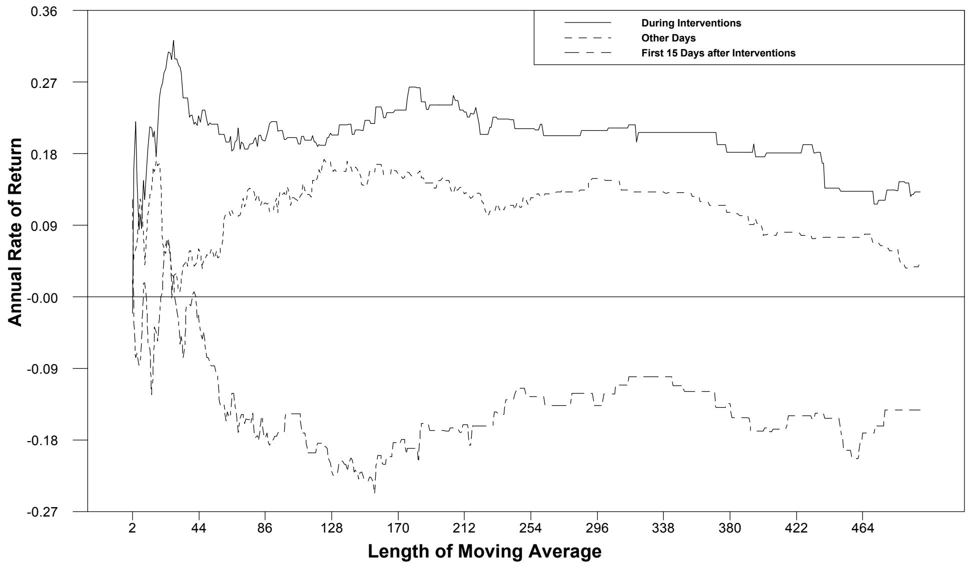

MA Trading Rule Returns During and After Intervention Periods

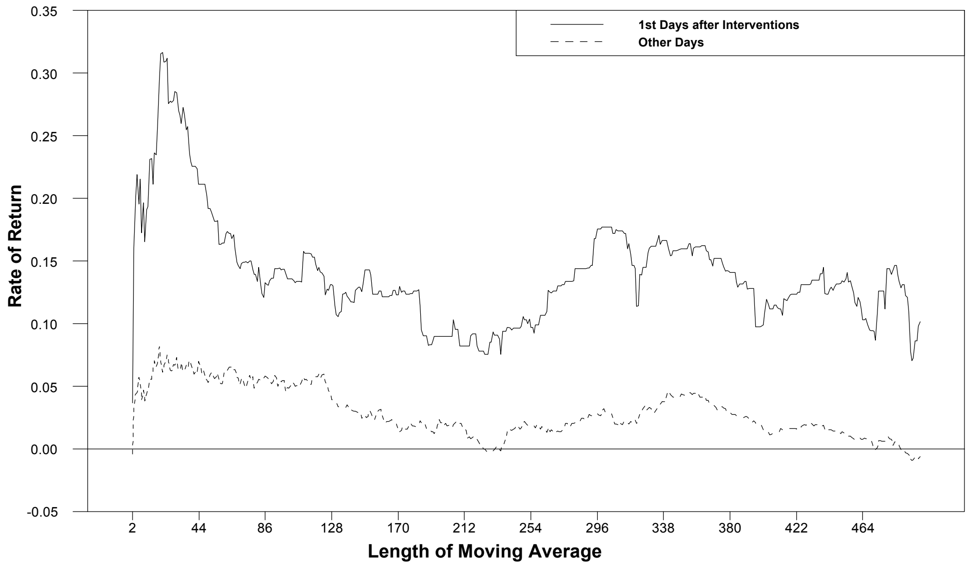

Profitability of Using MA Trading Rules to Try to Exploit Fed Interventions

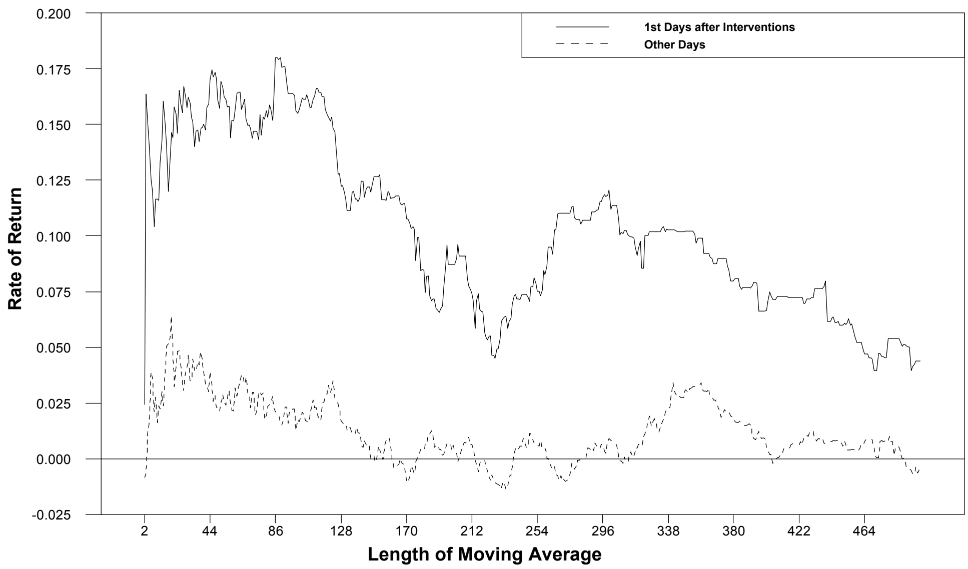

Profitability of Using MA Trading Rules to Try to Exploit Bundesbank Interventions

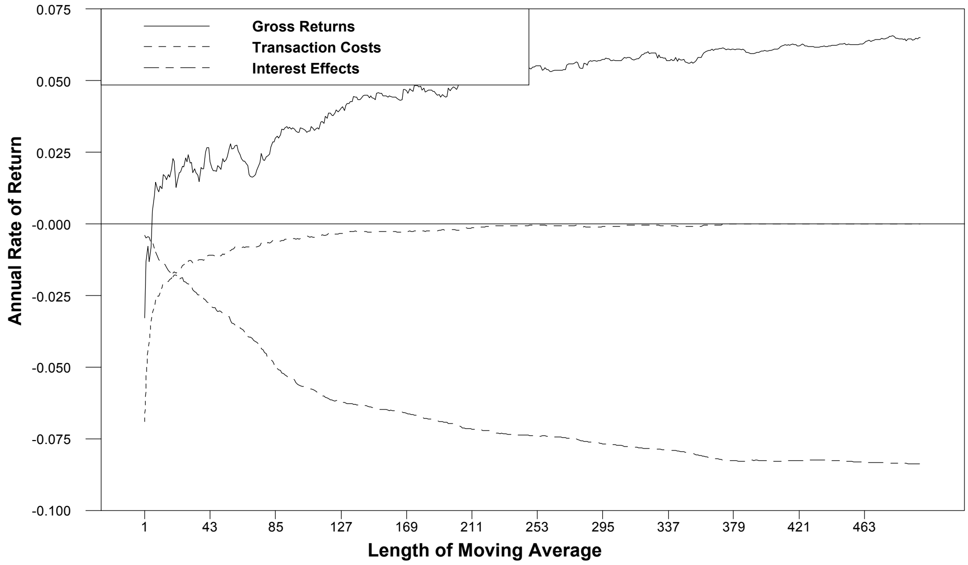

Components of Moving Average Trading Rule Returns for the DEM/ITL Exchange Rates 1979-1986

Appendix B

Omitted Tables

Table B.1: Autocorrelation Function of Daily JPY/USD Exchange Rate Changes

Laglength	AC	PAC	Q-Stat	Prob
1	0.015	0.015	0.8115	0.368
2	0.016	0.015	1.6704	0.434
3	0.034	0.034	5.8192	0.121
4	0.005	0.004	5.9009	0.207
5	0.037	0.036	10.774	0.056
6	-0.004	-0.006	10.834	0.094
7	0.006	0.005	10.949	0.141
8	0.021	0.019	12.549	0.128
9	0.021	0.021	14.147	0.117
10	0.023	0.020	16.035	0.099
11	0.011	0.009	16.476	0.124
12	0.014	0.011	17.181	0.143
13	0.022	0.018	18.848	0.128
14	0.020	0.017	20.222	0.123
15	0.024	0.021	22.297	0.100
16	0.003	0.000	22.339	0.133
17	-0.014	-0.018	23.033	0.148
18	0.001	-0.002	23.039	0.189
19	0.026	0.024	25.510	0.144
20	-0.007	-0.009	25.662	0.177

N.B. The asympt. S.E. of the corr. coeffs. is 0.0168.

Table B.2: Autocorrelation Function of Daily JPY/DEM Exchange Rate Changes

Laglength	AC	PAC	Q-Stat	Prob
1	0.080	0.080	22.744	0.000
2	0.006	-0.001	22.862	0.000
3	0.010	0.009	23.186	0.000
4	0.012	0.011	23.691	0.000
5	-0.028	-0.030	26.386	0.000
6	-0.001	0.004	26.390	0.000
7	0.036	0.036	31.007	0.000
8	0.040	0.035	36.678	0.000
9	0.013	0.008	37.302	0.000
10	0.013	0.010	37.932	0.000

N.B. The asympt. S.E. of the corr. coeffs. is 0.0168.

Table B.3: Parameter Estimates for GARCH(1,1) Model of Daily JPY/USD Exchange Rate Changes

	Coeff.	Std. Error	T-Stat.	p-value
α_3	0.0216	0.0168	1.2832	0.1994
α_5	0.0522	0.0162	3.2260	0.0013
β_0	3.84E-06	5.09E-07	7.548	0.0000
β_1	0.8184	0.0180	45.55	0.0000
β_2	0.0933	0.0088	10.55	0.0000

Table B.4: Parameter Estimates for GARCH(1,1) Model of Daily JPY/DEM Exchange Rate Changes

	Coeff.	Std. Error	T-Stat.	p-value
α_1	0.0677	0.0175	3.860	0.0001
α_3	0.0228	0.0173	1.313	0.1893
α_5	-0.0271	0.0174	-1.558	0.1193
β_0	7.78E-07	1.28E-07	6.075	0.0000
β_1	0.8728	0.0091	96.31	0.0000
β_2	0.1077	0.0081	13.36	0.0000

Table B.5: Day-of-the-Week Effects of Exchange Rate Changes

			Mon	Tue	Wed	Thu	Fri
DEM/USD	Exchange	Coeff	0.0003	-0.0002	-0.0004	-0.0002	0.0004
	Rate	Std Error	0.0003	0.0003	0.0003	0.0002	0.0003
	Changes	T-Stat	0.9511	-0.8921	-1.5410	-0.7192	1.4047
JPY/USD	Exchange	Coeff	0.0002	-0.0003	-0.0003	-0.0003	0.0002
	Rate	Std Error	0.0003	0.0002	0.0002	0.0002	0.0002
	Changes	T-Stat	0.5724	-1.3753	-1.1614	-1.4191	0.6477
JPY/DEM	Exchange	Coeff	-0.0001	-0.0001	0.0001	-0.0001	-0.0002
	Rate	Std Error	0.0002	0.0002	0.0002	0.0002	0.0002
	Changes	T-Stat	-0.5697	-0.4204	0.6017	-0.7407	-1.0377

Table B.6: Autocorrelation Function of Daily NLG/USD Exchange Rate Changes

Laglength	AC	PAC	Q-Stat	Prob
1	-0.075	-0.075	26.365	0.000
2	0.029	0.023	30.266	0.000
3	-0.002	0.002	30.292	0.000
4	0.041	0.040	38.167	0.000
5	0.031	0.037	42.676	0.000
6	-0.017	-0.014	44.008	0.000
7	0.083	0.080	76.539	0.000
8	-0.056	-0.046	91.497	0.000
9	0.017	0.003	92.881	0.000
10	-0.020	-0.017	94.812	0.000

N.B. The asympt. S.E. of the corr. coeffs. is 0.0142.

Table B.7: Autocorrelation Function of Daily FRF/USD Exchange Rate Changes

Laglength	AC	PAC	Q-Stat	Prob
1	0.004	0.004	0.0914	0.762
2	-0.085	-0.085	34.171	0.000
3	0.010	0.011	34.657	0.000
4	0.034	0.027	40.053	0.000
5	-0.006	-0.004	40.200	0.000
6	0.002	0.007	40.225	0.000
7	0.007	0.006	40.465	0.000
8	0.019	0.019	42.179	0.000
9	0.000	0.001	42.179	0.000
10	-0.015	-0.012	43.248	0.000

N.B. The asympt. S.E. of the corr. coeffs. is 0.0142.

Table B.8: Autocorrelation Function of Daily ITL/USD Exchange Rate Changes

Laglength	AC	PAC	Q-Stat	Prob
1	0.014	0.014	0.9705	0.325
2	-0.003	-0.003	1.0130	0.603
3	-0.012	-0.012	1.6437	0.650
4	0.056	0.057	16.489	0.002
5	0.002	0.000	16.502	0.006
6	-0.018	-0.018	18.031	0.006
7	-0.028	-0.026	21.690	0.003
8	0.001	-0.001	21.697	0.006
9	0.027	0.026	25.064	0.003
10	0.054	0.055	38.663	0.000

N.B. The asympt. S.E. of the corr. coeffs. is 0.0142.

Table B.9: Autocorrelation Function of Daily GBP/USD Exchange Rate Changes

Laglength	AC	PAC	Q-Stat	Prob
1	0.015	0.015	1.0376	0.308
2	0.007	0.007	1.2913	0.524
3	0.031	0.031	5.7451	0.125
4	0.006	0.005	5.9087	0.206
5	0.005	0.004	6.0086	0.305
6	0.012	0.011	6.6476	0.355
7	-0.003	-0.003	6.6798	0.463
8	0.018	0.018	8.1966	0.415
9	0.015	0.014	9.2863	0.411
10	0.061	0.060	26.573	0.003

N.B. The asympt. S.E. of the corr. coeffs. is 0.0142.

Appendix C

Derivation of Critical Sharpe Ratio

Proposition 1 *Suppose index returns come from a population with mean, μ, and standard deviation, σ, and:*

$$\sqrt{260} * \frac{\mu}{\sigma} = 0.32 \tag{C.1}$$

If $\{X_1, ..., X_N\}$ is a random sample drawn from this population, then the following approximation holds for large N: $P\left(\sqrt{260}\frac{\overline{X}}{S} > 0.32 + \frac{1.645}{\sqrt{\frac{N}{260}}}\right) \cong 0.05$

Proof. *One version of the central limit theorem states that:*[1]

$$\frac{\overline{X} - \mu}{\frac{\sigma}{\sqrt{N}}} \xrightarrow{D} N(0,1)$$

$$\Rightarrow \sqrt{260}\left(\frac{\overline{X} - \mu}{\frac{\sigma}{\sqrt{N}}}\right) \xrightarrow{D} N(0,260) \tag{C.2}$$

$$\Rightarrow \sqrt{N}\left(\sqrt{260}\frac{\overline{X}}{\sigma} - \sqrt{260}\frac{\mu}{\sigma}\right) \xrightarrow{D} N(0,260) \tag{C.3}$$

$$\Rightarrow \sqrt{N}\left(\sqrt{260}\frac{\overline{X}}{\sigma} - 0.32\right) \xrightarrow{D} N(0,260) \tag{C.4}$$

$$\Rightarrow \sqrt{N}\left(\sqrt{260}\frac{\overline{X}}{\sigma}\right) \xrightarrow{D} N(0.32,260) \tag{C.5}$$

Now, one of the Slutzky Theorems states that for sequences of random variables U_n and

[1] *See, e.g., Goldberger (1991), p.99.*

V_n, if $U_n \xrightarrow{P} c$ and V_n has a limiting distribution, then the limiting distribution of $(U_n V_n)$ is the same as that of cV_n.[2] *If we let*

$$V_n = \sqrt{N}\left(\sqrt{260}\frac{\overline{X}}{\sigma}\right) \tag{C.6}$$

and

$$U_n = \frac{\sigma}{S}, \tag{C.7}$$

then from (C.5) and because $U_n \xrightarrow{P} 1$ it follows that

$$U_n V_n = \xrightarrow{D} N(0.32, 260). \tag{C.8}$$

$$\therefore \sqrt{N}\left(\sqrt{260}\frac{\overline{X}}{S}\right) \xrightarrow{D} N(0.32, 260) \tag{C.9}$$

$$\Rightarrow \left(\sqrt{260}\frac{\overline{X}}{S}\right) \xrightarrow{A} N(0.32, \frac{260}{N}) \tag{C.10}$$

$$\Rightarrow P\left(\sqrt{260}\frac{\bar{X}}{S} > 0.32 + \frac{1.645}{\sqrt{\frac{N}{260}}}\right) \cong 0.05 \tag{C.11}$$

■

[2] *See, e.g., Goldberger (1991), p.102.*