

WORKING PAPER SERIES

Technical Trading Rules in the European Monetary System.

Christopher Neely and Paul Weller

Working Paper 1997-015C
<http://research.stlouisfed.org/wp/1997/97-015.pdf>

PUBLISHED: The Journal of International Money and Finance, June 1999.

FEDERAL RESERVE BANK OF ST. LOUIS
Research Division
411 Locust Street
St. Louis, MO 63102

The views expressed are those of the individual authors and do not necessarily reflect official positions of the Federal Reserve Bank of St. Louis, the Federal Reserve System, or the Board of Governors.

Federal Reserve Bank of St. Louis Working Papers are preliminary materials circulated to stimulate discussion and critical comment. References in publications to Federal Reserve Bank of St. Louis Working Papers (other than an acknowledgment that the writer has had access to unpublished material) should be cleared with the author or authors.

Photo courtesy of The Gateway Arch, St. Louis, MO. www.gatewayarch.com

Technical Trading Rules in the European Monetary System*

by

Christopher J. Neely†

Paul A. Weller‡

November 17, 1998

Primary Subject Code: G0 - Financial Economics

Secondary Subject Code: G14 - Information and Market Efficiency

Keywords: technical analysis, genetic programming, trading rules, exchange rates, European Monetary System, target zones

† Senior Economist, Research Department
Federal Reserve Bank of St. Louis
St. Louis, MO 63102
(314) 444-8568 (o), (314) 444-8731 (f)
Neely@stls.frb.org

‡ Department of Finance
College of Business Administration,
University of Iowa
Iowa City, IA 52242
(319) 335-1017 (o), (319) 335-3690 (f)
Paul-Weller@uiowa.edu

* The authors thank Kent Koch for excellent research assistance, Rob Dittmar for assistance with the computer code, their colleagues at the Federal Reserve Bank of St. Louis for generously sharing their computers at night and over weekends, and two anonymous referees for helpful comments. Paul Weller thanks the Research Department of the Federal Reserve Bank of St. Louis for its hospitality during his stay as Visiting Scholar when much of this work was completed. The views expressed are those of the authors and do not necessarily reflect official positions of the Federal Reserve Bank of St. Louis, the Federal Reserve System or the Board of Governors

Technical Trading Rules in the European Monetary System – by Christopher J. Neely
Paul A. Weller

Abstract:--97-015C

Using genetic programming, we find trading rules that generate significant excess returns for three of four EMS exchange rates over the out-of-sample period 1986-1996. Permitting the rules to use information about the interest rate differential proved to be important. The reduction in volatility resulting from the imposition of a narrower band may reduce trading rule profitability. Our results cannot be duplicated by commonly used moving average rules, filter rules or by two rules designed to exploit known features of target zone rates. There is no evidence that the excess returns are compensation for bearing systematic risk.

1. Introduction

Technical analysis – the use of current and past prices as inputs to trading strategies in a financial market – has a long history among investment professionals. However, only relatively recently have its claims found support within the academic community. The earliest evidence to suggest that simple technical trading rules were consistently profitable appeared in studies of the foreign exchange market (Dooley and Shafer, 1984; Sweeney, 1986). More recently Brock, Lakonishok and LeBaron (1992), using ninety years of daily data for the Dow Jones Industrial Average, showed that certain commonly used technical trading rules would have delivered economically significant excess returns over long periods. Further work on the foreign exchange market has confirmed and extended the original results (Levich and Thomas, 1993; Neely, Weller and Dittmar, 1997, henceforth NWD).

Despite the accumulating evidence on the value of technical analysis in dollar exchange rate markets, little is known about its effectiveness in other currency markets. This is particularly true in the case of currencies that are members of the Exchange Rate Mechanism (ERM) of the European Monetary System (EMS). One possible explanation for this neglect is the widely held belief that trading rules are most successful in markets that display high volatility. As Krugman and Miller (1993) point out, the original justification for constraining EMS exchange rates within a currency band was to reduce exchange rate fluctuations that were seen as discouraging trade and investment between member countries. To the extent that this objective has been achieved, it would be expected to have had a negative impact on the profitability of any trading rule that takes

positions in the currencies of member countries. Indeed, Lee and Mathur (1996) find little evidence of profitable trading rules for European cross-rates.

In this paper we are interested in examining a number of related questions. First, we ask whether it is possible to identify profitable trading rules *ex ante* for EMS currencies. We use the methodology based on genetic programming developed in NWD (1997) to avoid the bias inherent in selecting specific trading rules for investigation *ex post*. Since we find evidence for the existence of significant excess returns to trading rules, we then consider what the source of the excess returns might be. We show that information on past interest differentials is an important input to the trading rules for all currencies for which the rules are profitable. We look to see whether our results can be reproduced by moving average or filter rules that are commonly used by technical analysts. We also compare the performance of trading rules designed to exploit certain special features of exchange rates confined to a band. Finally we investigate whether the observed excess returns can be explained as compensation for bearing systematic risk.

We find strong evidence of economically significant excess returns for three of four deutschemark exchange rates. The performance of the rules is much more consistent across currencies than that of the moving average and filter rules, and it is also more robust to movements of the currencies in and out of the EMS and to bands of differing widths. There is no indication either that the trading rules exploit mean reversion, or that they are able to predict realignments. We find no support for the position that the excess returns can be interpreted as a risk premium. There is some suggestion that the reduction

in volatility caused by the imposition of a band does reduce trading rule profitability. The only currency for which there was no evidence of significant excess returns was the Dutch guilder, which was also the only currency that remained within a band of $\pm 2.25\%$ throughout our sample period. For the three remaining currencies the risk-return tradeoff of a benchmark portfolio is substantially improved when the trading rule returns are added, because of the low variability and correlation of the trading rule returns.

In Section 2 we outline our methodology. In Section 3 we describe the data and in Section 4 we present our results. Section 5 investigates various possible sources for the observed excess returns. In Section 6 we draw conclusions.

2. Methodology

We use genetic programming as a search procedure to identify trading rules that use information on current and past values of the exchange rate and the interest differential.¹ Koza (1992) developed the technique as an extension of the genetic algorithm of Holland (1975). The genetic algorithm applies the principle of Darwinian natural selection to complex problem solving, randomly generating a set of possible solutions and allowing them to “evolve” in successive generations to higher levels of “fitness”.

¹ For previous applications of the genetic programming procedure in the foreign exchange market, see Neely, Weller and Dittmar (1997) and Neely and Weller (1997). For an application to the equity market, see Allen and Karjalainen (1998).

The basic features of the genetic program used in this paper are (a) a means of encoding trading rules so that they can be built up from separate functions of the data, (b) a performance criterion, referred to as a measure of “fitness”, and (c) an operation that splits and recombines existing rules in order to create new rules. Before we describe these features, let us first introduce some notation. We adopt the convention that the deutschemark (DEM) is the domestic currency, and denote the exchange rate at date t (DEM per unit of foreign currency) by S_t . The domestic (foreign) overnight interest rate at date t is i_t (i_t^*). A trading rule can be thought of as a mapping from past exchange rates and interest rates to a binary variable, z_t , that takes the value +1 for a long position in foreign exchange at time t , and -1 for a short position. Trading rules are represented as trees whose nodes consist of various arithmetic functions, logical operators and constants. Examples of some of the arithmetic functions used are “average”, “max”, “min”, and “lag”. These functions take a numerical argument and return the value of the function (e.g., the moving average) of the data series (e.g., the exchange rate) taken over the period specified by the numerical argument, rounded to the nearest whole number. A function is also identified according to the series on which it operates. Thus $\max_s(3)$ at time t is equivalent to $\max(S_{t-1}, S_{t-2}, S_{t-3})$, $\text{lag}_i(3)$ at time t is equal to $(i_{t-3}^* - i_{t-3})$ and $\text{average}_s(3)$ is equal to the arithmetic average of S_{t-1} , S_{t-2} , S_{t-3} . Logical operators include “and”, “or”, “not”, “if-then” and “if-then-else”.

Figure 1 presents an example of a simple trading rule. The rule signals a long position at date t if the difference between the current exchange rate and its minimum over the last

five days is greater than the ten-day moving average of the interest differential. The function “rate” returns the spot exchange rate at a given time on date t .

The fitness criterion we use in the genetic program is the excess return to a fully margined long or short position in the foreign currency. The continuously compounded (log) excess return is given by $z_t r_t$ where z_t is the indicator variable described above, and r_t is defined as:

$$(1) \quad r_t = \ln S_{t+1} - \ln S_t + \ln(1 + i_t^*) - \ln(1 + i_t).$$

The cumulative excess return from a pair of round-trip trades (go long at date t , go short at date $t + k$), with round-trip proportional transaction cost c , is²

$$(2) \quad r_{t,t+k} = \sum_{i=0}^{k-1} r_{t+i} + \ln(1 - c) - \ln(1 + c).$$

Therefore the cumulative excess return r for a trading rule giving signal z_t at time t over the period from time zero to time T is:

$$(3) \quad r = \sum_{t=0}^{T-1} z_t r_t + \frac{n}{2} \ln \left(\frac{1 - c}{1 + c} \right)$$

where n is the number of trades. This measures the fitness of the rule.

² Each trade incurs a round-trip transaction cost because it involves closing a long (short) position and opening a short (long) one.

Figure 2 illustrates the reproduction operation used to create new rules. Two rules, the “parents”, are selected at random from a given generation of rules. The probability of selection is weighted according to the value of the fitness criterion evaluated over the training period. The rules in the population are ranked according to their fitness and each rule is assigned a probability weight proportional to $\frac{1}{1 + rank}$, where *rank* is the ordinal fitness rank of the rule in the selection period (see below). The two parent rules are then split at random and recombined as indicated to produce the new “offspring” rule, subject to the restriction that the resulting rule must be well-defined, and that it may not exceed a specified size (10 levels and 100 nodes).³

To implement the genetic programming procedures we formed 3 separate subsamples, known as the training, selection and validation periods. The training and selection periods were used to generate rules in the manner described below. The validation period was used to assess the out-of-sample performance of the rules. The distinct time periods for all currencies were chosen as follows: training period, 3/13/79 to 1/2/83; selection period, 1/3/83 to 1/1/86; validation period, 1/2/86 to 6/21/96.

The separate steps involved in implementing the genetic program are detailed below.

1. Create an initial generation of 500 randomly generated rules.
2. Measure the excess return of each rule over the training period and rank according to excess return.

³ As in our earlier work, we chose not to implement a mutation operator, because of the evidence that its impact is insignificant.

3. Select the highest ranked rule and calculate its excess return over the selection period. If it generates a positive excess return, save it as the initial best rule. Otherwise, designate the no-trade rule as the initial best rule, with zero excess return.
4. Select two rules at random from the current generation, using weights attaching higher probability to more highly-ranked rules. Apply the reproduction operator to create a new rule, which then replaces an old rule, chosen using weights attaching higher probability to less highly-ranked rules. Repeat this procedure 500 times to create a new generation of rules.
5. Measure the fitness of each rule in the new generation over the training period. Take the best rule in the training period and measure its fitness over the selection period. If it outperforms the previous best rule, save it as the new best rule.
6. Return to step 4 and repeat until 50 generations have been produced or until no new best rule appears for 25 generations.

This procedure generates a single trading rule. For each currency, we generated 100 rules in this way, and examined their performance over the validation period.

3. The Data

Daily U.S. dollar spot exchange rates and ECU central rates for the German mark (DEM), French franc (FRF), Italian lira (ITL), Dutch guilder (NLG) and British pound

(GBP) were collected by the European Commission.⁴ From the dollar rates, four exchange rate series – DEM/FRF, DEM/ITL, DEM/NLG and DEM/GBP – were constructed assuming the absence of triangular arbitrage. Overnight interest rates were recorded by *BIS*. The interest rates and exchange rates were collected at 10:00 a.m. and 2:15 p.m. Swiss time respectively. All exchange rate and interest rate data begin on 3/13/79 and end on 6/21/96.

Table 1 presents summary statistics on the exchange rate.⁵ We see that mean returns are very close to zero. The DEM/NLG is the least variable and the DEM/GBP the most variable of the four rates.⁶ All four rates are highly leptokurtotic. The DEM/FRF and DEM/ITL, which were “low credibility” ERM rates, are significantly skewed as well. The DEM/GBP is much less skewed and kurtotic than the other three rates. These results are consistent with those reported by Neely (1998), who found that weekly exchange rate changes for currencies in the ERM were much less variable but much more skewed and

⁴ Before 10/8/90, the ECU central rates for the British Pound were “notional,” that is, not binding on the British government. During this period the central rates were regularly changed by the European Commission to reflect current spot rates.

⁵ The mean return in Table 1 cannot be interpreted as the return to holding a long position in the foreign currency over the whole time period because the interest rate return accruing over weekends and holidays is not included. In all subsequent tables the reported figures include returns over weekends and holidays.

⁶ The DEM/GBP spent only 2 years, or 1/5 of the validation period, in the narrow band ($\pm 2.25\%$), whereas the DEM/NLG spent the whole period in the narrow band.

leptokurtotic than floating exchange rates.

Figure 3 shows the cumulative exchange rate return and interest rate differential for a long position in the DEM/FRF, including weekends and holidays. Over the whole sample period the cumulative interest differential has roughly matched the depreciation of the franc.

4. Results

We applied the genetic programming techniques explained in Section 2 to the four exchange rates (DEM/NLG, DEM/FRF, DEM/ITL, DEM/GBP) using a training period of 3/13/79 to 1/2/83 and a selection period of 1/3/83 to 1/1/86.⁷ We ran 100 independent trials for each exchange rate, thus obtaining 100 rules for each currency. We then examined the excess returns generated by these sets of rules over the validation period (1/2/86 to 6/21/96).

The measure of excess returns in (3) includes the round-trip transaction cost c . As in our previous work (NWD, 1997) we use a higher level of transaction cost in the training and

⁷ As noted in the previous section, our exchange rate series were constructed from dollar series. We checked our results against true cross-rate series obtained from DRI and found the results to be essentially identical. To ensure stationarity, the exchange rates used as information for the genetic program were normalized by dividing by respective ECU central rates.

selection periods (in-sample) than we do in the validation period (out-of-sample). In the latter interval, we choose c to reflect the costs faced by a large institutional trader (0.05%). However, in the in-sample period, we use a higher value for c (0.1%) to penalize rules that trade too often (see NWD, 1997, p. 414). This procedure, combined with the use of selection period performance as a check on the robustness of the in-sample profitability of a rule, has proved to be an effective means of reducing the dangers of overfitting the data.

During the validation period some of the currencies moved in and out of bands of varying widths. We divide the validation period for each currency into two subperiods, according to whether the currency was in a band of $\pm 6\%$ or less, or in a wide band of $\pm 15\%$ or outside the ERM. Table 2 describes these subperiods that we label S1 and S2.⁸ Panel A of Table 3 provides information on excess returns over the whole of the validation period. The mean annual excess return was positive for all four currencies.⁹ It ranged from six

⁸ The only currency for which these subperiods leads to aggregating returns for different bandwidths is the DEM/ITL. We examine below a further subdivision of subperiod S1 for this currency into the times when it was in 2.25% and 6% bands (see Table 6).

⁹ A referee has pointed out to us that we introduce a small upward bias to the reported excess return by neglecting the overnight interest rate spread. When a trader takes a short position, he must borrow the foreign currency and invest in the domestic currency, and so loses the spread. The size of the bias depends on the proportion of time a rule takes a short position, but we estimate that it is unlikely to exceed ten basis points per year for any currency over our sample period.

basis points for the DEM/NLG, which spent the whole of the period within a 2.25% band, to 2.75 percent for the DEM/GBP, which spent two years in a 6.0% band, and for the rest of the period was outside the ERM. For all currencies except the DEM/NLG over 90% of the rules generated positive excess returns. For the DEM/GBP only one rule earned a negative excess return.¹⁰

To assess the likelihood that these results could have been obtained by chance, in addition to reporting t-statistics, we calculate Bayesian posterior probabilities. We do this because classical test procedures have significant disadvantages in the present circumstances. They are concerned with determining the probability of rejecting the null hypothesis given that it is true. They do not balance that concern against the danger of falsely accepting the null, namely that the excess return to a trading rule is nonpositive. It is evident that the consequences of not identifying profitable trading rules are as costly as those of using unprofitable rules. In contrast, Bayesian posterior odds ratios summarize the evidence in favor of one hypothesis relative to another. A ratio greater than 0.5 means that the evidence favors the hypothesis of positive excess returns.

The second row shows the average t-statistic over all 100 rules calculated with a Newey-West correction for serial correlation — for the null hypothesis that mean return from each rule equals zero. We caution that these t-statistics are indicative only of the average statistical significance of individual rules, and not of the statistical significance of the mean return to a portfolio constructed from the 100 rules. Those statistics are shown in

¹⁰ The mean annual excess return for this rule was - 2.09%.

Table 4. The third row of Table 3 shows the average Bayesian posterior probability, across the 100 rules, that the mean excess return is greater than zero.¹¹ For all exchange rates except the DEM/NLG this figure exceeds 0.8 for the full out-of-sample period.

In panels B and C of Table 3 we report performance of the rules over the subperiods defined in Table 2. For the DEM/FRF we find, as one would expect, that both volatility and excess return increase when the exchange rate is not confined to the narrow band. However, there is a greater proportionate increase in excess return. In the case of the lira the differences are particularly striking. Volatility increases by a factor of four and excess return by a factor of fifteen. The sterling rules, however, perform noticeably better in the 6.0% band than outside it. This is puzzling in view of the fact that the rules were generated over training and selection periods when sterling was outside the ERM, and that volatility was also substantially higher over the validation subperiod when sterling was outside the ERM.

Trading rules generated from distinct trials are generally not identical, although they are often very similar. This variability in the structure of the rules produces differing performance levels out-of-sample. For this reason examining the performance of individual rules is less useful than looking at composite trading rules. Table 4 shows the

¹¹ The Bayesian posterior probabilities were calculated assuming a normal likelihood function, a uniform improper prior on the autoregression coefficients and an autoregressive process of order zero-to-ten lags for the trading rule returns. The order of the autoregressive process was chosen by the Schwarz criterion for each case.

returns from the median and uniform portfolio rules. The median portfolio rule dictates a long position if 50% or more of the individual rules signal a long position, and a short position otherwise. The uniform portfolio rule gives an equal weight to each of the 100 individual rules. The mean returns for both the portfolio rules are always positive for every exchange rate in every subsample.

4.1 The Performance of Portfolio Rules

We have found in our earlier work (NWD, 1997) that the median portfolio rule returns generally exceed those of the uniform rule. We confirm this finding in the present study.¹² The mean annual excess return over all four currencies rises from 1.62 percent for the uniform rule to 2.16 percent for the median rule. However, this is accompanied by an increase in the variability of returns, reflected in the fact that there is no improvement in Sharpe ratios. The explanation for this improvement in performance lies in the stochastic nature of the search procedure embodied in the genetic program. Even when the majority of rules consistently identify similar patterns in the data, occasionally the outcome of a trial will produce an inferior rule. Whereas the uniform portfolio rule attaches equal weight to all rules, the median portfolio rule attaches zero weight to those

¹² In general one must be wary of using particular weighting schemes in investigations such as the present one, since there is an obvious danger of data mining. However, in this case our decision to use the median portfolio rule was made *ex ante* on the basis of evidence from our earlier study (NWD, 1997).

rules whose signals are in the minority. This proves to be an effective means of screening out the influence of inferior rules.

The returns to the uniform rule are similar to the average returns reported in Table 3, but rise by between 5 and 15 basis points because we calculate net transaction costs for the uniform rule. Thus if two rules with opposite positions each reverse their positions on the same date they generate no net change in position for the uniform rule and so incur no transaction cost. In the computation of average return in Table 3 each rule would have incurred the round-trip transaction cost c .

For all currencies except the DEM/NLG we find posterior odds ratios above 0.9 for the whole out-of-sample period. This evidence provides strong support for the existence of economically and statistically significant excess returns for the DEM/FRF, DEM/ITL, and DEM/GBP. The volatility of returns varies considerably across currencies, affected mainly by the length of time spent in the bands of varying width. However, returns for the DEM/NLG are much less variable than for the DEM/FRF even over the periods when both currencies were in the 2.25% band. This is a clear reflection of the fact that the DEM/NLG band has had very high credibility. The mean annual trading frequency for the different currencies is quite similar, with approximately one trade every two months.

Table 5 presents the results of joint tests that the mean returns across exchange rates to the median and uniform portfolio rules are greater than zero. The probability values for the joint test across all four exchange rates for the median and uniform portfolio rules are

0.106 and 0.074, respectively. Of course one might have suspected, *ex ante*, that the stable DEM/NLG would provide the least scope for profitable technical trading rules. Therefore, in the third and fourth rows we present statistics excluding the DEM/NLG returns from the calculations. This exclusion lowers the significance levels to 0.055 and 0.038, respectively.

We note that the uniform and median portfolio rules for the DEM/NLG deliver annual returns of 0.03 and 0.12 percent, respectively, with 5.82 and 5.73 trades per year. With zero transactions costs, these portfolio rules would have returned 0.32 and 0.41 percent per year, which would have produced Sharpe ratios of approximately 0.5 and 0.67. The rules were evidently successful in finding patterns in the DEM/NLG data, as is confirmed by the timing tests reported later, but were unable to exploit these patterns profitably.

We have already noted evidence to suggest that there is a relationship between variation in the exchange rate and excess returns. We now turn to examining this in greater detail. We define a variable %XS as a measure of the proportion of total possible excess return captured by a trading rule:

$$(4) \quad \%XS = 100 * \left(\frac{\sum_{t=0}^{T-1} z_t r_t}{\sum_{t=0}^{T-1} |r_t|} \right)$$

The results of calculating the values of this measure for the portfolio rules are shown in Table 4. There are some differences depending on whether we consider the uniform or median rule. The uniform rule for DEM/ITL captures more of the total return when outside the narrow band, thus magnifying the effect of the increase in volatility on excess

return. However, as we will discuss below, this is in considerable part accounted for by the negative excess return earned in the 2.25% band. This depresses the average return in the narrow band subperiod, which aggregates periods when the currency was in a 6% band and a 2.25% band. The reverse is true of the DEM/GBP. Both uniform and median rules capture substantially more of the total excess return when sterling was in the 6% band than when it was out of the ERM. Overall no clear pattern emerges.¹³

Each sample subperiod corresponds to a single bandwidth for all currencies except the DEM/ITL. From Table 2 we see that the period S1 lumps together the separate periods when the lira was in a 6% and a 2.25% band. To determine whether the performance of the trading rules for the lira was dependent on the bandwidth we present the results of further subdividing the sample in Table 6. Both the uniform and the median portfolio rules performed poorly when the currency was in a 2.25% band, but did well in the 6% band. The difference is particularly striking in the case of the median portfolio rule, which had an excess return of 5.34% in the 6% band, and of - 4.10% in the 2.25% band.¹⁴ Since the training and selection periods for the lira occurred when the lira was in a 6% band, this suggests strongly that the currency behaved in an importantly different way during the period when it was confined to a 2.25% band.

¹³ The %XS return statistics for the DEM/NLG are much smaller than for the other currencies because the numerator is calculated net of transactions costs. Figures that are broadly comparable across currencies are obtained if transactions costs are set to zero.

¹⁴ The p-value for the test of the null of equality between the means is 0.07.

4.2 *Market Timing Tests*

As a separate check on the performance of the median portfolio rule we carry out the market timing tests described in Cumby and Modest (1987). They propose a test of the ability of a trading rule to predict the sign of an excess return one period ahead. It involves regressing the excess return to a long position on the indicator variable z_t , which takes the values +1 or -1 depending upon whether the rule signals a long or a short position. The results are presented in Table 7. There is strong evidence to indicate that the rules do possess predictive ability. The joint test of the null of no predictability for any of the exchange rates has a p-value of 0.003.¹⁵ However, the fact that the smallest p-value (0.01) is recorded for the DEM/NLG confirms the suggestion above that a trading rule can possess significant predictive ability without necessarily being profitable. Evidently the impact of low volatility coupled with the effect of transactions costs leads to economically insignificant excess returns. Conversely, the p-value for the DEM/GBP is 0.17 but the mean excess return to these rules is 3.14 percent.

The reader may wonder about differences between the t-statistic for the return to the median portfolio rule and the Cumby-Modest regression coefficient test statistic. These test statistics may provide different inferences for three reasons. First, because the Cumby-Modest regression includes a constant, the Cumby-Modest test statistic will depend explicitly on the product of the relative proportion of the time spent long and the

¹⁵ The Cumby-Modest test statistics are virtually identical if standard errors are calculated with the White heteroskedasticity-consistent covariance matrix.

excess return to a long position held throughout the sample period. This term may be positive or negative, making the Cumby-Modest test statistic larger or smaller than the t-statistic for the median portfolio rule. Second, the t-statistic for the median portfolio rule shown in Table 4 is adjusted for transaction costs while the returns used in the Cumby-Modest test are not. Finally, the t-statistic is calculated with a fifth-order Newey-West correction for serial correlation, while the Cumby-Modest statistic is not.

5. The Source of Excess Returns

We have shown that trading rules identified by a genetic program can generate economically significant excess returns for currencies in an exchange rate band, even when a narrow band is maintained. We now turn to considering the source of these excess returns. We first examine the structure of some of the simplest rules obtained for each currency. A feature that emerges consistently across all currencies is that the interest differential and not the past exchange rate series is the most important informational input to the trading rule. For example, a DEM/ITL rule which was the twelfth best during the out-of-sample period and earned an excess return of 4.88 per cent per year took the form “Take a long position if the minimum interest differential (Italian minus German) over the last four days exceeds 3.88.” The signals generated by this rule had a correlation of 0.96 with those of the median portfolio rule. An even simpler rule was found for the DEM/GBP. It was ranked 28th in terms of performance over the validation period, had an excess return of 3.32 per cent per year and a correlation of 0.95 with the median rule. It signaled a long position if the interest differential (British minus German) was greater

than 4.42 per cent and a short position otherwise. In fact the five simplest rules i.e. those with the smallest number of nodes, for DEM/ITL and DEM/GBP all used information only about the interest differential. In the case of the DEM/FRF four of the five simplest rules used information only about the interest differential. We confirmed the importance of the interest differential information by performing the following experiment. We ran the trading rules for all currencies over the validation period with the interest differential set to zero in the signal calculations, but included in the return calculations. The mean annual return fell to -0.24% and annual trading frequency dropped to 0.24.

The striking feature of the simple rules based on the interest differential is that they indicate a tendency for a higher interest rate for a particular currency today to be associated with a positive excess return to the currency tomorrow. However, as the form of the rules described above shows, and as we confirm below, the relationship does not imply that the rule “Take a long position in the currency with the higher interest rate” will necessarily perform well.

5.1 A Comparison with the Performance of Moving Average and Filter Rules

There is good evidence that in dollar currency markets various moving average (MA) and filter rules are able consistently to earn significant excess returns (Dooley and Shafer, 1984; Sweeney, 1986; Levich and Thomas, 1993). We therefore compare the performance of some simple MA and filter rules with the median portfolio rule for each currency. The MA rule (x, y) gives a long signal when the x -day moving average exceeds

the y -day moving average, and a short signal otherwise. The filter rule with filter x gives a long signal when the current exchange rate has risen by more than $100*x$ percent from its minimum over the last five days. It gives a short signal if the exchange rate has fallen by more than $100*x$ percent from its maximum over the last five days. Otherwise the current position is maintained. We examine four MA rules, (1,10), (1, 50), (5, 10), and (5, 50) and four filter rule sizes ($x = 0.005, 0.01, 0.015$ and 0.02).

When we compare the results from Table 4 with those from Tables 8 and 9, the general picture that emerges is that both the uniform and the median portfolio rules had a much more consistent performance, both during the overall sample period, and during the subperiods S1 and S2. Neither of the rules produced a negative excess return for any currency during any sample subperiod, whereas this is not true for any MA or filter rule. The annual excess return for the MA rules averaged over all currencies was - 0.05 per cent, and for the filter rules was - 0.19 per cent. As reported above, comparable figures were 1.62 per cent for the uniform rule, and 2.16 per cent for the median rule. In summary, six of the eight conventional trading rules performed rather poorly, and the two MA rules which did well still did not match the performance of the median rule averaged over all currencies.¹⁶

¹⁶ As a referee has pointed out to us, we cannot attach too much significance to the observed margin of superiority in the case of MA and filter rules, because we have not identified the best (in-sample) such rules. However, we would certainly expect the qualitative result to be preserved, because of our finding that the best genetic programming rules rely predominantly on interest rate data.

5.2 *A Comparison with the Performance of Floating Rate and EMS-specific Rules*

We next take the set of trading rules obtained for the DEM against the dollar in our earlier paper (NWD, 1997) and examine their performance for the present set of currencies. These rules were generated with training and selection periods between 1975 and 1980, and produced an average annual excess return of over 6 per cent when run on data for the DEM/USD over the period 1981-95. The results are presented in Table 10. The rules perform comparably well for the DEM/ITL and DEM/GBP, the two currencies that spent part of the sample period out of the ERM. In fact they do slightly better for the DEM/GBP than the rules obtained in the present investigation. Since the DEM/GBP spent all but two years of the sample period outside the ERM, this is the currency one would have expected to behave most like a “typical” floating exchange rate. It is interesting that the features of the data identified for a different currency using earlier training and selection periods are clearly present in the DEM/GBP data. In contrast, the rules do very poorly both for DEM/FRF and DEM/NLG, currencies that spent most, or all of the sample period in the 2.25% band.¹⁷

The results of these comparisons support the observation made above, that the structure of trading rules identified by the genetic program in this study is significantly different

¹⁷ We note a minor caveat here. The DEM/USD rules were generated with a genetic program that did not incorporate information about the interest differential. We do not think this difference significantly affects the conclusions.

from those that have been shown to perform well in dollar markets, and as a consequence they perform better.

It is possible that the genetic programming rules are taking advantage of simple trading strategies that are more appropriate than MA or filter rules for exchange rates within a band. To investigate this issue we consider two such strategies. The first assumes that the interest differential is the dominant source of any excess profits, and dictates that the trader should take a long position in the currency which has the higher interest rate. The second is founded on evidence that ERM currencies have displayed mean reversion, or a tendency to move towards the midpoint of the band, and so prescribes taking a long position in a currency when it is in the weak half of the band. The excess returns generated by these trading strategies are presented in Table 11. If we compare the returns from the high interest rate rule to those produced by the median portfolio rule we see that in every case the median rule does better.¹⁸ For the DEM/ITL and DEM/GBP the difference is comfortably over 3% per annum. The mean reversion rule also does consistently more poorly than the median rule. This demonstrates that the genetic

¹⁸ We analyze these simple rules to provide some information on the source of the profitability of the genetic programming rules. We do not interpret these results as a comprehensive comparison of different classes of trading rules. In particular, the simple interest rate rule used here should not be confused with more elaborate rules proposed for dollar exchange rate markets by Hodrick and Srivastava (1986) and Sweeney and Lee (1990).

programming rules take advantage of different features of the data from those conjectured above.

An important feature of the EMS has been the occurrence of periodic realignments of central parities. The return to being on the right side of the market at the time of a realignment can be substantial, and so it is of interest to determine whether the trading rules were able to profit from these episodes. Figure 4 shows the average annualized daily returns of the median portfolio rule from five days before to five days after a realignment for the DEM/FRF and the DEM/ITL.¹⁹ It is clear that the rules tend on average to take losing positions around realignments and that their overall performance is somewhat degraded by these events. The positive excess returns generated by the trading rules occur despite the losses incurred around realignments.

5.3 *Systematic Risk Exposure and Improvements to the Risk-Return Tradeoff*

So far we have succeeded in ruling out certain possible explanations for the observed excess returns to the trading rules we have identified. However, we have not directly addressed one of the most obvious, and also one of the most difficult questions to answer, namely, whether the excess returns represent compensation for bearing risk. Given the lack of a satisfactory model of the risk premium in the foreign exchange market, this is not something we can hope to resolve conclusively one way or the other. What we do in

¹⁹ During the validation period there were no realignments for the DEM/NLG, and only entry and exit observations for the DEM/GBP.

order to provide some useful evidence is to calculate the betas associated with the returns to the median portfolio trading rule, using two benchmark portfolios, the MSCI World Portfolio Index and the Commerzbank Index of German stocks. The results are presented in Table 12. The betas are not statistically significantly positive for any of the four exchange rates using either benchmark. Performing the regression with or without a constant does not change the inference. We conclude that there is no evidence that the excess returns earned by the technical trading rules are compensation for bearing systematic risk.

Sweeney and Lee (1990) advocate another method for adjusting the returns for risk. They assume that there is a constant risk premium (or discount) to a long position in the foreign currency. Because a short position earns the negative of the premium to a long position, the risk-adjusted return is given by

$$(5) \quad X^* = \sum_{t=0}^{T-1} z_t r_t + \frac{n}{2} \ln \left(\frac{1-c}{1+c} \right) - (p_1 - p_2) \cdot \sum_{t=0}^{T-1} r_t \quad ,$$

where p_1 is the proportion of the time spent in foreign currency and p_2 is the proportion of the time spent in the domestic currency ($p_1 + p_2 = 1$). To the extent that the rule is even-handed between long and short positions, or the return to the long position is close to zero, the X^* statistic will be close to the unadjusted return.

Table 13 shows that most of the (annualized) X^* statistics are close to the unadjusted mean annual portfolio returns that are shown in Table 4. In particular, the X^* adjustment does not systematically lower (or raise) the return. The largest differences between X^*

and the mean returns are observed for the GBP during the narrow target zone subsample, where the annualized X^* statistics for the two portfolio rules are a little over two percentage points lower than the mean annual return. The posterior probabilities still strongly favor the hypothesis that these two X^* statistics (2.35 and 3.21) are positive.²⁰ The X^* risk adjustment does not alter our previous conclusions.

To give some idea of the improvement in risk-return tradeoff potentially attainable with the use of these trading rules, we performed some simple calculations. We calculated the increase in excess return that a mean-variance investor would have earned if he had split his wealth optimally between a benchmark world portfolio and a given trading rule in such a way as to hold a portfolio with the same variance as the benchmark. We also calculated the portfolio weights to determine the extent to which the optimal portfolio was levered. In order to explain the calculations we define the following notation:

$$(6) \quad r = (r_b, r_t)$$

where r_b , r_t are the excess returns to benchmark and trading rule respectively. We denote by r_o the excess return to the optimal portfolio with standard deviation equal to σ_b , the standard deviation of the benchmark portfolio. V is the (2x2) covariance matrix of excess returns to benchmark and trading rule.

²⁰ The t statistic from the coefficient on the position variable in the Cumby-Modest test is virtually equivalent to the test statistic from the test of the significance of X^* proposed by Sweeney and Lee (1990), if transactions costs are ignored.

Then the vector of portfolio weights w on the risky assets in the optimal portfolio is given by:

$$(7) \quad w = \frac{r_o}{r'V^{-1}r} V^{-1}r$$

and the total excess return r_o on the optimal portfolio is given by:

$$(8) \quad r_o = \sigma_b \sqrt{r'V^{-1}r} .$$

The results of these calculations are given in Table 14. Our benchmark is the MSCI World Portfolio Index which had a mean annual excess return of 1.62 percent, measured in DEM, over the subsample 1988-1996. These risk-return calculations need to be interpreted cautiously because they assume that observed sample means and variances are known and equal to the true population values. Because the excess return of the median portfolio in the case of the DEM/NLG was slightly negative over the subsample, the optimal portfolio requires that a short position in the rule be taken. The optimal portfolio then uses a trading rule that takes precisely the opposite position from that generated by the genetic program. Clearly such figures are meaningless and we do not report them, but confine attention to the remaining currencies where we can assert with a high degree of confidence that excess returns are positive. In all cases the degree of leverage is quite high. This is a consequence of two important facts: first, that the variability of excess returns from the trading rules is much lower than from the benchmark, and second, that the correlation between excess returns is insignificantly different from zero. The excess return earned by the optimal portfolio is from five and a half to eight percentage points

above the excess return from the benchmark. Even taking into account the inevitable upward bias resulting from assuming that the population parameters of the excess return distributions are known, it is evident that the impact of the trading rules on the risk-return tradeoff is substantial.

6. Summary and Conclusions

We find that it is possible to identify profitable technical trading rules for three of the four EMS currencies we have examined. The rules are shown to produce economically significant excess returns after transaction costs over a ten-year out-of-sample period. There is some evidence that trading profits were higher during periods when the currencies were moving in broad rather than narrow bands. However, this finding is almost entirely attributable to the performance of the rules trading the Italian lira. Over the four-year period from 9/17/92 when the lira was outside the ERM, the median portfolio rule earned an annual excess return of 8.72 per cent. The excess return was 5.34 per cent when the lira was in a 6% band. In contrast, the annual excess return was - 4.10 per cent over the remainder of the sample period when the lira was confined to a 2.25% band.

There is evidence that currency volatility is an important determinant of trading rule profits. The Dutch guilder had by far the lowest variability against the German mark, and was the only currency where there were no significant profits to be earned. The lira had the highest variability of all currencies during the period when it had been withdrawn

from the narrow band, and as we have already noted, trading rule profitability was also very high.

In addition to examining the out-of-sample profitability of the trading rules, we perform the Cumby-Modest test to see whether the rules possess significant predictive power for the sign of the excess return one day ahead. We find strong evidence that the rules do possess such power.

These results confirm the effectiveness of the methodology based on genetic programming used in our previous paper (NWD, 1997). In this paper we have used a different set of currencies and different training, selection, and validation periods from those in NWD (1997). We have also extended the technique of the previous paper by allowing the rules to accept input from two distinct time series, exchange rates and interest differentials. This has proved to be particularly important in the present investigation. The great majority of all simple trading rules found for all currencies used information only about the interest differential. This fact points to a crucial difference between rules which work well for floating rates on the one hand, and those which work well for rates confined to a currency band on the other.

We consider various pieces of evidence in order to throw light on the source of the excess returns. We look at the returns to four moving average rules and four filter rules. The overall performance of the median portfolio rule generated by the genetic program is substantially better on average, and also much more consistent over currencies and

sample subperiods corresponding to times when currencies were moving within or outside narrow bands. We also examine two simple trading strategies designed to take advantage of certain special features of EMS data. Neither produces more than a small fraction of the excess returns earned by the median portfolio rule. To discover whether the returns to the trading rules can be interpreted as compensation for bearing systematic risk, we calculate betas for the returns to each portfolio of 100 rules over the period 1986-96 against two benchmarks - the Morgan Stanley Capital International (MSCI) world equity market index and the Commerzbank index of German equity. We find no indication that any of the estimated coefficients are significantly different from zero.

Finally, we present calculations to show that a mean-variance investor who takes advantage of the profitable trading rules can significantly improve his risk-return tradeoff in comparison to a benchmark world stock portfolio. The magnitude of the improvement is a consequence of the much lower variability of returns to the trading rules, and of their low correlation with the returns to the benchmark portfolio.

References

- Allen F, Karjalainen R, 1998. Using Genetic Algorithms to Find Technical Trading Rules. Working Paper, The Wharton School, University of Pennsylvania. Forthcoming in the Journal of Financial Economics.
- Brock W, Lakonishok J, LeBaron B, 1992. Simple Technical Trading Rules and the Stochastic Properties of Stock Returns. *Journal of Finance* 47, 1731-1764.
- Cumby RE, Modest DM, 1987. Testing for Market Timing Ability: A Framework for Forecast Evaluation. *Journal of Financial Economics* 19, 169-189.
- Dooley MP, Shafer JR, 1984. In: Bigman D., Taya T. (Eds.), *Floating Exchange Rates and the State of World Trade Developments*. Ballinger, Cambridge, pp. 43-69.
- Hodrick RJ, Srivastava S, 1986. The Covariation of Risk Premiums and Expected Future Spot Exchange Rates. *Journal of International Money and Finance* 5, S5-S21.
- Holland J, 1975. *Adaptation in Natural and Artificial Systems*. University of Michigan Press, Ann Arbor, MI.
- Kendall MG, Stuart A, 1958. *The Advanced Theory of Statistics*. Hafner Publishing, New York, NY.
- Koza JR, 1992. *Genetic Programming: On the Programming of Computers by Means of Natural Selection*. MIT Press, Cambridge, MA.
- Krugman P, Miller MH, 1993. Why Have a Target Zone? *Carnegie Rochester Conference Series on Public Policy* 38, 279-314.
- Lee CI, Mathur I, 1996. Trading Rule Profits in European Currency Spot Cross-Rates. *Journal of Banking and Finance* 20, 949-962.

- Levich R, Thomas L, 1993. The Significance of Technical Trading Rule Profits in the Foreign Exchange Market: A Bootstrap Approach. *Journal of International Money and Finance* 12, 451-474.
- Neely CJ, 1998. Target Zones and Conditional Volatility: The Role of Realignment. Federal Reserve Bank of St. Louis Working Paper 94-008C. Forthcoming in the *Journal of Empirical Finance*.
- Neely CJ, Weller P, 1997. Technical Analysis and Central Bank Intervention. Federal Reserve Bank of St. Louis Working Paper 97-002A.
- Neely CJ, Weller P, Dittmar R, 1997. Is Technical Analysis in the Foreign Exchange Market Profitable? A Genetic Programming Approach. *Journal of Financial and Quantitative Analysis* 32, 405-426.
- Sweeney RJ, 1986. Beating the foreign exchange market. *Journal of Finance* 41, 163-182.
- Sweeney RJ, Lee EJQ, 1990. In: Aggarwal R., and Lee C.F. (Eds.), *International dimensions of securities and currency markets. Advances in Financial Planning and Forecasting series vol. 4, part A.* JAI Press, Greenwich, CT, pp. 55-79.

Figure 1: An Arbitrary Rule




Figure 2: The Reproduction Operation




Figure 3: The Exchange Rate, Interest Rates and Cumulative Return from the DEM/FRF




Figure 4: Excess Returns to the Rules Around Periods of Realignment




Table 1: Descriptive statistics for the daily percentage returns including the interest rate differential, but excluding weekends and vacations: 1979:3:13 - 1996:06:21.

	FRF	ITL	NLG	GBP
N	4362	4358	4361	4356
Mean	0.0010	- 0.0030	0.0003	0.0003
SD	0.21	0.36	0.07	0.50
Skewness	- 8.98	- 2.59	- 1.15	- 0.71
Kurtosis	221.05	43.88	20.03	6.21
Min	- 5.80 ^a	- 5.42	- 1.13 ^b	- 4.97 ^c
Max	2.25	4.08	0.56	2.95 ^d

Notes: The first row gives the number of observations. The second and third rows give the mean and standard deviation of the daily percentage return. The mean return used here cannot be interpreted as the return to a long position because these data exclude the interest that would be earned on weekends and holidays. Skewness and kurtosis denote statistics that are standard normal if the data are normally distributed. See Kendall and Stuart (1958). Dates of extreme moves: a = realignment; b = day after realignment; c = exit from ERM; d = entry into ERM.

Table 2: Bandwidths during training, selection and validation periods for each currency

	Width of band	FRF	ITL	NLG	GBP
Width of band in training/selection		2.25%	6.00%	2.25%	Out
Subperiod 1 (S1)	2.25%	01/02/86 to 08/01/93	01/07/90 to 09/16/92	01/02/86 to 06/21/96	
	6.00%		01/02/86 to 01/06/90		10/08/90 to 09/16/92
Subperiod 2 (S2)	15%	08/02/93 to 06/21/96			
	Out		09/17/92 to 06/21/96		01/02/86 to 10/07/90 09/17/92 to 06/21/96

Notes: “Width of band” describes the permitted variation from central parity. The NLG was in a band of $\pm 2.25\%$ throughout the validation period. Under “width of band” the entry “out” means that the currency was not in the ERM for that period.

Table 3: Summary Statistics for Trading Rules

		FRF	ITL	NLG	GBP	
Panel A:	AR*100	0.86	2.45	0.06	2.75	
	t-statistic	1.06	1.16	0.35	1.18	
	post prob.	0.83	0.81	0.61	0.86	
	Entire Period	Sharpe ratio	0.33	0.35	0.10	0.33
		trades per year	7.68	5.30	6.76	6.51
		% XS	3.42	4.32	0.64	3.27
		% long	63.25	47.21	29.30	50.01
		# rules > 0	92	93	62	99
		MSD*100	0.74	2.03	0.16	2.37
# days in sample		3823	3823	3823	3823	
long return		1.03	0.53	0.28	-0.58	
		FRF	ITL	NLG	GBP	
Panel B:	AR*100	0.69	0.43	0.06	4.37	
	t-statistic	0.87	0.26	0.35	1.07	
	post prob.	0.80	0.47	0.61	0.86	
	S1	Sharpe ratio	0.32	0.18	0.10	0.97
		trades per year	7.23	6.25	6.76	6.23
		% XS	3.58	1.65	0.64	8.56
		% long	70.61	50.04	29.30	26.78
		# rules > 0	96	72	62	98
		MSD*100	0.62	0.68	0.16	1.30
# days in sample		2769	2450	3823	710	
long return		0.63	1.17	0.28	-4.67	
		FRF	ITL	NLG	GBP	
Panel C:	AR*100	1.31	6.04	NA	2.38	
	t-statistic	0.62	1.15	NA	0.88	
	post prob.	0.71	0.80	NA	0.79	
	S2	Sharpe ratio	0.40	0.61	NA	0.27
		trades per year	8.63	3.60	NA	6.61
		% XS	3.23	5.43	NA	2.60
		% long	43.91	42.16	NA	55.31
		# rules > 0	90	88	NA	99
		MSD*100	0.94	2.84	NA	2.50
# days in sample		1054	1373	NA	3113	
long return		2.08	-0.62	NA	0.36	

Notes: AR*100 is the mean annual return over the validation period for the 100 rules. The second row shows the average t-statistic, the mean of the 100 time-series t-statistics — calculated with a Newey-West correction for serial correlation — for the null hypothesis that each rule has a return equal to zero. Post prob is the Bayesian posterior probability that the mean excess return is greater than zero. The Sharpe ratio is the annual mean excess return divided by the annual standard deviation of the excess return. See equation (4) in the text for the definition of the variable % XS. % long is the percentage of the time the rule was long in the foreign (non-DEM) currency. The “# rules > 0” is the number of the 100 ex ante rules with positive mean validation period returns. MSD*100 is the monthly standard deviation in percentage terms. “# of days in sample” includes weekends and holidays. The long return is the mean annual return to a long position in the foreign (non-DEM) currency over the subsample. S1 and S2 refer to the subperiods identified in Table 2.

Table 4: Summary Statistics for Median and Uniform Portfolio Rules over the Entire Sample and Subperiods S1 and S2

		Median Portfolio Rule				Uniform Portfolio Rule			
Panel A:		FRF	ITL	NLG	GBP	FRF	ITL	NLG	GBP
Entire Period	AR*100	1.35	4.13	0.03	3.14	0.97	2.49	0.12	2.89
	t-statistic	1.69	1.99	0.15	1.32	1.59	2.04	0.81	1.37
	post prob.	0.96	0.98	0.55	0.91	0.93	0.98	0.81	0.90
	Sharpe ratio	0.53	0.65	0.05	0.37	0.50	0.65	0.27	0.38
	trades per year	4.87	4.39	5.82	1.24	5.67	4.53	5.73	3.90
	% XS	5.38	7.29	0.30	3.74	3.86	4.40	1.34	3.44
	% long	51.97	50.17	25.09	49.05	63.25	47.21	29.30	50.01
Panel B:		FRF	ITL	NLG	GBP	FRF	ITL	NLG	GBP
S1	AR*100	1.21	1.55	0.03	5.56	0.79	0.48	0.12	4.49
	t-statistic	1.54	1.03	0.15	1.36	1.18	0.52	0.81	1.18
	post prob.	0.94	0.58	0.56	0.92	0.90	0.31	0.81	0.91
	Sharpe ratio	0.56	0.64	0.05	1.23	0.43	0.38	0.27	1.07
	trades per year	5.01	5.36	5.82	1.54	5.35	5.35	5.73	3.69
	% XS	6.24	5.99	0.30	10.90	4.08	1.85	1.34	8.79
	% long	64.46	56.33	25.09	24.79	70.61	50.04	29.30	26.78
Panel C:		FRF	ITL	NLG	GBP	FRF	ITL	NLG	GBP
S2	AR*100	1.74	8.72	NA	2.59	1.45	6.08	NA	2.52
	t-statistic	0.82	1.70	NA	0.94	1.04	2.06	NA	1.03
	post prob.	0.76	0.95	NA	0.82	0.77	0.98	NA	0.85
	Sharpe ratio	0.55	1.10	NA	0.29	0.81	1.19	NA	0.32
	trades per year	4.16	2.66	NA	1.29	6.29	3.05	NA	4.00
	% XS	4.29	7.83	NA	2.83	3.58	5.46	NA	2.76
	% long	19.17	39.18	NA	54.58	43.91	42.16	NA	55.31

Notes: AR*100 is the mean annual return over the validation period for the 100 rules. The second row shows the Newey-West corrected t-statistic for the null hypothesis that each rule has a return equal to zero. Post prob is the Bayesian posterior probability that the excess return of the rule is greater than zero. The Sharpe ratio is the annual mean excess return divided by the annual standard deviation of the excess return. The posterior probability for the ITL uniform rule in period S1 is less than 0.5 despite the positive mean excess return because the $AR(p)$ process representing the series has a negative estimate for the mean. See equation (4) for the definition of the variable % XS. % long is the percentage of the time the rule was long in the foreign (non-DEM) currency.

Table 5: The Joint Significance of Portfolio Rule Returns

	Median	Uniform
statistic across all 4 exchange rates	7.624	8.538
p-value	0.106	0.074
statistic across all except DEM/NLG	7.585	8.452
p-value	0.055	0.038

Notes: The first two rows of the table present the test statistics and p-values of the joint significance of the median and uniform portfolio rules across all four exchange rates. The third and fourth rows present the same information for the test excluding the DEM/NLG.

Table 6: Summary Statistics for Uniform and Median Rule in Validation Period according to Bandwidth: DEM/ITL

Median Portfolio Rule

	Entire	6%	2.25%	Out
AR*100	4.13	5.34	-4.10	8.72
t-statistic	1.99	4.64	-1.25	1.70
post prob.	0.98	1.00	0.18	0.95
Sharpe ratio	0.65	2.26	-2.09	1.10
trades per year	4.39	6.72	3.34	2.66
% XS	7.29	21.09	-15.32	7.83
% long	50.17	87.46	9.87	39.18

Uniform Portfolio Rule

	Entire	6%	2.25%	Out
AR*100	2.49	2.59	-2.70	6.08
t-statistic	2.04	5.32	-1.24	2.06
post prob.	0.98	1.00	0.15	0.98
Sharpe ratio	0.65	2.41	-2.11	1.19
trades per year	4.53	6.11	4.22	3.05
% XS	4.40	10.22	-10.07	5.46
% long	47.21	67.85	23.51	42.16

Notes: The dates during the validation period when the currency was in each regime can be found in Table 2. The rows of the table are analogous to those in Table 4. See the notes to Table 4 for descriptions.

Table 7: Cumby-Modest Timing Tests for the Median Portfolio Rule - 1986-96.

	FRF	ITL	NLG	GBP
B	3.14	8.69	1.24	6.40
(s.e.)	(1.69)	(4.18)	(0.48)	(4.66)
p-value	0.06	0.04	0.01	0.17

Notes: B is the coefficient from a Cumby-Modest timing test regression of long returns to an asset on signals from a trading rule and (s.e.) denotes its standard error. No adjustment is made for transactions costs. A p-value less than .1 indicates that the test rejects the null hypothesis of no market timing ability for that exchange rate at the 10% level. The Wald test statistic for null of no predictability in all four exchange rates is 16.22, with a p-value of 0.003.

Table 8: Moving average rule results on ERM data

		Moving Average (1,10) Rule				Moving Average (1,50) Rule			
Panel A:		FRF	ITL	NLG	GBP	FRF	ITL	NLG	GBP
Entire Period	AR*100	-0.59	-0.40	-6.17	0.65	-0.10	4.22	-3.02	3.87
	t-statistic	-0.72	-0.18	-19.52	0.28	-0.11	1.98	-11.05	1.63
	post prob.	0.24	0.43	0.00	0.62	0.44	0.98	0.00	0.95
	Sharpe ratio	-0.20	-0.05	-5.14	0.08	-0.04	0.64	-2.58	0.45
	trades per year	43.26	48.90	76.97	43.74	15.28	16.33	36.10	14.61
	% XS	-2.34	-0.71	-71.07	0.77	-0.38	7.46	-34.76	4.61
	% long	50.95	48.78	52.42	49.02	45.17	38.77	55.98	44.81
Panel B:		FRF	ITL	NLG	GBP	FRF	ITL	NLG	GBP
S1	AR*100	-0.92	-1.93	-6.17	-1.07	0.65	0.14	-3.02	3.04
	t-statistic	-1.10	-1.24	-19.52	-0.26	0.82	0.09	-11.05	0.73
	post prob.	0.13	0.32	0.00	0.47	0.80	0.66	0.00	0.80
	Sharpe ratio	-0.38	-0.76	-5.14	-0.21	0.32	0.05	-2.58	0.65
	trades per year	44.31	52.31	76.97	42.73	13.45	21.16	36.10	12.87
	% XS	-4.76	-7.44	-71.07	-2.10	3.38	0.55	-34.76	5.96
	% long	49.55	49.51	52.42	41.97	42.36	37.18	55.98	40.00
Panel C:		FRF	ITL	NLG	GBP	FRF	ITL	NLG	GBP
S2	AR*100	0.28	2.31	NA	1.04	-2.06	11.50	NA	4.06
	t-statistic	0.14	0.41	NA	0.38	-0.91	2.19	NA	1.52
	post prob.	0.61	0.65	NA	0.64	0.18	0.98	NA	0.93
	Sharpe ratio	0.09	0.23	NA	0.14	-0.63	1.26	NA	0.46
	trades per year	40.56	42.83	NA	44.10	20.10	7.72	NA	14.90
	% XS	0.69	2.08	NA	1.14	-5.09	10.32	NA	4.44
	% long	54.65	47.49	NA	50.63	52.56	41.59	NA	45.90
		Moving Average (5,10) Rule				Moving Average (5,50) Rule			
Panel A:		FRF	ITL	NLG	GBP	FRF	ITL	NLG	GBP
Entire Period	AR*100	-1.07	-2.50	-2.52	-1.15	0.69	4.84	-0.91	3.41
	t-statistic	-1.26	-1.20	-11.62	-0.48	0.84	2.30	-4.67	1.44
	post prob.	0.09	0.12	0.00	0.34	0.79	0.99	0.00	0.92
	Sharpe ratio	-0.37	-0.34	-3.49	-0.13	0.27	0.73	-1.30	0.40
	trades per year	27.31	29.70	37.05	28.65	6.68	7.54	12.61	7.54
	% XS	-4.27	-4.42	-29.01	-1.37	2.75	8.55	-10.45	4.07
	% long	50.09	48.34	51.92	48.99	45.59	39.00	55.24	45.57
Panel B:		FRF	ITL	NLG	GBP	FRF	ITL	NLG	GBP
S1	AR*100	-0.49	-0.80	-2.52	-2.24	1.16	1.20	-0.91	3.29
	t-statistic	-0.60	-0.74	-11.62	-0.53	1.48	0.81	-4.67	0.79
	post prob.	0.27	0.13	0.00	0.44	0.93	0.79	0.00	0.79
	Sharpe ratio	-0.21	-0.37	-3.49	-0.44	0.62	0.51	-1.30	0.84
	trades per year	26.90	30.26	37.05	29.34	5.80	8.64	12.61	8.24
	% XS	-2.53	-3.09	-29.01	-4.39	6.00	4.63	-10.45	6.44
	% long	49.12	49.22	51.92	42.54	42.94	37.02	55.24	39.58
Panel C:		FRF	ITL	NLG	GBP	FRF	ITL	NLG	GBP
S2	AR*100	-2.61	-5.54	NA	-0.90	-0.53	11.33	NA	3.44
	t-statistic	-1.14	-1.01	NA	-0.33	-0.24	2.18	NA	1.28
	post prob.	0.14	0.15	NA	0.38	0.42	0.98	NA	0.89
	Sharpe ratio	-0.80	-0.54	NA	-0.10	-0.17	1.21	NA	0.40
	trades per year	28.42	28.73	NA	28.38	9.01	5.59	NA	7.39
	% XS	-6.43	-4.97	NA	-0.99	-1.31	10.17	NA	3.76
	% long	52.66	46.76	NA	50.47	52.56	42.53	NA	46.93

Notes: The moving average rule (x, y) gives a long signal when the x -day moving average exceeds the y -day moving average, and a short signal otherwise. The rows of the table are analogous to those in Table 4. See the notes to Table 4 for descriptions.

Table 9: Filter rule results on ERM data

		Filter 0.5% Rule				Filter 1.0% Rule			
Panel A:		FRF	ITL	NLG	GBP	FRF	ITL	NLG	GBP
Entire Period	AR*100	-1.05	0.72	-0.29	-2.59	-0.66	-0.25	-0.29	-1.14
	t-statistic	-1.26	0.35	-1.74	-1.09	-0.84	-0.13	-1.74	-0.49
	post prob.	0.12	0.66	0.03	0.13	0.18	0.43	0.03	0.32
	Sharpe ratio	-0.38	0.11	-0.57	-0.39	-0.26	-0.04	-0.57	-0.16
	trades per year	5.44	20.34	0.00	39.35	1.53	9.84	0.00	17.57
	% XS	-4.17	1.27	-3.33	-3.08	-2.64	-0.45	-3.33	-1.35
	% long	34.87	39.68	0.00	51.79	26.21	39.52	0.00	49.67
Panel B:		FRF	ITL	NLG	GBP	FRF	ITL	NLG	GBP
S1	AR*100	-0.83	-1.25	-0.29	-1.74	-0.55	-1.77	-0.29	-1.99
	t-statistic	-1.03	-0.81	-1.74	-0.44	-0.71	-1.67	-1.74	-0.48
	post prob.	0.15	0.46	0.02	0.39	0.25	0.01	0.03	0.35
	Sharpe ratio	-0.38	-0.54	-0.57	-0.34	-0.25	-0.76	-0.57	-0.52
	trades per year	1.85	5.96	0.00	23.17	0.26	0.30	0.00	7.72
	% XS	-4.31	-4.83	-3.33	-3.41	-2.86	-6.82	-3.33	-3.89
	% long	32.21	31.43	0.00	46.34	10.55	30.53	0.00	24.79
Panel C:		FRF	ITL	NLG	GBP	FRF	ITL	NLG	GBP
S2	AR*100	-1.62	4.23	NA	-2.78	-0.96	2.44	NA	-0.94
	t-statistic	-0.72	0.83	NA	-1.00	-0.46	0.47	NA	-0.35
	post prob.	0.22	0.80	NA	0.17	0.34	0.67	NA	0.37
	Sharpe ratio	-0.44	0.50	NA	-0.38	-0.29	0.23	NA	-0.12
	trades per year	14.91	46.02	NA	43.16	4.85	26.87	NA	19.94
	% XS	-4.00	3.80	NA	-3.04	-2.36	2.19	NA	-1.03
	% long	41.84	54.41	NA	53.04	67.36	55.57	NA	55.35
Panel A:		FRF	ITL	NLG	GBP	FRF	ITL	NLG	GBP
Entire Period	AR*100	-2.20	3.86	-0.29	-1.69	-0.13	2.26	-0.29	0.98
	t-statistic	-2.53	1.86	-1.74	-0.77	-0.16	1.08	-1.74	0.43
	post prob.	0.01	0.98	0.03	0.23	0.44	0.86	0.03	0.66
	Sharpe ratio	-0.81	0.55	-0.57	-0.24	-0.05	0.33	-0.57	0.15
	trades per year	0.95	4.11	0.00	7.93	0.10	2.77	0.00	2.86
	% XS	-8.74	6.82	-3.33	-2.02	-0.52	4.00	-3.33	1.17
	% long	24.40	18.57	0.00	47.37	10.62	16.77	0.00	60.79
Panel B:		FRF	ITL	NLG	GBP	FRF	ITL	NLG	GBP
S1	AR*100	-1.23	-1.18	-0.29	-7.38	-0.64	-1.18	-0.29	-4.70
	t-statistic	-1.54	-0.79	-1.74	-1.84	-0.81	-0.79	-1.74	-1.15
	post prob.	0.06	0.46	0.04	0.03	0.23	0.46	0.03	0.11
	Sharpe ratio	-0.56	-0.49	-0.57	-1.89	-0.29	-0.49	-0.57	-1.01
	trades per year	0.13	0.00	0.00	1.03	0.00	0.00	0.00	0.00
	% XS	-6.34	-4.54	-3.33	-14.46	-3.29	-4.54	-3.33	-9.20
	% long	11.12	0.00	0.00	76.90	0.00	0.00	0.00	100.00
Panel C:		FRF	ITL	NLG	GBP	FRF	ITL	NLG	GBP
S2	AR*100	-4.77	12.84	NA	-0.40	1.20	8.40	NA	2.28
	t-statistic	-2.06	2.54	NA	-0.15	0.55	1.63	NA	0.86
	post prob.	0.03	0.99	NA	0.44	0.73	0.93	NA	0.82
	Sharpe ratio	-1.28	1.62	NA	-0.05	0.34	0.87	NA	0.27
	trades per year	2.77	11.44	NA	9.27	0.35	7.72	NA	3.28
	% XS	-11.75	11.52	NA	-0.43	2.95	7.54	NA	2.49
	% long	59.30	51.71	NA	40.64	38.52	46.69	NA	51.85

Notes: The filter rule with filter x gives a long signal when the current exchange rate has risen by more than $100*x$ percent from its minimum over the last five days. It gives a short signal if the exchange rate has fallen by more than $100*x$ percent from its maximum over the last five days. Otherwise the current position is maintained. For the filter rules average return over the whole sample will not necessarily be equal to a weighted average of returns over the two subsamples, because the initial position in each subsample was arbitrarily chosen to be a long position. The number of trades in the two subsamples will not always sum exactly to the total number of trades because of these initial conditions. The rows of the table are analogous to those in Table 4. See the notes to Table 4 for descriptions

Table 10: USD/DEM trading rules run over EMS data: 1986-96

		FRF	ITL	NLG	GBP
Panel A: Entire Period	AR*100	-1.06	1.39	-1.27	3.01
	t-statistic	-1.28	0.66	-6.49	1.27
	post prob.	0.13	0.74	0.04	0.87
	Sharpe ratio	-0.40	0.19	-1.89	0.36
	trades per year	6.24	4.18	17.99	7.45
	% XS	-4.20	2.46	-14.60	3.58
	% long	38.53	23.55	51.43	41.78
	# rules > 0	5	93	4	95
	MSD*100	0.77	2.15	0.19	2.38
	# days in sample	3823	3823	3823	3823
long return	1.03	0.53	0.28	-0.58	
		FRF	ITL	NLG	GBP
Panel B: S1	AR*100	-0.68	-0.21	-1.27	0.39
	t-statistic	-0.86	-0.12	-6.49	0.09
	post prob.	0.21	0.60	0.04	0.57
	Sharpe ratio	-0.31	-0.09	-1.89	0.08
	trades per year	5.88	3.72	17.99	11.07
	% XS	-3.53	-0.82	-14.60	0.76
	% long	32.26	23.05	51.43	46.90
	# rules > 0	5	6	4	73
	MSD*100	0.63	0.69	0.19	1.36
	# days in sample	2769	2450	3823	710
long return	0.63	1.17	0.28	-4.67	
		FRF	ITL	NLG	GBP
Panel C: S2	AR*100	-2.05	4.26	NA	3.61
	t-statistic	-0.93	0.81	NA	1.32
	post prob.	0.21	0.76	NA	0.89
	Sharpe ratio	-0.73	0.39	NA	0.42
	trades per year	7.19	4.99	NA	6.72
	% XS	-5.05	3.82	NA	3.94
	% long	55.02	24.44	NA	40.62
	# rules > 0	6	91	NA	99
	MSD*100	0.81	3.13	NA	2.47
	# days in sample	1054	1373	NA	3113
long return	2.08	-0.62	NA	0.36	

Notes: 100 trading rules were obtained using data for the USD/DEM. The training period was 1975-77, and the selection period 1978-80. For more details see NWD (1997). The ERM exchange rate data were normalized by dividing them by a one-sided 250-day moving average. For explanation of the statistics, see notes to Table 3.

Table 11: Summary Table for “High Interest Rate” and “Mean Reversion” Trading Rules

		Mean Reversion Rule				High Interest Rate Rule			
Panel A:		FRF	ITL	NLG	GBP	FRF	ITL	NLG	GBP
Entire Period	AR*100	0.90	-1.02	0.10	-2.64	1.10	0.70	-0.51	-0.74
	t-statistic	1.11	-0.52	0.61	-1.13	1.35	0.33	-2.55	-0.31
	post prob.	0.86	0.32	0.79	0.14	0.92	0.63	0.01	0.40
	Sharpe ratio	0.35	-0.15	0.21	-0.32	0.43	0.09	-0.72	-0.09
	trades per year	2.86	5.06	14.13	5.35	3.25	1.72	23.20	17.76
	% XS	3.56	-1.80	1.15	-3.14	4.36	1.23	-5.90	-0.88
	% long	85.83	68.88	38.96	53.14	97.99	99.16	60.51	80.47
Panel B:		FRF	ITL	NLG	GBP	FRF	ITL	NLG	GBP
S1	AR*100	0.44	0.89	0.10	-1.50	0.90	1.44	-0.51	-5.25
	t-statistic	0.57	0.88	0.61	-0.37	1.16	0.95	-2.55	-1.21
	post prob.	0.72	0.85	0.80	0.32	0.87	0.57	0.01	0.13
	Sharpe ratio	0.20	0.38	0.21	-0.34	0.43	0.60	-0.72	-1.30
	trades per year	3.96	4.77	14.13	6.69	2.64	2.68	23.21	17.50
	% XS	2.30	3.45	1.15	-2.95	4.66	5.55	-5.90	-10.28
	% long	80.43	62.08	38.97	83.94	97.76	98.69	60.53	92.11
Panel C:		FRF	ITL	NLG	GBP	FRF	ITL	NLG	GBP
S2	AR*100	2.08	-4.43	NA	-2.89	1.62	-0.62	NA	0.29
	t-statistic	0.97	-0.86	NA	-1.08	0.75	-0.12	NA	0.11
	post prob.	0.84	0.20	NA	0.15	0.75	0.44	NA	0.54
	Sharpe ratio	0.68	-0.42	NA	-0.35	0.51	-0.06	NA	0.03
	trades per year	0.00	5.59	NA	5.16	4.85	0.00	NA	17.83
	% XS	5.05	-3.94	NA	-4.69	3.93	-0.55	NA	0.47
	% long	100.00	81.06	NA	46.10	98.58	100.00	NA	77.80

Notes: The “mean reversion” strategy takes a long position in the foreign currency if the spot price of foreign currency is below the actual or notional central parity. The “high interest rate” rule takes a long position in the currency with the higher interest rate each day. The rows of the table are analogous to those in Table 4. See the notes to Table 4 for descriptions.

Table 12: CAPM Regression Betas

		FRF	ITL	NLG	GBP
World Index	constant	1.789	3.721	-0.052	4.361
	(s.e.)	(0.785)	(2.041)	(0.163)	(2.327)
	B	-0.013	0.026	0.000	0.009
	(s.e.)	(0.017)	(0.044)	(0.004)	(0.050)
Commerzbank Index	constant	1.381	4.009	0.006	3.381
	(s.e.)	(0.800)	(1.991)	(0.168)	(2.700)
	B	-0.000	0.036	0.000	-0.002
	(s.e.)	(0.012)	(0.031)	(0.003)	(0.042)

Notes: The coefficients are those from a regression of monthly median portfolio rule excess returns on a constant and monthly excess returns for the benchmark over 7/29/86 to 5/21/96 for the MSCI World Index and 1/1/86 to 5/21/96 for the Commerzbank Index. Results from a regression with the constant constrained to equal zero also showed no evidence of systematic risk.

Table 13: Risk Adjustment with the X* Statistic

Panel A:		Median Portfolio Rule				Uniform Portfolio Rule			
		FRF	ITL	NLG	GBP	FRF	ITL	NLG	GBP
Entire Period	X*	1.32	4.12	0.18	3.14	0.70	2.52	0.24	2.89
	std err	0.79	2.08	0.16	2.38	0.50	1.24	0.13	2.12
	post prob	0.96	0.98	0.86	0.92	0.92	0.98	0.96	0.91
Panel B:		FRF	ITL	NLG	GBP	FRF	ITL	NLG	GBP
S1	X*	1.02	1.40	0.18	3.21	0.53	0.48	0.24	2.35
	std err	0.66	1.38	0.16	3.30	0.45	0.92	0.13	3.04
	post prob	0.95	0.67	0.84	0.86	0.88	0.31	0.96	0.80
Panel C:		FRF	ITL	NLG	GBP	FRF	ITL	NLG	GBP
S2	X*	3.08	8.59	NA	2.56	1.73	5.99	NA	2.49
	std err	2.54	5.34	NA	2.74	1.53	3.09	NA	2.43
	post prob	0.82	0.94	NA	0.83	0.81	0.98	NA	0.85

Notes: The standard error for X* is a generalization of Sweeney and Lee's (1990) measure which takes into account the possibility of serial correlation in returns. Post prob is the Bayesian posterior probability that X* is greater than zero.

Table 14: Characteristics of the optimal portfolio formed from MSCI World Index and median portfolio trading rules: 1988-96.

	FRF	ITL	GBP
Excess return	9.73	7.21	9.63
Portfolio weights trading rule	4.83	1.69	1.80
MSCI World	0.22	0.22	0.22

Notes: The calculations were performed over the period *01/01/88* to *05/21/96*. The first row of the table reports the annual percentage excess return from an optimally levered portfolio formed from benchmark MSCI World Index and currency trading rule, constructed to match the variance of the benchmark. The second and third rows of the table report the portfolio weights in the optimally levered portfolio.