

Do Technical Trading Rules Generate Profits? Conclusions from the Intra-Day Foreign Exchange Market

Riccardo Curcio*

Charles Goodhart**

Dominique Guillaume***

and

Richard Payne****

June 1997

Abstract. Several recent studies have demonstrated the profitability of technical analysis by simulating certain trading rules over a very long period of daily foreign exchange rates. In this paper, we use filter rules identified and supplied by technical analysts on the intra-daily foreign exchange market. We provide evidence that, although some profits could be made by following these rules in periods of trends, this was not the case on average. Our results are further strengthened when we incorporate transaction costs. We also simulate some of the rules used in previous studies and show that they would not be profitable when applied to our intra-daily data set.

Keywords: technical analysis, trading rules, foreign exchange markets.

—
* J.P. Morgan; ** London School of Economics, Department of Economics and Financial Markets Group; *** University of Oxford, C.S.A.E. - Institute of Economics and Statistics and London School of Economics, Financial Markets Group; **** London School of Economics, Department of Accounting and Finance and Financial Markets Group.

Any opinions expressed are those of the authors and not necessarily those of J.P. Morgan or the Financial Markets Group. Charles Goodhart and Richard Payne wish to thank the E.S.R.C. for financial support.

Dominique Guillaume is a Marie-Curie E.C. fellow. He also wishes to thank the E.S.R.C. for financial support.

1. Introduction

Ever since Dow created the first market indicators in the late 1800s, technical analysis has been at the centre of a debate between professionals operating in financial markets and academic economists. The latter argue that with a turnover of more than 1 trillion US dollars traded around the world every day, it is highly improbable that the foreign exchange market is not at least weakly efficient. Systematic opportunities to make profits based the analysis of past prices should be rapidly eliminated by arbitrage. More specifically, if a group of traders were to base their expectations on the extrapolation of past price trends, thereby reinforcing them, rational traders could always take a profitable speculative position based on the fundamental value, to which the exchange rate should ultimately return. In the long run, these rational traders would drive the others out of the market.

Recently, however, several researchers have overcome the natural skepticism of academics towards technical analysis in order to investigate its usefulness. Taylor and Allen (1992) surveyed more than 200 traders on the London foreign exchange market - by far the largest - and found that the use of technical analysis increases with the frequency of trading. In particular, chartism is the main forecasting tool of intra-daily traders. Several other studies, including Brock et al. (1992) and Levich and Thomas (1993), have tested the potential profitability of some technical rules. Using bootstrap methods to evaluate the significance of their results, they usually found that, with daily observations, the rules they simulated would have generated some profits. In addition to providing some rationale for the actual use of technical analysis in the markets, these results could be explained on the following grounds. First, from a statistical perspective, the existence of profit-making rules might be explained by the more complex, nonlinear dynamics observed in foreign exchange rates (Hsieh, 1989; Guillaume et al., 1995). Second, it is not necessarily true that sufficient fundamental traders willing or able to take speculative positions against those using technical analysis might exist. It might

well be that traders do not have sufficient liquidity or are not permitted by their institutions to take open positions for the long periods which could be needed in the event that exchange rates wander away from their fundamental values for long periods of time (I.M.F., 1993). Third, the existence of other types of traders, such as central bankers, who have quite different objectives, may be at the origin of inefficiencies in the foreign exchange markets (LeBaron, 1996).

However, these studies share a number of drawbacks. First, it is yet to be shown that the profits exhibited would remain if the tests were conducted on time periods shorter than a couple of years. But, traders cannot afford to make losses during a couple of months or years even if in the extremely long run their strategies make profits. Moreover, such long periods would be in contradiction with the study of Taylor and Allen (1992) referred to above, where it was shown that technical analysis was mainly used in the very short run. Second, and using the same line of reasoning, those studies make use of daily observations whereas the bulk of foreign exchange trading - more than 75 % according to the 1996 Bank for International Settlements survey - takes place within the day; that is, between dealers and/or brokers. Third, some of these studies such as the one by Brock et al. (1992) do not incorporate transaction costs in the computation of their profits. Finally, these studies involve ex-post simulations of trading rules potentially used by traders. There is however no evidence that these rules are the ones actually used by traders. They represent only a few of a very large set of possible trading rules.

Goodhart and Curcio (1993) tried to meet this last criticism by running the following experiment. They took some students whom they divided into two groups, the first group having to trade with no external technical advice whereas the second set could rely upon the service of an existing chartist program. Although the second group did not make excess profits, they had a considerably lower variance of profits. Although it yielded interesting insights, this experiment was only a single shot. In the present paper, we deal with the above short-comings by using buy and sell signals derived from technical rules reported by traders to Reuters. We then investigate whether intra-day trading according

to these technical indicators would yield profits. For comparative purposes, we also simulate some of the rules reported in previous studies using our intra-daily data set. Furthermore, the computation of profits is corrected to take account of transaction costs. Our main result is that even when transaction costs are not taken into account, it is *on average* not profitable to trade according to these rules, although excess returns can be made in periods of strongly trending exchange rates.

The remainder of this paper is divided as follows. In the next section, details on the data and methodology are given. We report our results in Section 3 and give some concluding remarks in Section 4.

2. Data and Methodology

2.1 The Data

Our data consists of two samples, the first covering the period from 10 April 1989 to 29 June 1989 and the second running from 31 January 1994 to 30 June 1994. For each sample we have a tick-by-tick FFX quotation series for the Mark (DEM), Yen (JPY) and Pound (GBP) against the US dollar (USD) plus observations on support-resistance and HIGH-LOW trading ranges downloaded from the Reuters FXNL screen. For the first sample we have approximately three FXNL observations per day whilst for the 1994 sample there is at most one observation per day. These trading ranges are constructed by Reuters from a daily survey of a number of established participants in the foreign exchange markets and represent bands within which the currency is expected to trade given information available to those participants at the beginning of the trading day.¹

Using the tick-by-tick quotations we construct an intra-daily bid, ask and midpoint quotation series for each of the three currencies, sampled at a fixed 1 hour calendar time interval.² We chose a 1-hour frequency rather than say, a tick-by-tick or 5 minutes frequency in order to avoid the price uncertainty due to the fact that our prices are quoted and not transaction prices.³ This gives a time-series length for each currency of around 1400 return observations for the first sample and over 2500 return observations for the second sample. These intra-day quotation series form the basis for our empirical trading rule applications.

We then combine the exchange rate data with the daily trading range data. For the 1989 sample this is straightforward. The downloaded trading range data is time-stamped to the

¹ Note that, although it could be argued that traders have no incentive to report their best trading range estimates, they also know that their estimates will be diluted in the aggregation of responses implying incentives to misreport are likely to be diminished or removed.

² Our analysis excludes weekends, defined as midnight on Friday to midnight Sunday GMT.

³ See Guillaume et al. (1997) for a discussion of this issue.

second and can hence be merged perfectly with the quotation series to give a single data set of contemporaneous trading range and quotation series.

For the 1994 sample, however, there are two minor problems with the trading range data. The first is that observations on all three currencies are not available for every trading day in our sample. This leaves us with around 90 daily observations for each currency (out of a possible maximum number of observations of approximately 110.) Second, the time at which the support-resistance and HIGH-LOW data was reported by Reuters is not specified in this data set.⁴ Fortunately, though, the FXNL page which shows the trading range data also presents the prevailing spot prices for all of the selected currencies. This allows us to time the 1994 trading range data by matching the spot prices from the FXNL screen to those from our tick-by-tick FAFX time series, a match being defined as a five minute interval in which all of the FXNL spot prices were simultaneously valid for the individual FAFX time-series.⁵ Once this was achieved the quotation and trading range data were merged as for the 1989 sample.

Note that we assume the trading range data to be valid only until midnight on the day on which it was observed in order to avoid generating spurious buy/sell signals from stale trading range data i.e. for all hours between midnight on a given day and the observation of the next days support-resistance data the trading signal process, s_t , is restricted to be zero. This can be justified in the following ways. First, intra-day foreign exchange traders are subject to strict overnight limits on their positions; in fact they are usually expected to completely close out all positions. The restriction we place on the validity of the trading range data implicitly reflects these limits. Note however that by closing our trader's position at midnight, we implicitly allowed him to roll over his position to his

⁴ Note that the trading range observations for all three currencies are simultaneously reported by Reuters.

⁵ In some cases there were multiple matches for a given daily observation. We countered this problem by running all our trading rules using daily trading range data defined from the first match-time in each day, the average and finally the last. Our results were almost identical across these three definitions and we present results based on the first possible match in the rest of this paper.

overseas counterpart when both markets business hours overlap.⁶ Second, a trader wishing to maintain an open position derived from a trading rule overnight will find it more costly than maintaining an intra-day position as he will have to pay overnight interest. This again implies it is more sensible to ensure that no positions can be left open overnight in our analysis.

2.2 The Basic Trading Rule

Define q_t to be the quote midpoint series for the exchange rate under consideration. As mentioned above, in our empirical analyses we sample the exchange rate series at a one hour frequency. Define r_t to be the percentage return on the exchange rate, calculated as the first difference of the logarithmic midquotes.

All of the technical rules we employ are based on the movement of the exchange rate outside a pre-defined trading range. We present results based on four classes of trading rule definition. A first range is derived from the support and resistance levels which appear on the Reuters FXNL screen. A second possibility is to use the HIGH-LOW data which also appears on the FXNL screen. We define a third range by using both the support-resistance and HIGH-LOW data, taking the minimum of the support and LOW data and the maximum of the resistance and HIGH figures as the range. Finally, we follow Brock, Lakonishok and LeBaron (1992) in order to construct a fourth type of trading range. This range (denoted Max-Min) is calculated, at each hourly observation point, as the local maximum and minimum of the exchange rate based on a given window of past hourly observations.⁷

⁶ Positions could be rolled over by a Japanese trader to his European counterpart or by an European trader to his American counterpart. We closed the position at 12.00 when there is no overlap between the American and Japanese markets.

⁷ Again following Brock, Lakonishok and LeBaron (1992) the horizons we employ are 50, 150 and 200 observations.

Once the trading range is defined, we can proceed to the application of the technical trading rule. Defining h_t to be the upper bound of the range and l_t to be the lower bound of the range, signals to buy and sell currency are generated according to the rule below,

$$s_t = \begin{cases} +1 & \text{if } q_t > h_t \\ 0 & \text{if } l_t < q_t < h_t \\ -1 & \text{if } q_t < l_t \end{cases} \quad (1)$$

where s_t represents the signal process which is unity for a buy signal and negative one for a sell signal. Note that a purchase (sale) will be maintained until the spot rate moves back in side the upper and lower bounds (hence, the position is closed without making a profit), or until midnight on the same day, (when it is assumed that all positions are closed out for institutional reasons as in the previous discussion,) whichever is the sooner. In the first sample however, we have several updates of the lower and upper bound per day, so that profits can also be made when the exchange rates is not trending (if, for example, the upper bound is adjusted to a higher range and crossed again; in this case, the position will be closed but with a profit). Dynamic percentage returns from following this rule are simply calculated as below.

$$\pi_t = s_t (\log(q_{t+1}) - \log(q_t)) = s_t r_{t+1} \quad (2)$$

Figure 1 gives an example of how such trading rule works.

In Section 3 we present data from the application of this rule to data on the DEM/USD, JPY/USD and USD/GBP. For each exchange rate, signals are generated according to the alternative range definitions and the mean return, number of buy/sell signals and a t -statistic for the significance of the mean rule return over the drift in the exchange rate return series are tabulated.⁸ The t -statistic represents the significance of the trading rule return over that from a ‘buy and hold’ strategy and is calculated as,

⁸ It is well known that the distribution of exchange rate returns deviates fairly significantly from the normal. This implies that a more appropriate way of testing the significance of trading rule returns would

$$t = \frac{\mu_T - \mu}{\left(\frac{\sigma^2}{N} + \frac{\sigma^2}{N_T}\right)^{\frac{1}{2}}} \quad (3)$$

where μ_T is the mean trading rule return, μ is the mean exchange rate return, N_T is the number of trading rule signals, N is the number of exchange rate observations and σ^2 is the exchange rate return variance. We also present results for a further modification of the range definitions. Indeed, traders usually modify the trading range by adding a band of variable size to the predefined range. This is the reason why this class of technical trading rules are often called filter rules. In this paper, we add a 0.01% band to each range definition and rerun the rules. Note that like the choice of the trading range, the choice of the band is necessarily arbitrary. However, rather than the value of the band, the crucial element in a filter rule is the choice of the trading range, which is the focus of this paper.

2.3 Accounting for Transaction Costs

The rule described above clearly takes account of neither the costs of opening and closing positions nor of the risk incurred by taking such a position. Since our trader closes his position at the latest at 12.00 p.m. G.M.T., he bears no interest costs on open positions. However, given the frequency at which one will change position in an intra-day analysis, trading costs due to bid-ask spreads are unlikely to be negligible. We (rather than using the midquote data as in the previous subsection) account for them by using the bid and ask quotation series. Defining a_t to be the prevailing ask quotation and b_t the bid quotation, our trading rule is redefined as,

be to use a bootstrap, as in Brock, Lakonishok and LeBaron (1992). However the results in Brock, Lakonishok and LeBaron (1992) are not qualitatively altered by using bootstrapped standard errors and hence we focus on traditional t-statistics to provide statistical inference.

$$s_t = \begin{cases} +1 & \text{if } a_t > h_t \\ 0 & \text{if } l_t < b_t, a_t < h_t \\ -1 & \text{if } b_t < l_t \end{cases} \quad (4)$$

such that a long position is established if the ask rises above the upper bound of the range and a short position taken on if the bid is below the lower bound. The real incorporation of transactions costs comes in the definition of the trading return, however. If one has, for example, taken a long position, the price at which currency was purchased was the market ask price. In closing out the position one sells back to the market, implying that one will receive the market bid. Hence, in opening and closing each position, the trader pays the quoted spread to the market maker as a transaction cost.⁹ If the trading rule then determines that a position should be opened, held open for k periods and then closed, the total holding period return is then,

$$\pi_t = \begin{cases} +(\log(b_{t+k}) - \log(a_t)) & \text{if } s_{t+j} = +1, \quad \forall j = 0, \dots, k-1 \\ -(\log(a_{t+k}) - \log(b_t)) & \text{if } s_{t+j} = -1 \quad \forall j = 0, \dots, k-1 \end{cases} \quad (5)$$

and per-period returns are based on the above profit calculation. Results from the specifications of the trading rule which incorporate transaction costs are given in the latter half of Section 3.

⁹ Note that the bid and ask data we employ are taken from Reuters FXFX screens. It has become a well recognized fact that the spreads quoted on these screens bracket the true inside spreads in the inter-dealer market, implying that our measure of transactions costs might be slightly overstated.

3. Results

In Table 1 summary statistics are given for our two samples (hereafter referred to as ‘a’ and ‘b’). Also, in Figures 2 to 7 we present plots of the hourly quotation series for the three currencies, for both of the sub-samples. Whilst in the second sub-sample there seems to be little evidence of global trends in the quotation series, the figures demonstrate that, for all three currencies, the first sub-sample series are subject to strong and significant trends. This observation is statistically reinforced by the estimated trends presented in the opening tables. Those for sample (a) are clearly greater than for their sample (b) counterparts.

Throughout the following section, for reasons of space, we present tables of results only for the JPY/USD exchange rate.¹⁰ Within the text, however, we discuss the broad results from the application of our trading rules to all three exchange rates.

3.1 Brief Summary of Results from Sample (a)

The first sample of intra-day exchange rates and technical rules was studied in Curcio and Goodhart (1992). Here we present a summary of their findings. Readers wanting a more complete description of the results from this sample should refer directly to Curcio and Goodhart (1992).

Results for the JPY/USD are presented in Tables 2 and 3 at the end of the paper. Examining the returns made from the application of the support-resistance and HIGH-LOW rules to these data (Table 2) demonstrates that in the majority of cases they are positive and are generally an order of magnitude greater than the returns from a buy and hold strategy. Examining the significance of the buy and sell sides separately, one finds that in 18 of the 36 cases the rules generate significant profit. Further, amalgamating the

¹⁰ Results for the other two currency pairs are available from the authors upon request.

buy and sell sides for each rule, of the 18 returns across the three currencies, 12 are significantly greater than zero.

Similar results obtain for the application of the Max-Min rules proposed by Brock, Lakonishok and LeBaron (1992) (presented for the JPY/USD in Table 3), although they are slightly weaker. Profits are positive for only 14 of the 36 separate buy and sell rules, and the composite buy-sell return is significantly greater than zero in 8 of the 18 cases.

Hence, it seems that there are rules which, when applied to our first intra-day exchange rate sample, could have made consistently positive profits. A more careful examination of the results, however, demonstrates that the significance of these rules is entirely asymmetric. Examining the support-resistance results, for the DEM/USD and JPY/USD all buy side rules are significant and none of the sell rules are, whilst the opposite holds for the USD/GBP. A similar conclusion holds when one examines the Max-Min results, with the sells being the only significantly profitable rules for the USD/GBP and the buys for the DEM/USD and JPY/USD.

These results can be linked to the trend in our exchange rate sample. As mentioned earlier, sample (a) covers a period of sustained USD appreciation, implying a strong upward trend in the DEM/USD and JPY/USD whilst the USD/GBP trended downward over our sample. As the filter rules which we employ are essentially trend-following rules, the significance and asymmetry of our results may be attributed to the strong trends in our exchange rate sample. This behaviour can also be seen in the number of signals generated. For the JPY/USD, for example, Table 2 demonstrates that there are far more buys than sells, a result which also holds for the DEM/USD whilst the converse is true for the USD/GBP. This, again, can be seen as indicative of the trend-chasing nature of our trading rules.

A further qualification to our result that application of these rules generates excess returns is that the analysis thus far does not account for transaction costs. As described in

Section 2, we embody these costs by incorporating the bid-ask spread into our return calculations. The results for the JPY/USD which correspond to our rules which include transaction costs are given in the final three columns of Tables 2 and 3. A general result which emerges from examination of these tables is that, once spreads are accounted for, virtually all trading rule profits are eliminated. For the Reuters trading range data, only 4 of the 36 buy or sell rules yield returns which are significantly greater than zero at conventional levels. Further, only 2 of the 18 composite bid and ask rules have significantly positive returns. As before, the returns from the Max-Min rules are poorer than the returns from the Reuters based rules; none are significantly greater than zero, whilst 7 of the 36 separate bid and ask rules and two of the composite rules yield significant losses.

3.2 Results from Sample (b)

Examining results for our second sample of exchange rates and technical rules gives a far more dismal picture of the profitability of these techniques. In comparison to those from sample (a), the results are sketchy and inconsistent. A source of this difference is likely to be the difference between the strength of the exchange rate trends across the two samples.

Specific observations on the application of the support-resistance and HIGH-LOW rules are as follows (see Table 4 for JPY/USD results). In only the USD/GBP case is there any indication of more signals being generated for one side of the market than for the other (in this case buy signals seem to predominate.) Only 15 of the 36 separate buy and sell returns are positive and of these only 4 are significant. In fact, there are almost as many significantly negative returns (3 of 36) as there are significantly positive returns. Further, only 2 of the 18 composite buy and sell rules yield significant profits. Again the results from the application of the Max-Min rules are weaker (see Table 5). Only 9 rules generate positive returns and there are far more significantly negative returns (7) than

significantly positive (2). This leads to the result that only one composite return yields excess returns significantly greater than zero, whilst 4 of the 18 yield significant losses.

For sample (b) all the significant profits accrue from the application of technical rules to the JPY/USD exchange rate. As can be seen from Figure 6, the JPY was subject to two sharp periods of appreciation, at the beginning and end of the sample. In line with the results from sample (a) we hypothesize that these episodes might be the source of the positive profits and test this hypothesis below.

Our test consists simply of dividing sample (b) into periods in which strong trends pertain in the exchange rates and periods of no trend. This type of exercise is far less easy for sample (a) where it seems that all three exchange rates are trended over the entire period. This is also true, to a lesser extent, for the DEM/USD in sample (b). The division leaves us with a ‘flat’ and ‘trended’ sub-sample for each of the JPY/USD and USD/GBP in sample (b).¹¹ Unfortunately, the ‘trend’ sub-sample for the USD/GBP contained too few signals for reliable statistical inference and we hence present tables which only cover the JPY/USD sub-samples.

Results from the application of the support-resistance rules to these sub-samples are shown in Tables 6 and 7 and they immediately bear out our intuition regarding the source of trading rule profits. The ‘flat’ JPY/USD sub-sample shows all buy and sell returns to be negative with, obviously, no significant trading rule profits. For the ‘trend’ sub-sample though, the appreciation of the JPY can be seen to generate significant profits for the sell-side rules. Hence it seems that these rules only generate consistent profits when there are strong trends in exchange rates. If one therefore desired to apply them practically, one would not only have to know the correct rule to implement, but also when to apply the rule in terms of the start points of possible exchange rate trends.

¹¹ The trended JPY/USD sub-samples cover observations 1-250 and 2200-2663 approximately. The trended USD/GBP sub-sample covers observations 1750-2663.

Finally, incorporating transaction costs into the analysis inevitably makes the results even more convincing of market efficiency. Simply incorporating the spread (as a transaction cost) into our calculations eliminates all significantly positive returns for the full sample (b) series. For the support-resistance rules only 10 applications have positive returns at all, whilst 14 returns are significantly negative. The Max-Min results all show negative returns, with 32 of the 36 being significantly negative.

Hence, the main points from all our results are as follows. Sample (a) indicates that significant profits can be made from the application of these trading rules, whilst sample (b) indicates the opposite. This difference in results can be explained in terms of the strength of the trend in exchange rates in the two samples. Trading rule application can be profitable in trending markets. This, however, does not make the task of generating profits any easier; in order to derive profitable returns from technical analysis one not only has to pick an appropriate rule, but must also be able to identify the start points of sustained trends in foreign exchange rates. Further results suggest that, on average, the support-resistance rules perform better than the Max-Min rules used in Brock, Lakonishok and LeBaron (1992) and that, regardless of the strength of the trend in the exchange rate, once transaction costs are accounted for no rule generates significant profits.

Conclusion

In this paper, we have evaluated the profitability of trading strategies based on the application of support and resistance levels provided by traders to Reuters in the intra-daily foreign exchange market. We also simulated some of the technical rules used in former studies. Our main result is that, on average, neither of these sets of rules generate profitable trading strategies. Our result is even more significant when we take transaction costs into account. However, this result requires a number of qualifications. First, as - quite naturally - some profitability could be achieved for the sub-periods where a trend could be observed in the data, one may argue that traders would actually act upon the trading signals reported on Reuters only in periods where such trends are present. Indeed, although the buy and sell signals are the ones reported by traders, we have no figures on the volume traded on these signals. Second, even though these signals are ones reported by traders, it is unlikely that they are the *only* ones that they use. As well as the filter rules applied in this paper there are at least two other main classes of chartist rules, namely rules based on the crossing of moving averages and rules based on so-called momentums. Third, Surajaras and Sweeney (1992) for example have argued that the use of portfolio strategies based on trading several currencies are likely to be more profitable. However, most intra-day foreign exchange traders, whether dealers or brokers, specialise in a single exchange rate. Bearing these words of caution in mind, our results are nevertheless consistent with efficiency in the foreign exchange market.

Bibliography

Brock W., Lakonishok J. and B. LeBaron (1992), "Simple technical trading rules and the stochastic properties of stock returns", *Journal of Finance*, 47, 1731-1764.

Curcio R. and C. Goodhart (1992), "When support/resistance levels are broken, can profits be made? Evidence from the foreign exchange market", *London School of Economics - Financial Markets Group Discussion Paper*, No 142, 1-22.

Curcio R. and C. Goodhart (1993), "Chartism: a controlled experiment", *The Journal of International Securities Markets*, 7, 173-186.

Goodhart C., Hall S., Henry S., and Pesaran B. (1993), "News effects in a high frequency model of the sterling-dollar exchange rate", *Journal of Applied Econometrics*, 8, 1-13.

Guillaume D., Pictet O., Dacorogna M. and Müller U. (1996), "Unveiling nonlinearities through time-scale transformations", *Olsen and Associates Discussion Paper*.

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

Hsieh D. (1989), "Testing for nonlinear dependence in daily foreign exchange rates", *Journal of Business*, 62, 339-368.

International Monetary Fund (1993), *International capital markets. Part I. Exchange rate management and international capital flows*, World Economic and Financial Surveys, Washington, 1-79.

LeBaron B. (1996), "Technical trading rule profitability and foreign exchange intervention", NBER working paper no. 5545, March.

Levich R. and L. Thomas (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.

Surajaras P. and R. Sweeney (1992), *Profit-making speculation in foreign exchange markets*, Westview Press, Inc.

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

Table 1: Summary Statistics for Hourly Return Series

Statistic	Sample a: 10/4/89-29/6/89			Sample b: 31/1/94-30/6/94		
	DEM/USD	JPY/USD	USD/GBP	DEM/USD	JPY/USD	USD/GBP
Mean	0.0000315	0.0000573	-0.0000651	-0.0000355	-0.0000412	0.0000120
Variance	0.0000025	0.0000029	0.0000029	0.0000013	0.0000023	0.0000009
Skewness	0.134	-0.753	-0.129	-0.184	0.031	-0.043
Kurtosis	11.047	13.181	7.886	11.394	7.995	9.608
No. of Obs.	1415	1415	1415	2663	2663	2663
Trend	0.000109*	0.011465*	-0.000150*	-0.000047*	-0.001773*	0.000019*

Note: The row denoted trend presents the linear trend estimated in the raw *midquote* series rather than from the returns. A * indicates that the estimated trend is significantly different from zero at the 5% level.

Table 2: JPY Trading Rule Results, sample a

Series	Side of Mkt	no transaction costs			with transaction costs		
		Number	Mean	t-stat	Number	Mean	t-stat
Returns	-	1415	0.005727	-	1415	5.68E-05	-
S-R	Buy	179	0.032337	1.96*	225	5.05E-05	-0.05
	Sell	111	0.005298	-0.03	108	7.39E-05	0.79
	Composite	290	0.021988	1.47	333	5.81E-05	0.01
H-L	Buy	188	0.038528	2.47*	215	0.000228	1.41
	Sell	121	0.010541	0.30	122	6.17E-05	0.75
	Composite	309	0.027569	2.85*	337	0.000168	1.10
Both	Buy	120	0.052711	2.88*	143	0.000291	1.61
	Sell	75	0.015384	0.48	73	0.000183	1.20
	Composite	195	0.038354	2.03*	216	0.000255	1.63
S-R+ 0.1%	Buy	121	0.066373	3.74*	147	0.000255	1.37
	Sell	86	0.007993	0.12	76	0.000105	0.83
	Composite	207	0.042119	1.74	223	0.000204	1.23
H-L+ 0.1%	Buy	127	0.051268	2.87*	149	0.000314	1.80
	Sell	100	-0.00371	-0.53	93	0.000129	1.04
	Composite	227	0.027049	2.49*	242	0.000243	1.61
Both+0.1%	Buy	88	0.063721	3.08*	100	0.00047	2.40*
	Sell	61	0.022457	0.75	53	0.000312	1.59
	Composite	149	0.046828	2.79*	153	0.000415	2.53*

Notes: In the first column, S-R represents SUPPORT-RESISTANCE, H-L represents HIGH-LOW and Both represents the combination of the two. A * indicates that the given t-statistic exceeds the 5% critical value such that the trading rule returns are significantly different from zero.

Table 3: JPY Max-Min Results, sample a

Series	Side of Mkt	no transaction costs			with transaction costs		
		Number	Mean	t-stat	Number	Mean	t-stat
Returns	-	1415	0.005727	-	1415	5.68E-05	-
50	Buy	145	0.032823	1.81	140	0.000163	0.72
	Sell	63	-0.01832	-1.09	53	-0.00031	-1.08
	Composite	208	0.017333	0.91	193	3.33E-05	-0.18
150	Buy	111	0.038047	1.91	101	0.000263	1.20
	Sell	21	0.004158	-0.04	12	-0.00012	-0.14
	Composite	132	0.032656	1.33	113	0.000222	1.02
200	Buy	102	0.0398	1.94	95	0.000264	1.18
	Sell	13	0.005589	0.00	4	-0.00036	-0.37
	Composite	115	0.035933	1.73	99	0.000239	1.05
50+ 0.1%	Buy	60	0.064038	2.58*	47	0.000378	1.31
	Sell	36	-0.02741	-1.15	24	-0.00064	-1.69
	Composite	96	0.029745	1.73	71	3.55E-05	-0.11
150+ 0.1%	Buy	48	0.063098	2.28*	37	0.000458	1.45
	Sell	16	-0.01487	-0.48	8	-0.00062	-0.96
	Composite	64	0.043606	1.83	45	0.000265	0.83
200+0.1%	Buy	48	0.063098	2.28*	37	0.000458	1.45
	Sell	16	-0.022	-0.56	4	-0.00036	-0.37
	Composite	64	0.041824	1.79	41	0.000378	1.22

Notes: The numbers in the first column represent The horizon of past returns over which local maximum and minimum exchange rate values are calculated. A * indicates that the given t-statistic exceeds the 5% critical value such that the trading rule returns are significantly different from zero.

Table 4: JPY Trading Rule Results, sample b

Series	Side of Mkt	no transaction costs			with transaction costs		
		Number	Mean	t-stat	Number	Mean	t-stat
Returns	-	2663	-4.12E-05	-	2663	-4.12E-05	-
S-R	Buy	103	-0.00014	-0.64	106	-0.00026	-1.37
	Sell	89	0.000362	1.96*	97	9.92E-05	0.35
	Composite	192	9.35E-05	1.22	203	-8.7E-05	-0.39
H-L	Buy	226	-9.58E-05	-0.52	242	-0.00025	-1.92
	Sell	226	0.000259	2.07*	234	0.000138	0.89
	Composite	452	8.16E-05	1.60	476	-5.7E-05	-0.20
Both	Buy	88	-0.0002	-0.98	93	-0.00037	-1.97*
	Sell	84	0.000422	2.26*	92	0.000147	0.63
	Composite	172	0.000102	1.23	185	-0.00011	-0.60
S-R+ 0.1%	Buy	69	-0.00022	-0.94	82	-0.00019	-0.85
	Sell	53	0.000216	0.83	68	0.000282	1.24
	Composite	122	-2.8E-05	0.11	150	2.29E-05	0.48
H-L+ 0.1%	Buy	149	-0.00011	-0.56	168	-0.00038	-2.72*
	Sell	184	0.000341	2.59*	197	0.000146	0.90
	Composite	333	0.000138	2.04*	365	-9.8E-05	-0.64
Both+0.1%	Buy	56	-0.00023	-0.90	67	-0.00036	-1.65
	Sell	52	0.000235	0.91	63	0.000366	1.61
	Composite	108	-4E-06	0.26	130	-1E-05	0.22

Notes: see Table 2.

Table 5: JPY Max-Min Results, sample b

Series	Side of Mkt	no transaction costs			with transaction costs		
		Number	Mean	t-stat	Number	Mean	t-stat
Returns	-	2663	-4.12E-05	-	2663	-4.12E-05	-
50	Buy	139	-0.00021	-1.29	128	-0.00084	-5.56*
	Sell	174	5.15E-06	0.30	167	-0.00028	-2.52*
	Composite	313	-9.1E-05	-0.56	295	-0.00052	-4.94*
150	Buy	66	-0.00022	-0.95	55	-0.00082	-3.59*
	Sell	110	0.000721	4.60*	105	-0.00016	-1.29
	Composite	176	0.000367	3.49*	160	-0.00039	-2.68*
200	Buy	48	-0.00019	-0.65	40	-0.00077	-2.87*
	Sell	109	5.82E-05	0.11	104	-0.00015	-1.19
	Composite	157	-1.6E-05	0.19	144	-0.00032	-2.05*
50 + 0.1%	Buy	47	-0.00058	-2.42*	46	-0.00161	-6.65*
	Sell	78	8.61E-05	0.26	78	-0.00045	-2.72*
	Composite	125	-0.00016	-0.95	124	-0.00088	-5.78*
150 + 0.1%	Buy	18	-0.00043	-1.08	16	-0.00121	-2.95*
	Sell	54	0.000121	0.38	55	-0.00036	-1.86
	Composite	72	-1.6E-05	0.07	71	-0.00055	-2.68*
200 + 0.1%	Buy	15	-0.0003	-0.66	13	-0.00107	-2.34*
	Sell	52	0.00016	0.56	53	-0.00031	-1.57
	Composite	67	5.65E-05	0.47	66	-0.00046	-2.10*

Notes: see Table 3.

Table 6: 'Flat' JPY sub-sample Trading Rule Results, sample b

Series		Number	Mean	t-stat
Returns	-	1929	1.34E-05	-
S-R	Buy	97	-0.00014	-1.03
	Sell	42	-0.00014	-0.57
	Composite	139	-0.00014	-1.22
H-L	Buy	208	-8.18E-05	-0.90
	Sell	142	-4.11E-06	-0.07
	Composite	350	-5E-05	-0.76
Both	Buy	85	-0.00018	-1.23
	Sell	37	-7.45E-05	-0.25
	Composite	122	-0.00015	-1.22
S-R+ 0.1%	Buy	63	-0.00023	-1.30
	Sell	18	-0.0006	-1.71
	Composite	81	-0.00031	-1.97*
H-L+ 0.1%	Buy	138	-9.62E-05	-0.86
	Sell	105	-5.47E-05	-0.29
	Composite	243	-7.8E-05	-0.93
Both+0.1%	Buy	56	-0.00023	-1.22
	Sell	17	-0.00059	-1.63
	Composite	73	-0.00031	-1.88


Notes: see Table 2.

Table 7: 'Trended' JPY sub-sample Trading Rule Results, sample b


Series		Number	Mean	t-stat
Returns		- 734	-0.00018	-
S-R	Buy	6	-9.16E-05	0.13
	Sell	47	0.000812	2.47*
	Composite	53	0.00071	3.72*
H-L	Buy	18	-0.00026	-0.18
	Sell	83	0.000712	2.70*
	Composite	101	0.00054	4.04*
Both	Buy	3	-0.00075	-0.58
	Sell	47	0.000812	2.47*
	Composite	50	0.000718	3.66*
S-R+ 0.1%	Buy	6	-9.16E-05	0.13
	Sell	35	0.000634	1.54
	Composite	41	0.000528	2.63*
H-L+ 0.1%	Buy	11	-0.00033	-0.28
	Sell	78	0.000877	3.44*
	Composite	89	0.000728	4.82*
Both+0.1%	Buy	-	-	-
	Sell	35	0.000634	1.54
	Composite	35	0.000634	1.54

Notes: see Table 2.


Figure 1: Illustration of the trading rule



**Figure 2: DEM/USD Hourly Spot Exchange Rate
10/4/89-29/6/89**




**Figure 3: JPY/USD Hourly Spot Exchange Rate
10/4/89-29/6/89**




**Figure 4: USD/GBP Hourly Spot Exchange Rate
10/4/89-29/6/89**



**Figure 5: DEM/USD Hourly Spot Exchange Rate
31/1/94-30/6/94**



**Figure 6: JPY/USD Hourly Spot Exchange Rate
31/1/94-30/6/94**



**Figure 7: USD/GBP Hourly Spot Exchange Rate
31/1/94-30/6/94**

