

WORKING PAPER SERIES

Intraday Technical Trading in the Foreign Exchange Market

Christopher J. Neely and Paul A. Weller

Working Paper 1999-016B

<http://research.stlouisfed.org/wp/1999/99-016.pdf>

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to Federal Reserve Bank of St. Louis Working Papers (other than an acknowledgment that the writer has had access to unpublished material) should be cleared with the author or authors.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

Intraday Technical Trading in the Foreign Exchange Market

Christopher J. Neely

Paul A. Weller

January 10, 2001

Abstract: This paper examines the out-of-sample performance of intraday technical trading strategies selected using two methodologies, a genetic program and an optimized linear forecasting model. When realistic transaction costs and trading hours are taken into account, we find no evidence of excess returns to the trading rules derived with either methodology. Thus, our results are consistent with market efficiency. We do, however, find that the trading rules discover some remarkably stable patterns in the data.

Keywords: technical trading rules, genetic programming, exchange rates
JEL subject numbers: F31, G15

* Senior Economist, Research Department
Federal Reserve Bank of St. Louis
P.O. Box 442 St. Louis, MO 63166
(314) 444-8568, (314) 444-8731 (fax)
neely@stls.frb.org

† Department of Finance
Henry B. Tippie College of Business
Administration
University of Iowa
Iowa City, IA 52242
(319) 335-1017, (319) 335-3690 (fax)
paul-weller@uiowa.edu

The authors thank Kent Koch for excellent research assistance and Olsen and Associates for supplying the data. The authors also thank Charles Goodhart and Richard Payne for assistance in obtaining additional data. Paul Weller would like to thank the Research Department of the Federal Reserve Bank of St. Louis for its hospitality while he was a Visiting Scholar, when this work was initiated. The views expressed are those of the author(s) and do not necessarily reflect official positions of the Federal Reserve Bank of St. Louis or the Federal Reserve System.

There has been a recent resurgence of academic interest in the claims of technical analysis. This is largely attributable to accumulating evidence that technical trading can be profitable over long time horizons (Brock, Lakonishok and LeBaron, 1992; Levich and Thomas, 1993; Neely, Weller and Dittmar, 1997)¹.

However, academic investigation of technical trading in the foreign exchange market has not been consistent with the practice of technical analysis. The majority of foreign exchange traders who use technical analysis are intraday traders who transact at high frequency and aim to finish the trading day with a net open position of zero. But, due to data limitations, most academic studies have evaluated the profitability of trading strategies that allow trades to be executed at most once a day. For example, in our earlier study using daily data (Neely, Weller and Dittmar, 1997), we provide strong evidence for the existence of profitable trading rules for a variety of currencies over a fifteen-year time horizon. The mean trading frequency for the rules we identify ranges from once every two weeks to once every three months. Evidently, these are not the trading strategies being used by the foreign exchange dealers in the London market surveyed in Taylor and Allen (1992). They documented the fact that technical analysis was widely used for trading at the shortest time horizons, namely days and weeks, and was used in some form by over 90 per cent of their respondents.²

This paper links trading practice with research more closely by investigating the performance of trading rules using high frequency data that allow the rules to change position within the trading day. There are two potential goals for such a paper: 1) to investigate the efficiency of foreign exchange markets at high frequencies; or 2) to find profitable technical

¹ Ready (1998) questions the evidence relating to the equity market.

rules, per se. Our primary focus is to use the performance of the trading rules we identify to provide a measure of market efficiency, an approach first advocated in Brock, Lakonishok and LeBaron (1992). We use an in-sample period to search for ex ante optimal trading rules and then assess the performance of those rules out-of-sample. We use two distinct methodologies: the first is a genetic program that can search over a very wide class of (possibly non-linear) trading rules; the second consists of linear forecasting models, which provide natural benchmarks against which to compare the genetic programming results. Our analysis does not specify the type of trader who might use such rules, but does assume that the trader faces reasonably low transactions costs. We find strong evidence of predictability in the data as measured by out-of-sample profitability when transaction costs are set to zero. However, the excess returns earned by the trading rules are very sensitive to the level of transaction costs and to the liquidity of the markets. When transaction costs are taken into account and trading is restricted to periods of high market activity, there is no evidence of profitable trading opportunities. Thus, our results are consistent with the efficient markets hypothesis.

2. Previous Work on Trading Rules in the Foreign Exchange Market

Most of the work analyzing technical trading rules in the foreign exchange market has used daily or weekly data, and has examined the profits to be earned by employing a particular rule or class of rules suggested by practicing technical analysts. The trading rules that have been most intensively investigated use filters and moving averages. A simple filter rule takes the form: “Take a long position in foreign currency when the exchange rate (dollar value of foreign

² A more recent survey of foreign exchange traders in the United States (Cheung and Chinn, 1999), although it did not explicitly address the issue of trading horizon, found that 30 per cent of the sample used technical analysis as the dominant guide to trading, compared to 25 per cent using fundamental analysis.

currency) rises by $x\%$ above its previous minimum over the last y days; take a short position when the exchange rate falls $x\%$ below its previous maximum over the last y days. Otherwise maintain the current position.” A double moving average rule compares a short- to a long-run moving-average, changing from a sell to a buy signal when the short-run moving average exceeds the long-run moving average by a given amount. Even these simple rules can take a great variety of different forms. The moving-average rules will vary depending on the time windows chosen for each moving average and the amount by which the short moving average must exceed or fall below the long moving average. The filter rules will depend on the size of the filter and the time window over which the previous high or low is calculated. Both classes of rules seek to identify changes in a trend.

A number of studies have examined the performance of trading rules using daily foreign exchange data (Dooley and Shafer, 1983; Sweeney, 1986; Levich and Thomas, 1993; Osler and Chang, 1999). The general conclusion is that the trading rules are able to earn significant excess returns net of transaction costs, and that this cannot be easily explained as compensation for bearing risk. For example, Neely, Weller and Dittmar (1997) found out-of-sample annual excess returns in the one to seven percent range in currency markets against the dollar during the period 1981-95. The highest trading frequency was observed in the rules found for the DEM/JPY, and was between two and three trades per month. This does not resemble the technical trading strategies used by most foreign exchange traders. We therefore seek to discover if the profit opportunities that exist over medium- to long-term horizons are also present at the short horizons typically employed by traders.

Although there has been much work investigating the statistical properties of high frequency exchange rate data—see Goodhart and O’Hara (1997)—there has been relatively little

work on high frequency trading rules. Goodhart and Curcio (1992) consider the usefulness of support and resistance levels published by Reuters. Osler (2000) looks at support and resistance levels published by six trading firms over the period 1996-1998, and finds evidence of significant power to predict intraday trend reversals. But she does not investigate whether it is possible to trade profitably on the basis of the signals net of transaction costs. Acar and Lequeux (1995) examine the profitability of a class of linear forecasting rules fitted to a sample of half-hourly data while Curcio et al. (1997) examine the performance of filter rules that have been identified by practitioners. None of these papers finds evidence of profit opportunities. However, by focusing on narrow classes of rules they are not able to rule out the possibility that a search over a broader class would reveal profitable strategies. Pictet et al. (1996) employ a genetic algorithm to optimize a class of exponential moving average rules. They run into serious problems of overfitting, and their rules perform poorly out of sample. Gençay et al. (1998) report 3.6 to 9.6 per cent annual excess returns, net of transactions costs, to proprietary real-time Olsen and Associates trading models using seven years of exchange rate data at 5-minute frequency. It is difficult to compare other results to theirs, given that their models are not publicly available.

3. The Genetic Program

Genetic algorithms are computer search procedures based on the principles of natural selection. These procedures were developed by Holland (1975) and extended by Koza (1992). They have been applied to a wide variety of problems in many fields and are most useful in situations where the space of possible solutions to a problem consists of decision trees or programs and thus cannot be handled by hill-climbing search routines that require differentiability. Our use of the genetic program follows an approach first applied to the foreign

exchange market in our earlier paper. Our description of the procedures used here will follow that of the previous paper.

In genetic programming, the individual candidate solutions are hierarchical character strings of variable length. These structures can be represented as decision trees, whose non-terminal nodes are mathematical functions, operators or constants. We make use of the following function set:

- arithmetic operations: “plus”, “minus”, “times”, “divide”, “norm”, “average”, “max”, “min”, “lag”;
- Boolean operations: “and”, “or”, “not”, “greater than”, “less than”;
- conditional operations: “if-then”, “if-then-else”;
- random numerical constants picked uniformly from (0,6);
- Boolean constants: “true”, “false”.

“Norm” returns the absolute value of the difference between two numbers. “average”, “max”, “min”, and “lag” respectively return the moving average, local maximum, local minimum and lagged value of a data series over a time window specified by the argument of the function, rounded to the nearest whole number.

An important advantage of genetic programming in constructing trading rules is that they can use (or ignore) additional information to construct technical rules (Neely and Weller, 1999 and 2001). In this exercise, we use three information variables as input to the genetic program. The first is the normalized value of the exchange rate, the exchange rate divided by its moving average over the previous two weeks.³ The second summarizes information on the interest

³ The normalization provides the rules with similar magnitudes of data both in and out of sample. For example, a rule comparing the exchange rate to a constant in the in-sample period could perform poorly because the constant was of inappropriate magnitude out of sample. This could come about as a result of non-stationarities in the data.

differential, and is defined in the section describing the data. The third variable is the hour of the day. We include this last variable because of the large and consistent intraday fluctuation in trading volume in foreign exchange markets. This is known to be associated with volatility, but may also have an effect on the first moment of the exchange rate series.

The genetic program searches for good solutions to problems of interest using the principles of natural selection. The program first randomly creates a population of arbitrary rules and allows the members of that population to reproduce and recombine their components over subsequent generations. Profitable rules are more likely to have their components reproduced in subsequent populations. In this way, the genetic program searches through the space of rules, concentrating on those parts of the space that have been shown to produce profitable rules. The basic features of the genetic program are (a) a means of encoding trading rules so that they can be built up from separate subcomponents (b) a measure of excess return or “fitness” (c) an operation which splits and recombines existing rules in order to create new rules.

We denote the exchange rate at period t (dollars per unit of foreign currency) by S_t , the short-term interest differential by D_t and time of day at period t by the variable, T_t . A trading rule is a mapping from past exchange rates and interest differentials indexed by time of day to a binary variable, z_t , which takes the value +1 for a long position in foreign exchange at time t , and -1 for a short position. Trading rules may be represented as trees, whose nodes consist of various arithmetic functions, logical operators and constants. The functions are distinguished by the data series on which they operate. Thus $\max_s(3)$ at time t is equivalent to $\max(S_t, S_{t-1}, S_{t-2})$, $\text{lag}_t(3)$ at time t is equal to T_{t-3} , and $\text{average}_s(3)$ is equal to the mean of S_t, S_{t-1} and S_{t-2} . In choosing the function set, we sought to maximize the potential power of the genetic program by including

simple functions known to be useful from previous work but avoided including additional arbitrary functions that would reduce the efficiency of the search.

Figure 1 presents an example of a trading rule that makes use of both exchange rate and time of day data. The rule signals a long position in foreign currency if the current exchange rate is greater than the 48-period moving average and the time of day (GMT) is between 0800 and 1600, and a short position otherwise. This example illustrates a simple, time-dependent rule. The function “rate” returns the average of bid and ask quotes for the exchange rate at half-hourly intervals.

The fitness criterion for the genetic program is the continuously compounded excess return to the trading rule over the given time period. We train rules under two assumptions about when they can trade. The first scenario permits trading 24 hours a day, 7 days a week. The second scenario—called restricted trading—only permits trading during 12-hour periods of heavy trading in the particular currency on business days. After the 12 hours of trading, the rule earns the overnight interest rate in the currency in which it is long—losing the overnight interest rate in the other currency. The continuously compounded (log) excess return over a half-hour is given by $z_t r_t$ where z_t is the indicator variable described above, and r_t is defined as:

$$r_t = \ln S_{t+1} - \ln S_t . \tag{1}$$

Each trade involves switching from a long to a short position or vice versa, and so incurs a round trip transaction cost. In other words, trading from a position long x units of foreign currency to one short the same amount requires a sale of $2x$ units, incurring a proportional

transaction cost of $2c$. Therefore the cumulative excess return r for a 24-hour trading rule giving signal z_t at time t over the period from time zero to time T is:

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

where n is the number of trades. This measures the fitness of the rule. Returns to rules subject to restricted trading would be computed using the interest differential for overnight positions as well as the exchange rate return.

Figure 2 illustrates the crossover and reproduction operation. A pair of rules is selected at random from the population, with a probability weighted in favor of rules with higher fitness. Then subtrees of the two parent rules are selected randomly. One of the selected subtrees is discarded, and replaced by the other subtree, to produce the offspring rule.⁴

To implement the genetic programming procedures we define 3 separate subsamples, referred to as the training, selection and out-of-sample test periods. The first two periods are equivalent to an in-sample estimation period. The third, the out-of-sample test period, is used to measure the performance of the rules trained and selected in the first two periods. The distinct time periods for all currencies were chosen as follows: training period, 02/01/96 to 03/31/96; selection period, 04/01/96 to 05/31/96; test period, 06/01/96 to 12/31/96. The first month of data was used to calculate starting values for moving averages and other functions taking lagged values as arguments.

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

1. Create an initial generation of 1000 randomly generated rules.
2. Measure the excess return of each rule over the training period and rank according to excess return.
3. Select the top-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 initial 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 1000 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 of the second generation.
6. Return to step 4 and repeat until we have produced 40 generations or until no new best rule appears for 10 generations.

4. The Linear Forecasting Model

We estimate an autoregressive model for each exchange rate over the training and selection periods on 24-hour data, including weekends, using only own lagged values of the log exchange rate. We restrict the maximum number of lags to 10. We then combine each estimated forecasting model with a filter to produce a trading rule. The purpose of the filter is to reduce trading frequency and accompanying transactions costs for those periods in which only a small

⁴ The operation is subject to the restriction that the resulting rule must be well-defined, and that it may not exceed a

change in the exchange rate is predicted. Denoting the one-period-ahead forecast of the log exchange rate at time t by $E_t(\ln(S_{t+1}))$ and the filter by f , trading signals are determined in the following way:

$$\begin{aligned}
 \text{If } z_{t-1} = +1, \quad z_t = -1, & \text{ if } E_t(\ln(S_{t+1})) < \ln S_t - f, \\
 & = +1, \text{ if } E_t(\ln(S_{t+1})) \geq \ln S_t - f. \\
 \\
 \text{If } z_{t-1} = -1, \quad z_t = +1, & \text{ if } E_t(\ln(S_{t+1})) > \ln S_t + f, \\
 & = -1, \text{ if } E_t(\ln(S_{t+1})) \leq \ln S_t + f.
 \end{aligned} \tag{3}$$

For example, the first two conditional equations above say that, if the rule has a long position at $t-1$, it will only switch to a short position at t if the exchange rate is forecast to fall by more than the size of the filter from t to $t + 1$. If the forecast change in the exchange rate is greater than or equal to the negative size of the filter, the rule will maintain a long position. Trading rules with filters ranging from zero to 0.0005 in steps of 0.0001 and estimated lag coefficients from one to ten are run on the data from the training and selection periods, and excess return is calculated assuming the following three values of one-way transaction cost: 0, 0.0001, and 0.0002. The trading rule with the highest excess return for each of the three levels of transactions cost is then run on the out-of-sample test period.

We also estimate an expanded model in which in addition to the lagged values of the log exchange rate we include the lagged normalized exchange rate, the hour of the day and the interest futures differential. Trading rules are then formed in the same way, tested in-sample over the same range of lags and filters and run on the test period.

specified size (10 levels and 100 nodes).

5. The Data

We use half-hourly bid and ask quotes for spot foreign exchange rates during 1996 from the HFDF96 data set provided by Olsen and Associates. Half-hourly quotes provide a useful tradeoff between the desires to accurately approximate the information set of an intraday trader and to limit the size of data sets and computational costs. They also substantially reduce the risk of introducing microstructural artifacts. We examine four currencies against the dollar – the German mark (DEM), the Japanese yen (JPY), the British pound (GBP) and the Swiss franc (CHF). We use three variables as input to the genetic program. The first is the normalized half-hourly exchange rate series, constructed by calculating a simple average of bid and ask quotes and dividing by a two-week moving average. The second is the difference (U.S. minus foreign contract) in the transaction prices for the short-term interest rate futures contract whose expiry is closest to the time stamp of the exchange rate data. Because Japanese futures data were unavailable, only the U.S. futures price was used for the JPY exchange rate. The U.S. contract is traded on the Chicago Mercantile Exchange. Data for the foreign contracts comes from the London International Financial Futures Exchange (LIFFE). The third variable is the time of day (GMT).

We present summary statistics for the distributions of half-hourly log exchange rate changes in Table 1. Standard deviations are quite similar across currencies, and all exchange rates display very high kurtosis. In the top panel of Figure 3 we plot autocorrelations for the log returns, using all hours, and find highly significant negative first order autocorrelation for all currencies.⁵ This significant first-order autocorrelation is also present in both bid and ask prices

⁵ It is interesting to compare this result with the findings of Bollerslev and Domowitz (1993). Using observations on percentage changes at five-minute intervals in the USD/DEM market they report *positive* third-order autocorrelation when the series is constructed either from the average of bid and ask quotes or from an algorithm approximating

and it is robust to excluding outliers in the bid-ask spread. We checked to see whether the summary statistics or autocorrelation patterns were sensitive to the omission of weekends or off-peak trading hours. The bottom panel of Figure 3 shows that measuring autocorrelation only during business hours reduces mean first-order autocorrelation to -0.12, from -0.17 when measured during all hours. There was also a decline in kurtosis as more periods of low market activity were omitted. However the kurtosis still remained highly significant in all cases.

Baillie and Bollerslev (1991) note the existence of significant negative first order serial correlation in hourly exchange rate series, and suggest that it is a spurious consequence of two features of the data collection process. They used a data set in which each observation consisted of the average of the five most recent bid quotes, a procedure known to induce serial correlation. However this is not a feature of the data set that we use, and so we consider the second reason they propose—non-synchronous trading. If there are periods during which no trade occurs, and a zero return is recorded, this may not be an accurate reflection of the movement of the true underlying return process. When a trade occurs after a period of inactivity, the observed return is a sum of the accumulated returns over the periods of no trade. If the series has a non-zero mean, this will induce mean reversion in the observed series. The first point to make is that quotes may adjust even when no trade takes place, so it is unclear to what extent the argument applies to our data set. However, we do find a rather high proportion of zero returns in the full sample, ranging across the four currencies from 22 to 26 per cent of the total number of observations. This is almost entirely due to the presence of weekends, and the figures fall to a range of 3.9 to 5.9 per cent when weekends are excluded and further when only business hours are considered.

transaction prices. Thus the pattern of momentum and reversal documented at longer time horizons in equity markets appears here at much shorter horizons.

Lo and MacKinlay (1990) derive a formula for the level of serial correlation induced by non-synchronous trading if the true return series follows a random walk with drift. It depends on the mean and variance of the return series, and the probability of no trade occurring. We use the sample proportion of zero return observations as a proxy for the probability of no trade occurring. If true returns are generated by the model

$$r_t = \mu + \varepsilon_t \quad (4)$$

where ε_t is a noise term with zero mean and variance σ^2 independent at all leads and lags, and π is the probability of no trade, then $\rho(i)$, the induced correlation in observed returns r_t^o at lag i is given by

$$\rho(i) = \frac{-\mu^2 \pi^i}{\sigma^2 + \frac{2\pi}{1-\pi} \mu^2}. \quad (5)$$

We obtain a figure for $\rho(1)$ of -0.00000706 for the DEM even when the probability of no trade is set to 0.227. This is to be compared with the observed value of -0.14 . We conclude that the magnitude of serial correlation observed in the data cannot be explained by non-synchronous trading, and treat it as a true feature of the data.

6. Results

We consider first the unrestricted, benchmark case in which trading is allowed to take place twenty-four hours a day, seven days a week. For each currency we generated twenty-five rules from the genetic program under each of three assumptions about transactions costs in training and selection periods. We used one-way transaction costs of zero, one and two basis

points ($c = 0, 0.0001, 0.0002$).⁶ From those twenty-five rules we selected those which had a positive excess return during the selection period and also traded at least once.

We aggregate the signals from the sets of individual rules by constructing an equally weighted portfolio rule. The equally weighted portfolio rule assumes that the trader permits each rule an equal share in the position taken by the portfolio. Table 2 presents results for this rule. To investigate pure predictability—as opposed to profitability—in column three we report annual returns assuming zero transaction costs in the out-of-sample period. To indicate the potential profitability (or lack thereof) of these rules, column six of Table 2 reports the level of transaction cost measured in basis points that would reduce the excess return to zero (break-even transaction cost). The rules trained with zero transactions costs in-sample produce returns that are very high, in three of the four cases over 100 per cent per annum. This provides strong evidence of a predictable component in the exchange rate series. But the rules trade very frequently, approximately once an hour on average, or every other period. Because of this, the break-even transaction cost is low. The highest figure among the four exchange rates, that for the British pound, is 1.01 basis points for a one-way trade. This is largely attributable to the somewhat lower trading frequency of these rules.

As the transaction cost in training and selection periods is increased from zero to one and then two basis points, both annual excess returns before transaction cost and trading frequency fall sharply. But break-even transaction cost rises uniformly to levels close to the level that a large institutional trader would face. It also becomes more difficult to find good rules according to our in-sample selection criteria, most notably in the case of the GBP, where only five of the

⁶ We chose not to compute rules for higher levels of transaction cost because of the increasing difficulty of finding rules that were profitable in-sample using higher levels of costs. In addition, estimates of foreign exchange

twenty-five rules satisfied the criteria for $c = 0.0002$. One of the most striking features of Table 2 is the steady rise in break-even transaction cost as the in-sample value of c is increased.⁷ Since the break-even transaction cost can be interpreted as the average excess return per zero-cost trade, this demonstrates the ability of the search procedure to identify rules that can successfully predict not just the direction but also the *magnitude* of a price change. It also shows that there are remarkably stable patterns in the high frequency data. Although a purely speculative trader cannot exploit these patterns, they nevertheless represent important information for foreign exchange dealers. A dealer who takes account of the predictability in the exchange rate in setting quotes will trade more profitably than one who does not.

We can investigate more systematically the role played by serial correlation in the data by comparing the performance of the linear (autoregressive) forecasting model with that of the genetic program. Table 3 reports the estimated coefficients of the models with the highest excess return (net of one-way transaction costs of 0.0001) over training and selection periods. As the statistics on serial correlation would lead one to expect, the first two lags in the data are much the most important in all cases. Only the model for the GBP has more than three lags and the coefficients on lags four and higher are small. The implied models for the log exchange rates are (barely) stationary. The optimally selected filters for all currencies are 0.0001, matching the chosen level of transaction cost. When we consider the out-of-sample performance of the autoregressive forecasting model (see Table 4) we see a similar pattern of improvement as the in-sample transaction cost increases. However, the results are clearly superior to those derived from

transaction costs suggest that 2-2.5 basis points for a one-way trade is realistic for recent large transactions (Neely, Weller and Dittmar, 1997).

⁷ The number of trades and breakeven transaction cost for the equally weighted rule are not simple averages over all rules. We correct for the fact that if two rules simultaneously trade in opposite directions, this has no effect on the net open position and so does not generate a trade.

the genetic program at the highest level of transaction cost. In all cases the break-even transaction cost is higher, dramatically so in the case of the DEM, where it is 24.4 basis points. If we take 2.5 basis points as an estimate of the one-way transaction cost faced by a large institutional trader, then the trading rules earn excess returns net of transaction cost which in all cases exceed twenty per cent per annum. However, it is not clear that this is a reasonable thing to do given that we have assumed that trading takes place twenty four hours a day and during weekends. There are periods during the week when the major markets are closed and trading activity is much reduced. The fall in liquidity is very likely to be associated with an increase in transaction cost.

For this reason we generate a new set of rules under the assumption that trading is restricted to occur during a twelve-hour period on weekdays only. Such rules were able to observe both business and non-business data but were only permitted to change positions during business hours. During non-business hours the rules earned or lost the appropriate interest differential. We selected the business hours to coincide with the time of the most active trading in the particular currency (see Melvin, 1997 for figures on the DEM). They were chosen as follows: DEM 0600-1800 GMT, JPY 0400-1600 GMT, CHF 0500-1700 GMT, and GBP 0500-1700 GMT.

The results for the genetic program with restricted trading are presented in Table 5. The annual excess returns with zero transactions costs are reduced in close proportion to the reduction in trading time for all currencies except the DEM. There is still strong evidence of predictability for these currencies. Break-even transaction costs are generally reduced to a level below that which an institutional trader would face. The only exception to this is the GBP, where for $c = 0.0002$ we find a break-even transaction cost of 4.2 basis points. One should be cautious about

reading too much into this finding. There were a relatively small number of good rules identified in sample, they traded infrequently and tended to be skewed towards short positions.⁸

In Table 6 we show the results of imposing restricted trading on the autoregressive-forecasting model. Again, the models are estimated on 24-hour, in-sample data but are only permitted to change positions during business hours. During non-business hours, the models earn or lose the appropriate interest differential. The model with the highest excess return net of transaction costs is then tested out of sample. The picture changes dramatically when compared to the figures in Table 4. In all cases in which the rules trade, the break-even transaction costs fall to a level below that which even a large trader would face. This demonstrates conclusively that the apparent profitability of the trading rules obtained with $c = 0.0002$ is solely attributable to trading during periods of reduced market activity when transaction costs are likely to be substantially higher than the benchmark figure of 2.5 basis points that we have chosen.

The results from the rules derived from the extended linear forecasting model with additional variables were not significantly different from the autoregressive results reported in Table 4, and so we omit them. This indicates that there is no additional forecasting power contained in the variables added in the extended model—exchange rate normalized by a two-week moving average, time of day and interest differential.

Tables 7 and 8 summarize the extent to which the genetic program rules find common patterns during the out-of-sample test period. For each observation we calculate the proportion of rules which signal a long position and then count how often the proportion of long rules lies in each quintile. For example, the first entry in the third column of Table 7 indicates that for 50.6

⁸ We also computed results for the case where rules trained on 24-hour data were used for restricted trading out of sample. Thus, whatever position was signaled by the rule at the beginning of the no-trade period was held until trade

percent of the observations, 0 to 20 percent of the all-day DEM rules with $c = 0$ were long in the DEM. That is, more than 80 percent of the rules were simultaneously short the DEM over half the time. High numbers in the first and last quintiles indicate consensus among the rules. If there were no predictable patterns in the data, the trading rules would switch randomly between long and short positions and we would tend to observe a high percentage of observations in the middle quintile. We observe the highest degree of consensus in the all-day trading scenario with zero transaction cost. There is a general tendency for consensus to decline as transaction costs are increased.

The fact that the trading rules identified by the genetic program generally perform less well than those generated by the autoregressive-forecasting model deserves some comment. This is likely to be attributable to two factors. First, the variables in addition to the exchange rate series that were provided as input to the genetic program proved not to be informative. This is suggested by the fact that the inclusion of these variables in the forecasting model did not make any difference.⁹ We have found in our previous work that the inclusion of uninformative data can degrade the efficiency of the genetic program. Second, if the relevant information enters the model in a linear fashion, then confining the search to the set of linear models will be a more efficient procedure.

was allowed to start again. The returns of the rules were in almost all cases inferior to those reported in Table 5 and we do not include them.

⁹ We confirmed this fact for the case of the genetic program rules by conducting various experiments in which the separate data series were randomized separately and changes in out-of-sample performance for the rules were recorded. No significant impact was observed for any series but the exchange rate.

7. Discussion and Conclusion

Our findings demonstrate that there are very stable predictable components to the intraday dollar exchange rate series for all the currencies we consider, German mark, Japanese yen, Swiss franc and British pound. But neither the trading rules identified by the genetic program nor those based on the linear forecasting model produce positive excess returns once reasonable transaction costs are taken into account and trade is restricted to take place during times of higher market activity. The rules based on the autoregressive forecasting model perform at least as well as those found by the genetic program and the extended linear model, indicating that our results are largely attributable to the low order negative serial correlation in the data.

A number of authors have found negative first order autocorrelation in exchange rate returns at various high frequency horizons and some have offered explanations. Baillie and Bollerslev (1991) claim that non-synchronous trading and zero trading are responsible for the negative autocorrelation in hourly data . We have shown that this explanation is implausible for our data set. Using the model of Lo and Mackinlay (1990) to relate the induced negative autocorrelation to the probability of no trade, we find that the figure implied by the data is several orders of magnitude smaller than that actually observed. Goodhart and Figliuoli (1991) find negative autocorrelation in minute-by-minute data. Zhou (1996) explains negative autocorrelation in tick-by-tick data with “errors in data” and “screen fighting.” Again, we find neither explanation convincing in our (lower-frequency) half-hourly data. Screen fighting effects are unlikely to persist for so long. Another potential explanation is provided by Danielsson and Payne (2001), who document differences between indicative quotes of the type used here and firm interdealer quotes at very high frequencies. These differences disappear, however, as one samples returns at five- (or more-) minute intervals. Therefore, it appears to be an unlikely

explanation for autocorrelation in 30-minute returns. We obtained a five-day sample of USD/DEM transactions data from Reuters (D2000-2) but were unable to estimate the autocorrelation in half-hourly returns with sufficient precision to reject any hypothesis of interest.¹⁰ It is not clear to us, however, that one should prefer firm interdealer quotes for a study of technical analysis. Dealers make their money on spreads, and indicative quotes may provide a more accurate picture of the terms on which a non-dealer can trade. A further problem is that long spans of firm interdealer quotes are not available for study.

What we have shown is that it is unnecessary to assume that the negative autocorrelation is an artifact. If it is a true feature of the data, it is too small to be exploited by non-dealers to make speculative excess returns. This conclusion highlights the importance of going beyond simple evidence of predictability in order to assess market efficiency.

A striking feature of our results is that the break-even transaction costs generally converge to a level close to that faced by a large institutional trader, namely two to three basis points per one-way trade. These conclusions are based on an analysis of round-the-clock trading. If we restrict trading to occur during a twelve-hour window of high volume, break-even transaction costs are considerably reduced.

Our findings are consistent with those of Lyons (1998). He examined the trading behavior of a foreign exchange dealer over the course of a week, using data that enabled him to decompose profits into speculative and non-speculative components. He found that he could attribute less than ten per cent of profits to speculation and that the vast majority came from trading off the spread.

¹⁰ The authors thank Charles Goodhart and Richard Payne for their assistance in obtaining the D2000-2 data.

It is interesting that the foreign exchange market seems to display quite different characteristics depending on the trading horizon. At weekly and monthly horizons there is strong evidence to indicate significant and persistent trends, but, as we show here, this is not the case at intraday horizons.¹¹ This may be a consequence of the uneven division of capital allocated to financing trade at different horizons. Although no precise figures are available, there is little doubt that a much greater volume of transactions is accounted for by traders who close their positions at the end of each day than by those who take open positions with horizons of weeks or months.

¹¹ Another possibility is that the 1996 exchange rate data were anomalous in being characterized by a lot of “range-trading” and low volatility.

References

- Acar, E., Lequeux, P., 1995. Trading Rules Profits and the Underlying Time Series Properties. Unpublished manuscript, Banque Nationale de Paris, London Branch.
- Baillie, R.T., Bollerslev, T., 1991. Intra-day and inter-market volatility in foreign exchange rates. *Review of Economic Studies* 58 (3), 565-585.
- Bollerslev, T., Domowitz, I., 1993. Trading patterns and prices in the interbank foreign exchange market. *Journal of Finance* 48 (4), 1421-1443.
- Brock, W., Lakonishok, J., LeBaron, B., 1992. Simple technical trading rules and the stochastic properties of stock returns. *Journal of Finance* 47 (5), 1731-1764.
- Cheung, Y.W., Chinn, M.D., 1999. Macroeconomic implications of the beliefs and behavior of foreign exchange traders. Working Paper, University of California, Santa Cruz.
- Curcio, R., Goodhart, C., Guillaume, D., Payne, R., 1997. Do technical trading rules generate profits? Conclusions from the intra-day foreign exchange market. *International Journal of Finance and Economics* 2 (4), 267-280.
- Danielsson, J., Payne, R., 2001. Real trading patterns and prices in spot foreign exchange markets. Forthcoming, *Journal of International Money and Finance*.
- Dooley, M.P., Shafer, J.R., 1983. Analysis of Short-run Exchange Rate Behavior: March 1973 to November 1981. In Bigman, T., Taya, T. (Eds.), *Exchange Rate and Trade Instability: Causes, Consequences, and Remedies*. Cambridge, MA, pp. 43-69.
- Gençay, R., Ballochi, G., Dacorogna, M., Olsen, R., Pictet, O., 1998. Real-time trading models and the statistical properties of foreign exchange rates. Unpublished manuscript, Olsen & Associates.

- Goodhart, C., O'Hara, M., 1997. High frequency data in financial markets: issues and applications. *Journal of Empirical Finance* 4 (2-3), 73-114.
- Goodhart, C.A.E., Figliuoli, L., 1991. Every minute counts in financial markets. *Journal of International Money and Finance* 10 (1), 23-52.
- Goodhart, C., Curcio, R., 1992. When support/resistance levels are broken, can profits be made? Evidence from the foreign exchange market. Discussion Paper 142, London School of Economics, Financial Markets Group.
- Holland, J., 1975. *Adaptation in Natural and Artificial Systems*. University of Michigan Press, Ann Arbor, MI.
- Koza, J.R., 1992. *Genetic Programming: On the Programming of Computers by Means of Natural Selection*. MIT Press, Cambridge, MA.
- Levich, R.M., Thomas, L.R., 1993. The significance of technical trading rule profits in the foreign exchange market: a bootstrap approach. *Journal of International Money and Finance* 12 (5), 451-474.
- Lo, A.W., MacKinlay, A.C., 1990. An Econometric Analysis of Nonsynchronous Trading. *Journal of Econometrics* 45 (1/2), 181-211.
- Lyons, R.K., 1998. Profits and position control: a week of FX dealing. *Journal of International Money and Finance* 17 (1), 97-115.
- Melvin, M., 1997. *International Money and Finance*. Addison Wesley Publishing, Reading, MA.
- Neely, C.J., Weller, P.A., 1999. Technical trading rules in the European Monetary System. *Journal of International Money and Finance* 18 (3), 429-458.

- Neely, C.J., Weller, P.A., 2001. Technical analysis and central bank intervention. Federal Reserve Bank of St. Louis Working Paper 97-002B. Forthcoming in the Journal of International Money and Finance.
- Neely, C.J., Weller, P.A., Dittmar, R., 1997. Is technical analysis in the foreign exchange market profitable? A genetic programming approach. Journal of Financial and Quantitative Analysis 32 (4), 405-426.
- Osler, C.L., Chang, P.H.K., 1999. Methodical madness: technical analysis and the irrationality of exchange-rate forecasts. Economic Journal 109 (458), 636-661.
- Osler, C.L., 2000. Support for resistance: technical analysis and intraday exchange rates. Federal Reserve Bank of New York Policy Review 6 (2), 53-68.
- Pictet, O.V., Dacorogna M.M., Dave, R.D., Chopard, B., Schirru, R., Tomassini, M., 1996. Genetic algorithms with collective sharing for robust optimization in financial applications. Unpublished manuscript, Olsen & Associates.
- Ready, M.J., 1998. Profits from technical trading rules. Unpublished manuscript, University of Wisconsin, Madison.
- Sweeney, R.J., 1986. Beating the foreign exchange market. Journal of Finance 41 (1), 163-182.
- Taylor, M.P., Allen, H., 1992. The use of technical analysis in the foreign exchange market. Journal of International Money and Finance 11 (3), 304-314.
- Zhou, B., 1996. High-frequency data and volatility in foreign-exchange rates. Journal of Business and Economics Statistics 14 (1), 45-52.

Table 1
Summary statistics

	<i>Mean</i>	<i>Std. Dev.</i>	<i>Skew</i>	<i>Kurt</i>	$\rho(1)$	$\rho(2)$	$\rho(3)$	<i>Min</i>	<i>Max</i>
DEM	0.00022	0.07189	-0.07	25.64	-0.14	-0.03	-0.01	-0.93	0.97
JPY	0.00050	0.07939	-0.05	14.16	-0.17	-0.02	0.00	-0.90	0.92
CHF	0.00063	0.09339	-0.23	31.73	-0.17	-0.01	-0.01	-1.59	1.62
GBP	-0.00071	0.07041	0.27	34.14	-0.19	-0.03	-0.02	-1.20	1.22

Note: the table presents statistics for log exchange rate changes constructed from the full data set, consisting of 16080 half-hourly observations (average of bid and ask) taken 24 hours a day, seven days a week for the year 1996. Mean and standard deviation are multiplied by 100. The skewness and kurtosis statistics are standard normally distributed. $\rho(i)$ records the autocorrelation coefficient at lag i . Min and max record the smallest and largest half-hourly percentage changes over the sample period.

Table 2
Out-of-sample trading rule performance for the equally weighted portfolio rule:
All-day trading

	<i>c</i>	<i>Annual return</i>	<i>Number of rules</i>	<i>Number of trades</i>	<i>Break-even transaction cost</i>	<i>% long</i>	<i>Long return</i>
DEM	0.0000	66.92	25	4908.76	0.40	45.69	2.30
DEM	0.0001	46.09	21	887.57	1.51	60.19	
DEM	0.0002	6.30	19	88.58	2.08	54.46	
JPY	0.0000	130.56	25	4164.44	0.91	48.40	11.72
JPY	0.0001	43.28	23	451.57	2.80	45.30	
JPY	0.0002	16.30	13	144.69	3.28	60.33	
CHF	0.0000	127.48	25	4846.88	0.77	50.02	11.51
CHF	0.0001	92.40	25	1773.96	1.52	50.46	
CHF	0.0002	30.99	15	388.60	2.33	45.98	
GBP	0.0000	132.34	25	3830.92	1.01	49.62	-15.80
GBP	0.0001	111.18	25	1920.96	1.69	48.49	
GBP	0.0002	31.59	5	412.00	2.24	63.60	

Note: the equally weighted portfolio rule attaches a weight ($1/\#$ of rules) to each rule satisfying the selection criteria. Column 2 records the value of c , the one-way transaction cost used in training and selection periods. Column 3 gives the annualized per cent excess return over the seven-month out-of-sample test period calculated assuming zero transaction cost. Column 4 reports the number of rules out of the twenty-five obtained for each case that produced a positive excess return before transactions costs and also traded. These were the rules used for the out-of-sample test. Number of trades reports the number of trades corrected for double counting. Break-even transaction cost is the one-way transaction cost (in basis points) which reduces the annual excess return during the test period to zero. The break-even cost is computed as $100 \cdot \left(\frac{7}{12}\right) \text{annual return} / (2 \cdot \text{number of trades})$. % long is the average percentage of the test period the rules held a position long the foreign currency. Long return gives the annualized excess return to a long position in the currency held throughout the out-of-sample test period (buy-and-hold return).

Table 3
Estimated coefficients for the optimal linear forecasting model: $c = 0.0001$

<i>Lag</i>	<i>1</i>	<i>2</i>	<i>3</i>	<i>4</i>	<i>5</i>	<i>6</i>	<i>7</i>	<i>8</i>	<i>9</i>	<i>const</i>	<i>f</i>
DEM	0.89	0.11								0.0003	0.0001
JPY	0.85	0.15								0.0115	0.0001
CHF	0.77	0.18	0.05							0.0001	0.0001
GBP	0.78	0.16	0.05	0.00	0.01	0.02	-0.03	-0.01	0.02	-0.0009	0.0001

Note: columns 2 to 10 give the estimated lag coefficient for the best performing model over training and selection periods when one-way transaction cost was 0.0001. Column 11 records the constant and column 12 the optimal filter. Presenting more digits would show that all of the models imply stationary ARMA processes for the log exchange rate.

Table 4
Out-of-sample trading rule performance for the linear forecasting model
All-day trading

	<i>c</i>	<i>Annual Return</i>	<i>Number of trades</i>	<i>Break-even transaction cost</i>	<i>% long</i>
DEM	0.0000	92.68	3847	0.71	0.42
DEM	0.0001	79.30	640	3.63	0.37
DEM	0.0002	30.75	37	24.37	0.40
JPY	0.0000	94.28	2304	1.20	0.16
JPY	0.0001	73.03	811	2.64	0.12
JPY	0.0002	61.50	335	5.38	0.11
CHF	0.0000	137.64	4021	1.00	0.39
CHF	0.0001	161.28	1926	2.46	0.42
CHF	0.0002	111.48	996	3.28	0.40
GBP	0.0000	121.63	2898	1.23	0.79
GBP	0.0001	93.72	1024	2.68	0.81
GBP	0.0002	56.20	408	4.04	0.82

Note: Column 2 records the value of c , the one-way transaction cost used in training and selection periods. Column 3 gives the annualized per cent excess return over the seven-month out-of-sample test period calculated assuming zero transaction cost. Break-even transaction cost is the one-way transaction cost (in basis points) which reduces the annual excess return during the test period to zero. % long is the average percentage of the test period the rule held a position long the foreign currency.

Table 5
Out-of-sample trading rule performance for the equally weighted portfolio rule:
Restricted trading

	<i>c</i>	<i>Annual return</i>	<i>Number of rules</i>	<i>Number of trades</i>	<i>Break-even transaction cost</i>	<i>% long</i>
DEM	0.0000	3.60	13	591.38	0.18	52.59
DEM	0.0001	1.04	13	126.00	0.24	53.45
DEM	0.0002	-0.47	17	45.65	-0.30	52.39
JPY	0.0000	55.59	25	1952.84	0.83	47.63
JPY	0.0001	25.76	20	409.60	1.83	68.98
JPY	0.0002	8.06	18	123.00	1.91	64.40
CHF	0.0000	50.56	25	1750.60	0.84	43.57
CHF	0.0001	0.51	8	182.13	0.08	40.70
CHF	0.0002	6.35	8	109.38	1.69	65.61
GBP	0.0000	50.24	25	1608.32	0.91	54.75
GBP	0.0001	35.56	24	744.50	1.39	47.12
GBP	0.0002	9.60	10	67.30	4.16	27.87

Note: Trading was restricted to a twelve-hour period on weekdays. Periods for each currency were: DEM 0600-1800 GMT, JPY 0400-1600 GMT, CHF 0500-1700 GMT, and GBP 0500-1700 GMT. For further explanation see notes to Table 2.

Table 6
Out-of-sample trading rule performance for the linear forecasting model
Restricted trading

	<i>c</i>	<i>Annual Return</i>	<i>Number of trades</i>	<i>Break even transactions cost</i>	<i>% long</i>
DEM	0.0000	0.17	15	0.33	0.56
DEM	0.0001	0.17	15	0.33	0.56
DEM	0.0002	-6.10	13	-13.62	0.57
JPY	0.0000	30.98	1012	0.89	0.21
JPY	0.0001	1.06	128	0.24	0.10
JPY	0.0002	1.92	14	3.97	0.12
CHF	0.0000	44.09	1927	0.66	0.44
CHF	0.0001	42.37	984	1.25	0.41
CHF	0.0002	12.05	0	NA	1.00
GBP	0.0000	21.61	1360	0.46	0.69
GBP	0.0001	17.83	516	1.00	0.70
GBP	0.0002	-7.19	46	-4.54	0.84

Note: Trading was restricted to a twelve-hour period on weekdays. Periods for each currency were: DEM 0600-1800 GMT, JPY 0400-1600 GMT, CHF 0500-1700 GMT, and GBP 0500-1700 GMT. For further explanation see notes to Table 4.

Table 7
Consensus of trading rules identified by the genetic program: All-day trading

	<i>c</i>	0-20%	20-40%	40-60%	60-80%	80-100%
DEM	0.0000	50.57	0.05	0.00	3.17	46.21
DEM	0.0001	6.17	25.43	18.02	16.19	34.20
DEM	0.0002	4.09	3.39	76.37	9.22	6.93
JPY	0.0000	36.09	7.16	15.77	9.29	31.68
JPY	0.0001	20.70	29.55	15.49	19.28	14.97
JPY	0.0002	1.40	8.03	46.42	25.47	18.67
CHF	0.0000	40.39	8.75	0.89	11.21	38.76
CHF	0.0001	38.26	8.66	6.83	8.84	37.41
CHF	0.0002	12.42	31.85	42.57	9.13	4.03
GBP	0.0000	43.48	7.20	2.57	5.68	41.07
GBP	0.0001	32.00	14.56	13.52	12.60	27.32
GBP	0.0002	7.91	10.57	41.64	33.83	6.04


Note: the table reports the quintiles of the distribution of the proportion of all trading rules giving a long signal over the out-of-sample test period.

Table 8
Consensus of trading rules: Restricted trading

	<i>c</i>	0-20%	20-40%	40-60%	60-80%	80-100%
DEM	0.0000	1.70	28.98	30.54	31.85	6.93
DEM	0.0001	1.32	16.88	39.85	40.68	1.27
DEM	0.0002	0.00	23.53	35.09	41.38	0.00
JPY	0.0000	53.99	0.00	0.00	0.00	46.01
JPY	0.0001	0.08	13.20	29.76	24.75	32.22
JPY	0.0002	0.29	14.13	16.90	50.01	18.66
CHF	0.0000	55.26	1.67	0.10	7.76	35.21
CHF	0.0001	8.63	53.35	21.72	15.78	0.53
CHF	0.0002	0.00	4.65	35.63	43.32	16.40
GBP	0.0000	36.69	8.25	1.28	5.26	48.52
GBP	0.0001	43.60	6.58	1.46	11.88	36.48
GBP	0.0002	44.42	33.80	17.31	4.48	0.00

Note: the table reports the quintiles of the distribution of the proportion of all trading rules giving a long signal over the out-of-sample test period, with trading restricted as described in the notes to Table 5.

Figure 1: A simple trading rule



Notes: The rule signals a long position in foreign currency if the current exchange rate is greater than the 48-period moving average and the time of day (GMT) is between 0800 and 1600, and a short position otherwise.

Figure 2: Crossover and reproduction

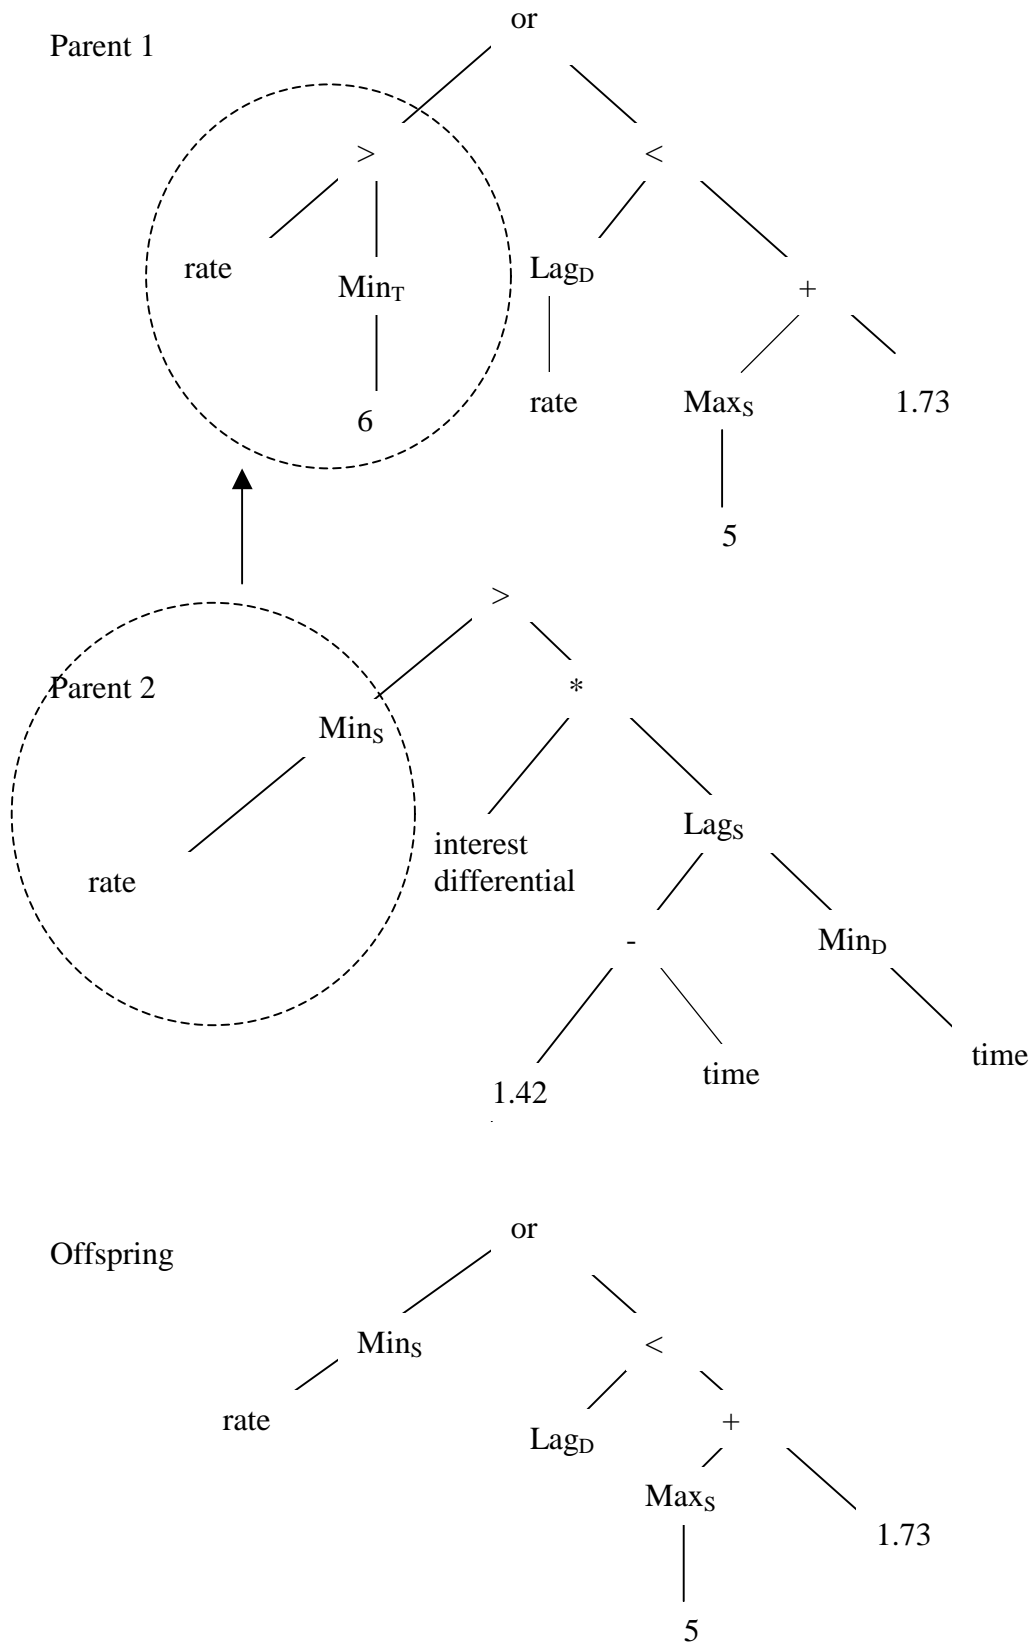
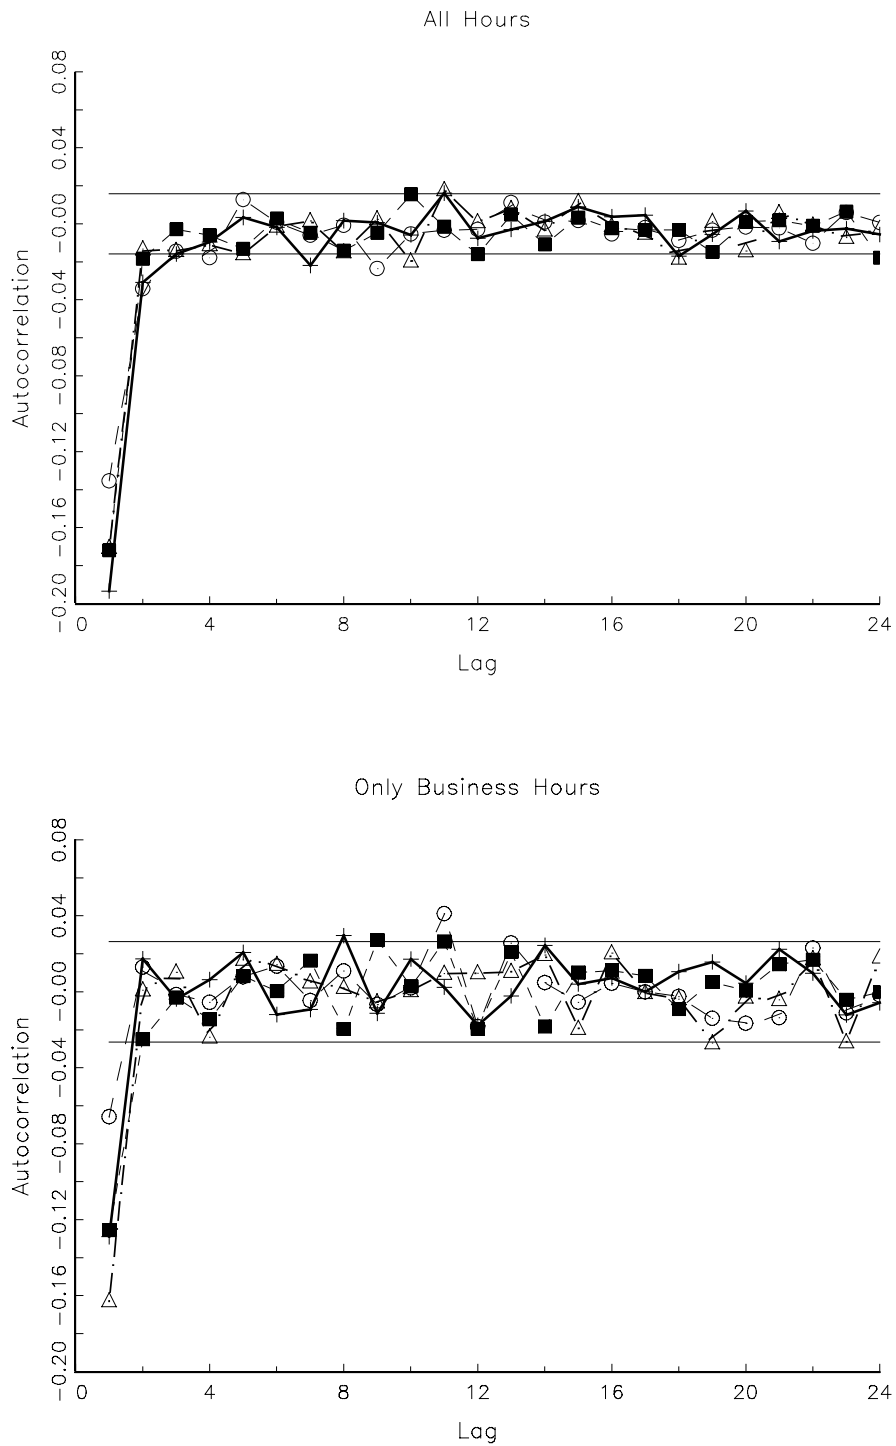


Figure 3: Autocorrelation coefficients for log exchange rate changes



Notes: The horizontal lines indicate the asymptotic 95% confidence interval for zero autocorrelation. The autocorrelation coefficients from the DEM, JPY, CHF and GBP are represented as circles, solid squares, triangles, and pluses, respectively.