Available online at www.sciencedirect.com
SCIENCE DIRECT
Journal of International Money and Finance
22 (2003) 223–237
Journal of International Money and Finance
www.elsevier.com/locate/econbase

Intraday technical trading in the foreign exchange market

C.J. Neely [a,*], P.A. Weller [b]

[a] *Research Department, Federal Reserve Bank of St. Louis, P.O. Box 442, St. Louis, MO 63166, USA*
[b] *Department of Finance, Henry B. Tippie College of Business Administration, University of Iowa, Iowa City, IA 52242, USA*

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 find, however, that the trading rules discover some remarkably stable patterns in the data.

JEL classification: F31; G15

Keywords: Technical analysis; Trading rule; Genetic programming; Exchange rate; High frequency

1. Introduction

There has been a recent resurgence of academic interest in the claims of technical analysis. This development is largely attributable to accumulating evidence that technical trading can be profitable over long time horizons. However, academic investigation of technical trading in the foreign exchange market has not been consistent with the practice of technical analysis. Most technical traders 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

* Corresponding author. Tel.: +1-314-444-8568; fax: +1-314-444-8731.
E-mail address: neely@stls.frb.org (C.J. Neely).

0261-5606/03/$ - see front matter © 2003 Elsevier Science Ltd. All rights reserved.
doi:10.1016/S0261-5606(02)00101-8

strategies on daily or weekly data (Dooley and Shafer, 1983; Sweeney, 1986; Levich and Thomas, 1993; Neely et al., 1997). These papers find that trading rules earn significant excess returns, net of transaction costs, which cannot be easily explained as compensation for bearing risk. The trading frequency for the rules studied in these papers typically ranges from 3 to 26 trades per annum. Evidently, these are not the trading strategies being used by the foreign exchange dealers surveyed by Taylor and Allen (1992), Cheung and Chinn (2000) and Cheung et al. (2000). These studies document the fact that technical analysis is widely used for trading at the shortest time horizons, namely, days and weeks, and that its use may be increasing.

But, despite their practical importance, there has been relatively little study of 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 firms over 1996–1998 and finds significant evidence of 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. Osler (Federal Reserve Bank of New York, unpublished, 2001) examines the potential importance of conditional orders for exchange rate dynamics and technical analysis. Acar and Lequeux (Banque Nationale de Paris, London Branch, unpublished, 1995) examine the profitability of a class of linear forecasting rules fitted to a sample of half-hourly data, whereas 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. Pictet et al. (Olsen & Associates, unpublished, 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. (Olsen & Associates, unpublished, 1998) report 3.6–9.6% annual excess returns, net of transaction costs, to proprietary real-time Olsen and Associates trading models using seven years of exchange rate data at a 5-minute frequency. It is difficult to compare other results with theirs, given that their models are not publicly available.

This paper follows trading practice more closely than most past research by investigating the performance of trading rules using high-frequency data that allow the rules to change position within the trading day[1]. We examine the performance of the trading rules to measure market efficiency, an approach first advocated in Brock et al. (1992), rather than to find profitable rules, per se. 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. Two distinct methodologies are employed: the first is a genetic program that can search over a very wide class of (possibly nonlinear) trading rules; the second consists of linear forecasting models, which provide natural benchmarks against which to compare the genetic programming results. The analysis does not specify the type of trader who might use such rules, but does assume that

[1] Of course, the papers discussed above—Goodhart and Curcio (1992), Osler (unpublished, 2001), Acar and Lequeux (unpublished, 1995), Curcio et al. (1997), Pictet et al. (unpublished, 1996) and Gençay et al. (unpublished, 1998)—do permit intraday trading of one form or another.

the trader faces reasonably low transaction costs. There is 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 reasonable 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.

We must qualify our results, however, by pointing out that failing to find profitable rules with these methods does not guarantee that such rules do not exist. Specifically, certain rules that are used in practice, such as those which exploit the tendency of support and resistance levels to cluster at round numbers, might be very difficult to generate using our methodology (Osler, 2001, unpublished). Indeed, even if it is theoretically possible that the genetic program could construct certain types of rules, experience using the technique on other problems has shown that lack of computational power or insufficient data may preclude the discovery of certain rules in practice.

2. 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). This use of the genetic program follows an approach first applied to the foreign exchange market by Neely et al. (1997). That paper and its working paper version (Neely and Weller, 2001a) provide more details on genetic programming.

An important advantage of genetic programming in constructing trading rules is that the method can use additional information to construct technical rules (Neely and Weller, 1999, 2001b). This exercise uses three information variables as input to the genetic program: 1) the normalized value of the exchange rate, which is the exchange rate divided by its moving average over the previous two weeks[2]; 2) the interest differential; and 3) the hour of the day. We include the hour because of the known intraday patterns in foreign exchange trading volumes. Such patterns are known to be associated with volatility, but may also affect the first moment of the exchange rate series.

The fitness criterion for the genetic program is the continuously compounded excess return to the trading rule. We train rules under two assumptions about when they can trade. The first scenario permits trading 24 hours a day, 7 days a week. Because of concerns about illiquidity during off-peak hours, the second scenario—called restricted trading—only permits trading during 12-hour periods of heavy trad-

[2] 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. Non-stationary data could produce such problems.

ing 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.

Each trade involves switching from a long to a short position or vice versa, and so incurs a round trip transaction cost ($2c$). 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:

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

where z_t is an indicator variable taking the value $+1$ (–1) for a long (short) position at t, r_t is the continously compounded return for holding a long position in foreign currency from t to $t+1$ and n is the number of trades from time zero to time T. Returns to rules subject to restricted trading include the interest differential for overnight positions as well as the exchange rate return.

The genetic program requires three separate subsamples, referred to as the training, selection and test periods. Training and selection are equivalent to an in-sample estimation period. The out-of-sample test period measures the performance of the rules trained and selected in the first two periods. The subsamples are as follows: training, 02/01/96–03/31/96; selection, 04/01/96–05/31/96; and test, 06/01/96–12/31/96. January 1996 was used to calculate starting values for functions requiring lagged data.

3. 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 first difference of the log exchange rate. The maximum lag length is 10. We then combine each estimated forecasting model with a filter to produce a trading rule. The filter reduces trading frequency and accompanying transaction costs for those periods in which only a small change in the exchange rate is predicted. Denoting the one-period-ahead forecast of the change in the log exchange rate at time t by $E_t(\ln(S_{t+1}))-\ln S_t$ and the filter by f, trading signals are determined in the following way:

$$\begin{aligned}
\text{If } z_{t-1} = +1,\ z_t &= -1, && \text{if } E_t(\ln(S_{t+1}/S_t)) < -f,\\
&= +1, && \text{if } E_t(\ln(S_{t+1}/S_t)) \geq -f.\\
\text{If } z_{t-1} = -1,\ z_t &= +1, && \text{if } E_t(\ln(S_{t+1}/S_t)) > f,\\
&= -1, && \text{if } E_t(\ln(S_{t+1}/S_t)) \leq f.
\end{aligned} \quad (2)$$

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 –*f*, the rule will maintain a long position.

Models with filters that range from 0 to 5 basis points in steps of 1 basis point and 1–10 lags of the independent variable are estimated on the training and selection periods. The in-sample excess return from the implied trading rule is calculated assuming the following three values of one-way transaction cost: 0, 1 and 2 basis points. The rule with the highest in-sample excess return for each level of transaction cost is then run on the out-of-sample test period[3].

4. 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. By excluding higher-frequency data, they also substantially reduce the risk of introducing microstructural artifacts (Lyons, 2001). 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 (US minus foreign contract) in the prices for the short-term interest rate futures contract closest to expiry. Because Japanese futures data were unavailable, only the US futures price was used for the JPY exchange rate. The US 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).

Table 1 presents summary statistics for the distributions of half-hourly log exchange rate changes. Standard deviations are quite similar across series, and all exchange rates display very high kurtosis. The top panel of Fig. 1 displays autocorrelations for the log returns, using all hours. There is highly significant negative first-order autocorrelation for all currencies, which is also present in both bid and ask returns and robust to excluding outliers in the bid-ask spread. Excluding off-peak hours—i.e., measuring autocorrelation only during business hours—reduces the mean first-order autocorrelation from –0.17 to –0.12 (see the bottom panel of Fig. 1). We shall return to discuss possible causes and implications of the autocorrelation in the conclusion.

[3] Including the lagged normalized exchange rate, the hour of the day, and the interest futures differential produced very similar results and so we omit them for brevity.

Table 1
Summary statistics[a]

	Mean	Std. Dev.	Skew	Kurt	ρ(1)	ρ(2)	ρ(3)	Min	Max
DEM	0.00022	0.0719	–0.07	25.64	–0.14	–0.03	–0.01	–0.93	0.97
JPY	0.00050	0.0794	–0.05	14.16	–0.17	–0.02	0.00	–0.90	0.92
CHF	0.00063	0.0934	–0.23	31.73	–0.17	–0.01	–0.01	–1.59	1.62
GBP	–0.00071	0.0704	0.27	34.14	–0.19	–0.03	–0.02	–1.20	1.22

[a] The table presents statistics for log exchange rate changes constructed from the full data set, consisting of 16,080 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 would be distributed as standard normal variables if the underlying series were normal. $\rho(i)$ records the autocorrelation coefficient at lag i. Min and max record the smallest and largest half-hourly percentage changes over the sample period.

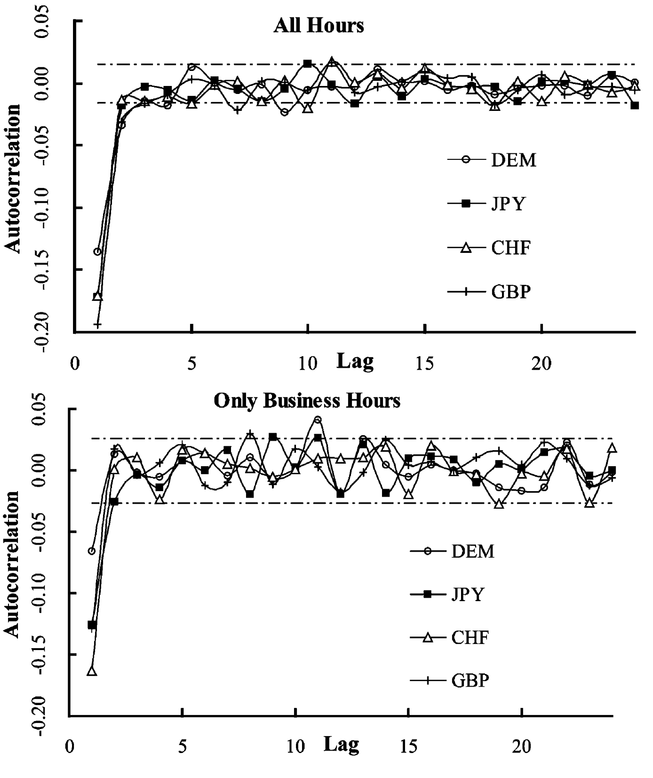

Fig. 1. The top panel shows the autocorrelation coefficients estimated using all data while the bottom panel shows the same statistics using only business hours. The horizontal lines indicate the asymptotic 95% confidence interval for zero autocorrelation. The autocorrelation coefficients from the DEM, JPY, CHF and GBP are represented as circles, solid squares, triangles, and pluses, respectively.

5. Results

We consider first the unrestricted case in which trading is allowed to take place 24 hours a day, seven days a week. We have strong doubts about whether such a trading strategy was achievable at the prices quoted, given that we permit trading when major markets are closed, trading activity is reduced, and transaction costs are higher. Nevertheless, we consider the 24-hour, seven-day trading rule results to be a useful benchmark to measure predictability in the data and with which to compare results from the restricted trading case.

For each currency we generated 25 rules from the genetic program under each of three assumptions about transaction costs in training and selection periods. We used one-way transaction costs of 0, 1 and 2 basis points[4]. From those 25 rules we selected those which had a positive excess return during the selection period and also traded at least once.

An equally weighted portfolio rule aggregates each set of 25 rules by apportioning 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 3 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 6 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 transaction costs in-sample produce very high returns, over 100% per annum in three of the four cases. 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 exchange rates, that for the British pound, is 1.01 basis points for a one-way trade. This value is largely attributable to the lower trading frequency of these rules.

As the transaction cost in the training and selection periods increases from 0 to 1 and then 2 basis points, both annual excess returns before transaction cost and trading frequency fall sharply. But break-even transaction cost rises uniformly 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 (positive returns and number of trades), most notably in the case of the GBP, where only five of the 25 rules satisfied the criteria for $c = 2$. 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[5]. Since the break-even transaction cost can be interpreted as the average excess return

[4] 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 transactions costs suggest that 2–2.5 basis points for a one-way trade is realistic for recent large transactions (Neely et al., 1997).

[5] The number of trades and break-even 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 (which has no effect on the net open position), the portfolio does not trade.

Table 2
Out-of-sample performance for the all-day-trading equally weighted portfolio rule[a]

	c	Annual return	Number of rules	Number of trades	Break-even transaction cost	% long	Long Return
DEM	0	66.92	25	4908.76	0.40	45.69	2.30
DEM	1	46.09	21	887.57	1.51	60.19	
DEM	2	6.30	19	88.58	2.08	54.46	
JPY	0	130.56	25	4164.44	0.91	48.40	11.72
JPY	1	43.28	23	451.57	2.80	45.30	
JPY	2	16.30	13	144.69	3.28	60.33	
CHF	0	127.48	25	4846.88	0.77	50.02	11.51
CHF	1	92.40	25	1773.96	1.52	50.46	
CHF	2	30.99	15	388.60	2.33	45.98	
GBP	0	132.34	25	3830.92	1.01	49.62	–15.80
GBP	1	111.18	25	1920.96	1.69	48.49	
GBP	2	31.59	5	412.00	2.24	63.60	

[a] The equally weighted portfolio rule attaches a weight (1/# of rules) to each rule satisfying the selection criteria. The columns provide the following information: *c* reports the one-way transaction cost used in training and selection periods—in basis points. Annual Return provides the annualized percent excess return over the seven-month out-of-sample test period, calculated assuming zero transaction cost. Number of rules shows how many rules (out of 25) that produced a positive in-sample excess return before transactions costs and also traded. These rules were used for the out-of-sample test. Number of trades reports the quantity of test-period trades. 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 (approximately) as $100\cdot(^{7}/_{12})$*annual return*/(2·*number of trades*). % long is the proportion of the test period the rules held a long position in the foreign currency. Long return gives the annualized excess return to a long position in the foreign currency held throughout the out-of-sample test period.

per zero-cost trade, this demonstrates that the search procedure can identify rules that can 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 and sets quotes accordingly 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 1 basis point) over the training and selection periods. The coefficients on the first lag of the return data are consistent with the statistics on serial correlation. The selected filters for three of the currencies are 1 basis point—matching the chosen level of transaction cost—the exception being the DEM, for which the chosen filter size was 2 basis points.

Table 3
Estimated coefficients for the optimal linear forecasting model: $c = 1$ basis point[a]

Lag	1	2	3	4	5	6	7	const x 10^{-4}	filter
DEM	–0.112							0.081	2
JPY	–0.154							0.044	1
CHF	–0.233	–0.053						0.132	1
GBP	–0.218	–0.061	–0.014	–0.015	–0.009	0.010	–0.012	–0.011	1

[a] Columns 2–8 give the estimated lag coefficient for the best performing model over the training and selection periods when one-way transaction cost was 1 basis point. Column 9 reports the constant. Column 10 reports the optimal filter in basis points.

When we consider the out-of-sample performance of the autoregressive forecasting model (see Table 4), there is a similar pattern of change in out-of-sample zero-cost returns, trading frequency and break-even costs as the in-sample transaction cost increases. For example, the DEM zero-cost annual return decreases from 102–40%, the number of trades declines from 3611 to 61 and the break-even cost increases from 0.83 to 19.15 basis points as the in-sample one-way transaction cost increases from 0 to 2 basis points. However, the autoregressive results are clearly superior to those derived from the genetic program at the highest level of transaction cost. Specifically, the break-even transaction cost is higher in all cases, dramatically so in the case of the DEM, where it is 19.15 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 "$c = 1$" and "$c = 2$" trading rules in Table 4 earn excess returns net of

Table 4
Out-of-sample performance for the all-day-trading linear forecasting model[a]

	c	Annual return	Number of trades	Break-even transaction cost	% long
DEM	0	102.35	3611	0.83	59.9
DEM	1	67.44	640	3.09	59.2
DEM	2	39.84	61	19.15	58.1
JPY	0	148.09	1504	2.89	52.4
JPY	1	148.09	1504	2.89	52.4
JPY	2	75.80	456	4.87	51.4
CHF	0	157.16	1958	2.35	54.7
CHF	1	157.16	1958	2.35	54.7
CHF	2	117.79	1036	3.33	53.9
GBP	0	158.76	4530	1.03	46.3
GBP	1	142.60	1534	2.73	50.2
GBP	2	84.42	574	4.31	50.9

[a] The notes to Table 2 describe the content of the columns.

transaction cost that exceed 37% per annum in all cases. However, this may be an unrealistic conclusion, given our previously stated concerns about illiquidity during off-peak hours.

For this reason we generate a new set of rules under the assumption that trading is restricted to a 12-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 transaction 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 $c = 2$ basis points, we find a break-even transaction cost of 4.16 basis points. One should be cautious about reading too much into this finding. There was a relatively small number of good rules identified in-sample; they traded infrequently and tended to be skewed toward short positions.

Table 5
Out-of-sample performance for the restricted-trading equally weighted portfolio rule[a]

	c	Annual return	Number of rules	Number of trades	Break-even transaction cost	% long
DEM	0	3.60	13	591.38	0.18	52.59
DEM	1	1.04	13	126.00	0.24	53.45
DEM	2	–0.47	17	45.65	–0.30	52.39
JPY	0	55.59	25	1952.84	0.83	47.63
JPY	1	25.76	20	409.60	1.83	68.98
JPY	2	8.06	18	123.00	1.91	64.40
CHF	0	50.56	25	1750.60	0.84	43.57
CHF	1	0.51	8	182.13	0.08	40.70
CHF	2	6.35	8	109.38	1.69	65.61
GBP	0	50.24	25	1608.32	0.91	54.75
GBP	1	35.56	24	744.50	1.39	47.12
GBP	2	9.60	10	67.30	4.16	27.87

[a] Trading was restricted to a 12-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. The notes to Table 2 describe the content of the columns.

In Table 6 we show the results of imposing restricted trading on the autoregressive-forecasting model under two methods of estimating coefficients. The first method constructs rules by estimating coefficients on 24-hour, in-sample data but only permits the rules to change positions during business hours. During non-business hours,

Table 6
Out-of-sample performance for the restricted-trading linear forecasting model[a]

Coefficients estimated using all data	c	Annual return	Number of trades	Break-even transaction cost	% long
DEM	0	1.92	14	3.98	69.9
DEM	1	1.92	14	3.98	69.9
DEM	2	1.92	14	3.98	69.9
JPY	0	58.90	1997	0.86	54.2
JPY	1	53.29	682	2.27	51.1
JPY	2	-3.35	17	-5.72	48.1
CHF	0	59.05	1963	0.87	51.4
CHF	1	37.21	954	1.13	53.8
CHF	2	16.37	180	2.64	46.9
GBP	0	47.67	1749	0.79	46.3
GBP	1	37.59	772	1.41	45.8
GBP	2	4.87	42	3.36	56.4

Coefficients estimated using only business hour returns	c	Annual return	Number of trades	Break-even transaction cost	% long
DEM	0	11.92	1997	0.17	64.6
DEM	1	−2.26	11	−5.96	46.8
DEM	2	−2.26	11	−5.96	46.8
JPY	0	41.13	1587	0.75	31.7
JPY	1	−4.54	20	−6.59	25.8
JPY	2	−4.54	20	−6.59	25.8
CHF	0	48.38	1814	0.77	52.3
CHF	1	34.24	1036	0.96	54.5
CHF	2	14.31	198	2.10	48.9
GBP	0	45.45	1962	0.67	49.0
GBP	1	33.33	867	1.12	45.3
GBP	2	−8.91	24	−10.77	61.6

[a] The notes to Tables 2 and 5 describe the content of the columns and the restricted trading hours, respectively. While both panels present results that restrict trading to business hours, the top panel uses coefficients estimated on all data and the bottom panel uses coefficients estimated using only business hour data.

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. Those results are shown in the top panel. Again, we see that, as in-sample transactions costs rise, annual zero-cost returns and numbers of trades fall and the break-even transaction cost tends to rise. In the case of the DEM, the choice of model was insensitive to the level of transaction cost. As with the unrestricted trading cases (Tables 2 and 4), in most cases the restricted-trading autoregressive model (top panel Table 6) has higher break-even transaction costs than the restricted-trading genetic program (Table 5).

The second method of imposing restricted trading on the autoregressive-forecasting model estimates coefficients only on business hour data and, again, only permits the rules to change positions during business hours. These results, shown in the bottom panel of Table 6, are generally worse in terms of break-even cost than either those of the restricted hour genetic programming (Table 5) or the autoregressive model using all hours (top panel of Table 6). As before, the zero-cost annual returns and number of trades fall as in-sample transactions costs rise. But the break-even transactions costs no longer covary predictably with the level of in-sample transactions costs.

The fact that the trading rules identified by the genetic program often perform less well—by the metric of break-even transaction cost—than those generated by the autoregressive-forecasting model deserves some comment, given that the genetic program could construct autoregressive rules or approximations of the rules. Two factors enable the autoregressive rules to outperform the more flexible genetic programming rules out-of-sample. First, the variables in addition to the exchange rate series that were provided as input to the genetic program proved to be uninformative. This is suggested by the fact that the inclusion of these variables in the forecasting model did not make any difference[6]. 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.

6. Discussion and conclusion

We find very stable predictable components in 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 times of normal market activity. Rules based on the autoregressive forecasting model per-

[6] Experiments in which the separate data series used as inputs to the genetic program were randomized support this supposition. Only randomizing the exchange rate series over the in-sample period produced any significant change in out-of-sample performance.

form about as well or better than those found by the genetic program, 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 nonsynchronous trading is responsible for the negative autocorrelation in hourly data. This explanation is implausible for our data set. For example, the nonsynchronous trading model of Lo and Mackinlay (1990) implies negative autocorrelation several orders of magnitude smaller than that actually observed. Zhou (1996) suggests that negative autocorrelation in tick-by-tick data is a consequence of "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[7].

Further, it is not clear to us 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, 2 to 3 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 12-hour window of high volume, break-even transaction costs are considerably reduced.

It is interesting that the foreign exchange market displays very different characteristics at different trading horizons. 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, at least for the sample period we consider. 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

[7] 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. The authors thank Charles Goodhart and Richard Payne for their assistance in obtaining the D2000–2 data.

their positions at the end of each day than by those who take open positions with horizons of weeks or months.

Acknowledgements

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. Michael Dempster, Clemens Kool, two anonymous referees and seminar participants at the Judge Institute, Cambridge University, the Seventh Annual Global Finance Conference at DePaul University, Maastricht University and the Seventh International Conference for Forecasting Financial Markets in London provided helpful comments. Paul Weller thanks 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.

References

Acar, E., Lequeux, P., 1995. Trading rules profits and the underlying time series properties. Unpublished manuscript, Federal Reserve Bank of New York.

Baillie, R.T., Bollerslev, T., 1991. Intra-day and inter-market volatility in foreign exchange rates. Review of Economic Studies 58 (3), 565–585.

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

Cheung, Y.W., Chinn, M.D., 2000. Currency traders and exchange rate dynamics: a survey of the US market. Journal of International Money and Finance.

Cheung, Y.W., Chinn, M.D., Marsh, I.W., 2000. How do U.K.-based foreign exchange dealers think their market operates? NBER Working Paper 7524.

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. 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. Harper Business, pp. 43–69.

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., 2001. The Microstructure Approach to Exchange Rates. MIT Press, Cambridge, MA.
Melvin, M., 1997. International Money and Finance. Addison Wesley Publishing, Reading, MA.
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.
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., 2001a. Intraday technical trading in the foreign exchange market. Federal Reserve Bank of St. Louis Working Paper 99-016B.
Neely, C.J., Weller, P.A., 2001b. Technical analysis and central bank intervention. Journal of International Money and Finance 20 (7), 949–970.
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.
Osler, C.L., 2001. Conditional orders and exchange-rate dynamics: Explaining the success of technical analysis. Unpublished manuscript, Federal Reserve Bank of New York.
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.
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.