Transactions Data Tests of Efficiency: An Investigation in the Singapore Futures Markets

MAHENDRA RAJ

The aim of this article is to test the profitability of technical trading rules in the intra-day currency futures market. A wide range of technical strategies are applied to tick data over a two-year period for two currency futures—Japanese Yen (JY) and Deutschemark (DM)—traded in the Singapore International Monetary Exchange. The study finds that after incorporating transactions costs and testing for the significance of the profits using a bootstrap methodology, none of the technical trading systems produce significant returns. © 2000 John Wiley & Sons, Inc. Jrl Fut Mark 20:687–704, 2000

INTRODUCTION

Evidence of the profitability of technical-trading rules in financial markets is documented in several recent studies. The basic premise underlying technical analysis is that markets move in trends and that these trends persist. This suggests some sort of serial dependency in prices. If

For correspondence, Dr. Mahendra Raj, Aberdeen Asset Management Professor of International Finance, Aberdeen Business School, Robert Gordon University, Viewfield Road, Aberdeen AB15 7AW, United Kingdom; e-mail:m.raj@rgu.ac.uk.

Received June, 1999; Accepted February, 2000

■ *Mahendra Raj is the Aberdeen Asset Management Professor of International Finance at Aberdeen Business School at Robert Gordon University in the United Kingdom.*

technical-trading rules can be used to obtain abnormal profits, then it would indicate that the markets are not weak-form efficient.

The aim of this research article is to look at the profitability of technical trading strategies in the foreign-exchange futures market, and by doing so, the article examines the efficiency of the market. Transaction prices of two currency futures contracts (Deutschemark and Japanese Yen) traded in the Singapore International Monetary Exchange over the period February 1, 1992 to December 31, 1993 were used for the study. Few studies have examined the intra-day efficiency of these markets. In this article, the buying and selling prices are taken to be the prices immediately following the generation of a signal by the trading system. The significance of the intra-day profits generated using numerous trading rules is tested using the bootstrap methodology developed by Efron (1979). In addition, the profitability of these rules is analyzed after adjusting for transaction costs and a risk premium. The existence of statistically significant profits from these trading rules would indicate that exchange rates are not serially independent drawings from a stationary distribution, i.e., exchange rates do not follow a random walk.

Hsieh (1988) examined the statistical properties of the closing bid prices of five foreign currencies from 1974 to 1983. He found that the distributions of exchange rates differ across days of the week, and that their means and variances also shift over time. His results suggest the non-normality of exchange rates as the distributions are unimodal and have high coefficients of kurtosis. However, he found that, once heteroskedasticity is accounted for, there is no evidence of serial correlation in the data. Engel and Hamilton (1990) also found little evidence of linear serial dependence in exchange-rate changes, supporting the conclusion that exchange rates follow random walks. However, they found evidence of nonlinear serial dependence characteristic of long swings. These long swings could possibly be exploited profitably by mechanical trading rules.

The statistical properties of foreign exchange futures have not been examined as thoroughly as have those of the spot markets. So (1987) showed that the distribution of currency futures returns is nonstationary. Glassman (1987) found no evidence of autocorrelation once limit days in futures markets were removed. Kao and Ma (1992) found evidence of short-term price dependence of a nonlinear form in foreign currency futures. Levich and Thomas (1993) found no evidence of serial correlation once heteroskedasticity was adjusted for.

Neftci (1991) investigated the statistical properties of technical analysis to determine whether these had any empirical basis. He suggested that the rules followed by technical analysts might capture information contained in higher-order moments of asset prices. On testing the moving-average rule on the Dow-Jones Industrials over the period 1911 through 1976, he found evidence supporting nonlinearity in the series.

Dooley and Shafer (1983) reported the profits using filter rules for nine currencies using daily spot rates over the period from 1973 through 1981. They incorporated the interest expense and interest income of long and short positions, as well as transaction costs. Their results indicate that small filters (x = 1, 3, 5%) would have been profitable for all currencies over the entire sample period, although the statistical significance of these profits was not examined.

Sweeney (1988) looked at the profitability of filter rules of sizes x = 0.5, 1, 2, 3, 4, 5, and 10% on daily exchange rates for ten currencies over the period of April 1973 through December 1980. He stated that these rules resulted in profits in 80% of the cases, with the profits being superior for the smaller filters.

Lucac, Brorsen, and Irwin (1988) examined twelve trading systems on commodity futures contracts. Seven of the systems generated positive gross profits and four exhibited positive net profits even after adjusting for risk.

Levich and Thomas (1993) tested filter and moving-average rules on five currency futures over the period from 1976 to 1990 and implemented a testing procedure based on bootstrap methodology. They concluded that these technical trading strategies were profitable and that these profits were significant.

Taylor (1994) found that channel rules could have been used successfully on currency futures from 1982 to 1990 to generate positive returns. Ntungo and Boyd (1998) compared trading systems based on neural networks with ARIMA models and concluded that both yielded positive returns of the same order.

This article examines the profitability of several moving-average, filter, and channel rules in the intra-day Singapore futures market using a bootstrap methodology. Though a number of rules generate positive gross returns, once adjustments are made for transactions costs and risk premia, none of the rules is found to be profitable.

DATA AND TRADING RULES

Data

The Singapore International Monetary Exchange (SIMEX) supplied the data used in this study. SIMEX is also linked to the Chicago Mercantile

Exchange (CME) enabling Eurodollar, Japanese Yen, Deutschemark, and British Pound futures traded on one exchange to be transferred to or liquidated at the other exchange.

The data series used are the Japanese Yen and Deutschemark futures prices from February 1, 1992 to December 31, 1993. Unlike previous studies, the data used in this study is intra-day data consisting of all trades occurring over the period of study. The number of price observations obtained was 9948 for the Japanese Yen contract and 14,124 for the Deutschemark.

In order to generate a single futures series for each of the currency futures, successive contracts were brought together. As both the Japanese Yen and Deutschemark contracts have four contract months—March, June, September, and December—the single futures series was generated by using the March Contract for December, January, and February, the June contract for March, April, and May, the September contract for June, July, and August, and the December contract for September, October, and November. Therefore, the series is constructed by combining the 3 months prior to delivery month for each contract. The use of futures contracts eliminates the need for overnight interest rates on spot interbank deposits.

Technical Trading Rules

The three mechanical trading rules used in this study to test the foreign exchange market efficiency are the filter rule, the moving-average rule, and the channel rule. These three rules as described in Murphy (1986) are easy to understand and commonly used. Variations of these three trading rules form nine of the twelve trading systems tested by Lucac et al. (1988). Ntungo and Boyd (1998) found that more advanced techniques like neural networks do not necessarily perform better than simpler ARIMA models. Implementation of these three techniques is straightforward and hence, in this article, they are selected over other more complicated rules. These three rules are defined below.

Filter Rule

When the current futures contract price rises above the most recent local minimum by a predetermined amount, the rule generates a buy signal taking a long position. When the current futures price falls below the most recent local maximum by a predetermined amount, the rule generates a sell signal taking a short position. A local minimum is taken to be a point where the two prices to either side of it are both greater than the price at that point, whereas a local maximum is defined as the point where the two prices to either side of it are both less than the price at that point.

The filter rules used in this study are the 0.5, 1, 2, 3, 4, 5, and 10% rules.

Moving-Average Rule

The moving-average rule (also known as the moving-average crossover rule) is a mechanical trading rule that generates buy and sell signals through the use of two moving averages—a long-period average and a short-period average. When the short-term moving average intersects the long-term moving average from below, a buy signal is generated, and when the short-term moving average intersects the long-term moving average from above, a sell signal is generated. A band also can be introduced to the moving-average rule to reduce the number of false signals caused by price volatility. When the band is introduced, signals are generated when the short-term moving average moves above or below the long-term moving average by an amount larger than the band. If the short-term moving average stays within the band, no signal is generated and the position remains the same.

Since transaction data is used in this study, the periods for this rule are the transaction periods. As a result, a 3/150 moving-average rule means that the short-period average is a 3-period one, whereas the long period average is a 150-period one. Incorporating a 1-% band implies that the 3-period short average has to be above or below the 150-period long average by 1%. The moving-average rules used in this study are 1/5, 1/50, 1/100, 1/150, 1/200, 3/150, 5/150, 2/200, 3/200, and 5/200. In addition, all of these rules were tested with the addition of a 1-% band.

Channel Rule

Under the channel rule, a long futures position is replaced by a short position when the price is less than the minimum price during the preceding L trades. Likewise, a short futures position is replaced by a long futures position when the price is higher than the maximum price during the preceding L trades.

The channel rules used in this study are with L = 5, 50, 150, and 200 trades.

All of the above rules imply that, once the initial signal is generated, an investor always will have an open position in the market, either long or short.

METHODOLOGY

Application of Trading Rules

The application of these trading rules in this study differs from prior investigations. In previous studies that have used closing prices, it was assumed that if a signal was generated by one of the rules, i.e., in the case of a 5-% filter rule, if the closing price moves 5% above the previous local minimum, a buy signal is generated, then it also can be assumed that the investor can buy or sell the contract at the closing price generating the signal. Returns then are measured from that signal until a reverse signal is made. It is unlikely, however, that traders would be able to buy at the exact closing price when the signal was generated rather than at some time following that signal. In this study, as the data used is intraday data, it can be assumed that once either a buy or a sell signal is generated, the actual buy or sell is initiated at the next traded price rather than the one initiating the signal.

"Returns to the trading rules" can be measured by:

$$R_{t,t+1} = \ln\,(F_{t+1}/F_t)$$

where F_t = currency futures price at time t, t = time of trade following signal.

Hence, these returns are measured from the time of actually purchasing or selling a futures contract, not from when the signal was generated. Returns can be measured this way as futures prices reflect the contemporaneous interest differential between the foreign currency and the U.S. dollar.

By applying the technical trading rules to the original currency futures series the overall *gross returns* to these rules can be measured. Total returns are measured as follows:

$$TR_{1,N+1} = \sum_{t=1}^{N} R_{t,t+1} = \sum_{t=1}^{N} d_t \ln\,(F_{t+1}/F_t)$$

where $d_t = 1$ is a long futures position and $d_t = -1$ is a short futures position.

These returns are the gross returns to the trading rules. To adjust the returns for a risk premium, the study has followed Sweeney (1986) in using a constant-risk premium over the analysis period. He estimated this constant-risk premium to be equal to the returns to a buy-and-hold strategy over the analysis period. This was calculated as follows:

$$RP_{1,N+1} = \sum_{t=1}^{N} \ln (F_{t+1}/F_t)$$

As the trading rules generate both buy and sell signals, the actual risk premium to be used in adjusting the returns to the rules depends on the time the rules are long or short. As this study uses intra-day data, the proportion of total trades long or short were used to determine the actual risk premium, rather than the time long or short. The adjusted risk premium is as follows:

$$ARP = (1 - f)\ RP - fRP$$

where f = proportion of trades short and $(1 - f)$ = proportion of trades long.

In order to determine the net returns of the trading rules, the gross returns also need to be adjusted for transaction costs. The transaction costs used are those estimated by Levich and Thomas (1993). They estimated that the cost of transacting in the currency futures market is 2.5 basis points per transaction for a large institution. However, they used a more conservative estimate of 4 basis points, which also will be used in this study. Transaction costs will have the greatest impact on the small trading rules where greater numbers of signals are generated. The *net returns* to these trading rules then can be calculated by deducting the adjusted-risk premium and transaction costs from the gross returns, as shown below:

$$TNR = TR - ARP - TC$$

Bootstrapping and Significance Testing

Bootstrapping provides a way of testing the significance of the profits from these trading rules without making any assumptions about the distribution of the process underlying exchange rates. As prior research has shown that exchange rates are not random independent drawings from a stationary distribution, significance testing through the use of t-tests is meaningless, as they assume normal stationary and time independent distributions. The use of a bootstrap methodology makes none of these assumptions.

The present study will use a bootstrap methodology to test the significance of the trading-rule profits.[1] In using the bootstrap methodology, new 'simulated' series of exchange-rate data are generated by taking random shufflings of the original series (sampling with replacement). By generating new series from random samples of the original data series, the simulated series have the same distributional properties as the original time series. For this study, 500 new 'simulated' series were generated for each of the trading rules.[2]

Each of the technical trading strategies then was applied to each of the 500 simulated series and the profits calculated, creating distributions of profits. The original profits then could be compared against the simulated profits in order to determine their significance. The null hypothesis and alternative hypothesis are as follows:

H_0: there is no information in the original time series
H_1: there is information in the original time series.

RESULTS

Summary Statistics

Table I presents the summary statistics for the Deutschemark (DM) and Japanese Yen (JY) Future contracts. Both of the mean tick returns are very close to zero, with the mean for the DM being negative as it fell over the period examined. The standard deviations for the two contracts are remarkably similar.

The coefficient of skewness for JY is slightly positive, giving the distribution a slight rightwards appearance. However, for the DM, it is significantly negative, pushing the distribution to the left. Both of the currency futures exhibited extreme positive excess kurtosis, making the distributions leptokurtic. This is consistent with prior analysis on the statistical properties of exchange rates. In comparison to a normal distribution with a coefficient of kurtosis of 0, these are extreme figures. The

[1]Levich and Thomas (1993) used a bootstrap methodology to test the significance of trading rule profits. However their simulated series were generated by sampling *without* replacement—better known as a jackknife. Bootstrapping involves sampling *with* replacement and provides better estimates of standard errors (see LePage & Billard, 1992).

[2]500 simulations were used due to the enormous amount of computing time involved. However this does not represent a limitation on the results as Brock, Lakonishok, and LeBaron (1992) show there is no significant difference between using 500 simulations or more than 500 simulations.

TABLE I

Summary Statistics of Intra-Day Returns (Logarithmic) on Selected Foreign Currency Futures Contracts

Currency Futures Contract	*Japanese Yen (JY)*	*Deutschemark (DM)*
Sample N	9497	14123
Mean	.000012449	−0.0000049392
Standard Deviation	0.0014648	0.0014924
Average Number of Trades per Day	19	29
Coefficient of Skewness	0.1485	−1.3995
-Standard Deviation	0.0251	0.0206
Coefficient of excess kurtosis	141.2949	107.3779
-Standard Deviation	0.0503	0.0412
Minimum Return	−0.038194	−0.039675
Maximum Return	0.038194	0.025967

Note: N = number of logarithmic returns.

maximum tick return for the JY contract was 0.038 or 3.9% and for the DM contract, 0.026 or 2.6%.

Autocorrelation Tests

Autocorrelation coefficients were estimated for 20 lags (not shown in the article) for both the DM and JY. All of the coefficients are close to zero, taking on both positive and negative values. None of these coefficients are significant at the 5% level using a chi^2 distribution with degrees of freedom equal to 1.

A further test for autocorrelation that was undertaken is the Ljung–Box–Pierce statistic. This statistic is calculated as follows:

$$Q = N(N + 2) \sum_{j=1}^{J} \frac{1}{N - j} \hat{p}_j^2 \text{ for } J = 1, \ldots, p$$

where p_j is the jth-order autocorrelation of the residuals, N is the number of observations, and j is a preselected number of autocorrelations. The Q statistics at lags 5, 10, 15, and 20 (not shown) are estimated. Again, none of these are significant at the 5% level. The Durbin h-statistic, a large sample test for first-order serial correlation, when lagged dependent variables were present, also was calculated. These statistics also were not significant at the 5% level, suggesting that serial correlation was not a problem in the data.

TABLE II

Gross Returns for Filter Rules on Original Exchange-Rate Series
(per cent per annum)

	Filter Size (%)							
Futures Contract	*0.5*	*1*	*2*	*3*	*4*	*5*	*10*	*Average Profit*
Japanese Yen	−2.7	2.8	−12.6	−8.0	0.0	0.0	0.0	−2.9
Deutschemark	−1.8	−2.6	−15.2	7.7	7.7	0.0	0.0	−0.6

Autocorrelation Tests Adjusted for Heteroskedasticity

In order to ensure the validity of these tests, we tested for heteroskedasticity. The results of these heteroskedasticity tests (not reported) showed highly significant heteroskedasticity for the JY, while heteroskedasticity was not present to a significant degree in the DM. Only the Goldfield–Quandt statistic was significant for the DM returns, whereas all three tests were highly significant for the JY returns.

The Ljung–Box–Pierce statistics were adjusted for heteroskedasticity using White's Heteroskedastic Consistent Matrix, which corrects the estimates for heteroskedasticity (Granger & Newbold, 1986). The result of adjusting the statistics for heteroskedasticity makes them even less significant. This is particularly noticeable for JY, where the estimate of the standard error had increased from 0.01023 to 0.066680, whereas for the DM, it increased from 0.008414 to 0.009261, having little effect on the results.

Returns of Trading Rules on Original Exchange Rates Series

Filter Rules

Results for the filter rules are presented in Table II. For both DM and JY, the filter rules generated negative returns in the majority of instances with the average profit for JY being −2.9% and −0.6% for DM. In 2 out of the 3 cases where the returns were positive, they were generated on only one signal by the filter rules. Thus, they were generating returns only equivalent to the market (buy-and-hold strategy). The small filters traded most frequently with the 0.5% rule, generating 127 trades for JY and 160

TABLE III

Net Returns for Filter Rules on Original Exchange-Rate Series

A. Japanese Yen *Filter Rule (%)*	*Gross Returns*	*Risk Premium*	*Transaction Costs*	*Net Returns*
0.5	−2.7%	0.8%	2.8%	−6.3%
1	2.8%	−0.2%	0.5%	2.5%
2	−12.6%	−2.7%	0.1%	−10.0%
3	−8.0%	−5.1%	0.0%	−2.9%
4	0.0%	0.0%	0.0%	0.0%
5	0.0%	0.0%	0.0%	0.0%
10	0.0%	0.0%	0.0%	0.0%

B: Deutschemark *Filter Rule (%)*	*Gross Returns*	*Risk Premium*	*Transaction Costs*	*Net Returns*
0.5	−1.8%	−0.4%	3.5%	−5.0%
1	−2.6%	0.1%	1.0%	−3.6%
2	−15.2%	1.8%	0.2%	−17.2%
3	7.7%	2.1%	0.0%	5.6%
4	7.7%	2.1%	0.0%	5.6%
5	0.0%	0.0%	0.0%	0.0%
10	0.0%	0.0%	0.0%	0.0%

for DM. The higher filters did not generate any trade, as the changes in tick prices were not large enough to generate signals.

Table III presents the net returns for the filter rules for JY and DM. The majority of the net returns remained negative for JY. The risk premium appears to get more negative the greater the size of the filter, suggesting that these rules were generating the wrong signals. The 3-% filter rule incurs a −5.11% risk premium, pulling its net return closer to zero. Transaction costs had the greatest effect on the 0.5-% rule, which incurred transaction costs of 2.77%. The results for the DM were similar to JY, with the majority being negative. Only the 3- and 5-% rules generated positive net returns. However, this is based only on one trade over the two-year period. The average profit for JY was −2.39% and −2.08% for DM.

Moving-Average Rules

Table IV presents the gross results for the moving-average rules. The first half of the table presents the returns from the moving-average rules without a band, while the second half of the table presents the returns from the moving-average rules using a 1-% band. All of the trading rules ap-

TABLE IV

Gross Returns for Moving-Average Rules on Original Exchange-Rate Series (%/annum)

	Moving-Average Rule (Short-Term (Trades)/Long-Term (Trades))										
Futures Contract	*1/5*	*1/50*	*1/100*	*1/150*	*1/200*	*3/150*	*5/150*	*2/200*	*3/200*	*5/200*	*Average Profit*
JY	8.5	3.0	4.6	5.5	5.0	4.3	5.5	3.7	4.7	4.9	5.03
DM	137	−7.5	2.0	1.0	1.5	1.1	1.0	−0.8	−0.6	0.8	0.995
	Moving-Average Rule with 1% Band (Short-Term (Trades)/Long-Term (Trades))										
Futures Contract	*1/5*	*1/50*	*1/100*	*1/150*	*1/200*	*3/150*	*5/150*	*2/200*	*3/200*	*5/200*	*Average Profit*
JY	−2.2	1.9	1.4	8.6	9.7	6.5	8.5	9.3	8.9	5.5	5.8
DM	9.2	10.3	0.3	−2.0	0.2	−1.0	−2.0	−0.1	−0.4	−2.0	1.2

peared to be profitable for the Japanese Yen, with an average gross profit of 5.03% for the 2-year period. However, these rules traded extremely frequently. The 1/5 rule traded over 2000 times in the 2-year period. For the DM, the rules generated both positive and negative returns, with an average profit of 0.995%. Again, these rules traded over 3000 times with the 1/5 rule.

Introducing the 1-% band had a significant effect on the number of trades generated by the rules. In the case of JY, the 1/5 rule generated only 9 trades compared to over 2000 without the band. Apart from the 1/5 rule, the moving-average rules generated positive returns for the JY, giving an average profit of 5.8%. Five of the ten rules generated positive returns for the DM, with an average profit of 1.2%.

The net returns for the strategies based on moving-average rules after incorporating transactions costs and risk premia are presented in Table V. The gross returns for JY for the rules without a band were reduced significantly once transactions costs were adjusted for. This was particularly noticeable for the 1/5 rule, with the net return being −38% compared to the gross return of 8.45%. The risk premium had little effect on the gross returns, as the rules were in both long and short positions for equivalent numbers of trades. The introduction of the 1-% band significantly reduced the impact of transaction costs. The risk premium also had little effect except for the 1/5 rule. The average profit of the moving average rules for JY with the 1-% band was 2.36% compared to −7.53% without the band.

TABLE V

Net Returns for Moving-Average Rules for Original Exchange-Rate Series

A: Japanese Yen (0-% band)					*(1-% band)*			
Moving Average Rule	*Gross Returns*	*Risk Premium*	*Transaction Costs*	*Net Returns*	*Gross Returns*	*Risk Premium*	*Transaction Costs*	*Net Returns*
1/5	8.5%	−0.1%	46.5%	−38.0%	−2.0%	3.5%	0.2%	−5.9%
1/50	3.0%	0.1%	9.6%	−6.6%	1.9%	0.0%	0.6%	1.3%
1/100	4.6%	0.5%	6.3%	−2.2%	1.4%	0.3%	0.6%	0.4%
1/150	5.5%	0.7%	2.9%	1.9%	8.5%	−0.2%	0.5%	8.3%
1/200	5.0%	0.6%	4.2%	0.1%	9.7%	−0.5%	0.5%	9.7%
3/150	4.3%	0.7%	3.4%	0.2%	6.5%	−0.2%	0.5%	6.1%
5/150	5.5%	0.7%	2.9%	1.9%	8.5%	−0.2%	0.5%	8.3%
2/200	3.7%	0.7%	3.4%	−0.4%	9.3%	−0.5%	0.5%	9.3%
3/200	4.7%	0.7%	3.0%	1.0%	8.9%	−0.5%	0.5%	9.0%
5/200	4.9%	0.7%	2.8%	1.4%	5.5%	−0.5%	0.5%	5.5%

VB: Deutschemark (0-% band)					*(1-% band)*			
Moving Average Rule	*Gross Returns*	*Risk Premium*	*Transaction Costs*	*Net Returns*	*Gross Returns*	*Risk Premium*	*Transaction Costs*	*Net Returns*
1/5	13.7%	0.0%	71.7%	−58.0%	9.2%	1.0%	0.5%	7.7%
1/50	−7.0%	0.1%	12.4%	−20.0%	10.3%	0.7%	0.8%	8.8%
1/100	−2.0%	0.1%	7.9%	−10.0%	0.3%	0.4%	1.0%	−1.1%
1/150	9.9%	0.0%	3.8%	−2.9%	−2.0%	0.5%	1.0%	−3.5%
1/200	1.5%	0.0%	5.8%	−4.3%	0.2%	0.4%	0.9%	−1.2%
3/150	1.1%	0.0%	4.6%	−3.6%	−1.0%	0.5%	0.9%	−2.4%
5/150	1.0%	0.0%	3.9%	−2.9%	−2.0%	0.5%	0.9%	−1.8%
2/200	−7.6%	0.0%	4.5%	−5.3%	−0.1%	0.4%	4.5%	−5.0%
3/200	−0.6%	0.0%	4.0%	−4.7%	−0.4%	0.4%	0.9%	−1.9%
5/200	0.8%	0.0%	3.6%	−2.8%	−2.0%	0.4%	0.9%	−3.3%

The results were similar for the DM, with transaction costs significantly reducing gross returns. Again, this was most significant for the 1/5 rule. All of the net returns for the DM without a band were negative. Once the 1-% band was introduced, transaction costs had a relatively minor effect on the gross returns, with the net returns being only slightly less than the gross returns. The average profit of the moving-average rules for DM with a 1-% band was 0.96% compared to −16.75% without the band.

TABLE VI

Gross Returns for Channel Rules on Original Exchange-Rate Series
(per cent per annum)

	Channel Rule (Trades)				
Future Contract	*5*	*50*	*150*	*200*	*Average Profit*
Japanese Yen	−5.9	5.4	10.0	4.6	3.5
Deutschemark	16.3	−1.5	4.5	7.3	6.6

TABLE VII

Net Returns for Channel Rules on Original Exchange-Rate Series
(per cent per annum)

A: Japanese Yen Channel Rule	*Gross Returns*	*Risk Premium*	*Transaction*	*Net Returns*
5	−5.9%	0.1%	22.4%	−28.5%
50	5.4%	0.4%	2.7%	2.3%
150	10.0%	0.9%	0.8%	8.3%
200	4.6%	0.8%	0.7%	3.1%

B: Deutschemark Channel Rule	*Gross Returns*	*Risk Premium*	*Transaction*	*Net Returns*
5	16.3%	−0.0%	33.0%	−16.7%
50	−1.5%	0.0%	3.8%	−5.4%
150	4.5%	0.1%	1.2%	3.1%
200	7.3%	−0.0%	0.9%	6.4%

Channel Rules

The results for strategies based on channel rules are presented in Table VI. The average profit for JY was 3.5% and 6.6% for DM. The channel rule based on 5 trades gave a very high return of 16.33% for JY. However, this rule generated 1028 trades over the 2-year period, incurring transaction costs far in excess of the return. The larger channel rules, which traded less frequently, also generated positive profits.

Table VII presents the net returns for the channel rules. The transaction costs and risk premium had little effect on the gross returns for both DM and JY except for the channel rule based on 5 trades and to a lesser extent the rule based on 50 trades. The average net return for JY was 4.6% without including the rule based on 5 trades (−3.7% with the

TABLE VIII

Ranking of Original Filter-Rule Profits in 500 Simulations

Futures Contract	*Filter Size (%)* 0.5	1	2	3	4	5	10	*Average Rank*
Japanese Yen	188	287	25	88	108	n/a	n/a	139
Deutsche mark	209	189	30	393	n/a	n/a	n/a	205

N.B. n/a means these filter rules did not generate any signals under the simulated distributions

TABLE IX

Ranking of Original Channel-Rule Profits in 500 Simulations

Futures Contract	*Channel Rule (Trades)* 5	50	150	200	*Average Rank*
Japanese Yen	108	356	427	341	308
Deutsche mark	456	220	323	354	338

rule). The average net return for DM is 1.4% without the rule based on 5 trades (−3.1% with the rule).

Significance Testing

The above section presented the gross and net returns to the technical trading rules over the period February 1, 1992 to December 31, 1993. This section looks at determining whether any of these returns, in fact, are significant by utilizing the bootstrap methodology. Tables VIII, IX, and X report the ranks of the trading-rule profits for the actual series in comparison to the 500 randomly generated series. Each of the trading rules was applied to the 500 randomly generated series. The profits generated by applying the trading rule to each of the 500 random series then were ranked. This was done for each of the trading rules. The ranking of the profits generated by applying the trading rule to the original series in relation to the profits of the rule applied to the 500 random series are shown in these tables. A rank greater than 475 would imply that the trading rule generates profit that is significant at the 5-% level while a rank greater than 495 would mean that the trading-rule profit is significant at 1-% level.

TABLE X

Ranking of Original Moving-Average-Rule Profits in 500 Simulations

	Moving-Average Rule (Short-Term (Trades)/Long-Term (Trades))										
Futures Contract	*1/5*	*1/50*	*1/100*	*1/150*	*1/200*	*3/150*	*5/150*	*2/200*	*3/200*	*5/200*	*Average Rank*
JY	453	123	220	271	274	286	273	231	226	266	262
DM	427	326	346	360	377	354	352	316	358	360	358
	Moving-Average Rule with 1-% Band (Short-Term (Trades)/Long-Term (Trades))										
Futures Contract	*1/5*	*1/50*	*1/100*	*1/150*	*1/200*	*3/150*	*5/150*	*2/200*	*3/200*	*5/200*	*Average Rank*
JY	393	406	268	215	255	393	417	243	409	360	308
DM	161	195	271	405	431	246	216	486*	237	206	338

*Significant at the 5-% level

Only one out of the entire set of technical trading rules generated returns was significant at 5% level. This was the 2/200 moving-average rule with a 1-% band, having a rank of 486. However, this was not significant at the 1% level (495 and above). In general, the ranks of the trading-rule profits vary over the possible range, indicating that the trading rules not only did not produce significant returns, they also did not produce consistent returns. Therefore, we cannot reject the null hypothesis that there is no information in the original series in favor of the alternative hypothesis at the 1-% level.

The results above indicate that the two exchange-rate series analyzed, DM and JY, follow random walks having serial independence. This serial independence was confirmed through tests of autocorrelation.

CONCLUSIONS

The intent of this article was to extend prior research on the profitability of technical trading rules in the foreign exchange futures market. A wide range of technical strategies was applied to intra-day data over a two- year period for two currency futures—Japanese Yen (JY) and Deutschemark (DM). Prior studies have focused only on closing prices, using these prices as both the signal and the price at which futures contracts can be bought and sold. Not only have intra-day data been used, but also buying and selling prices have been taken to be the prices following the signals generated by the trading systems. This is a more realistic assumption as it is unlikely that investors can transact at the same price as the trigger price.

The results of this article do not support the hypothesis that there is information in the original exchange-rate series that can be profited on through the use of technical trading strategies. None of the technical trading systems produced significant returns (with the exception of one rule, which was significant at the 5-% level), and there was no consistency in the rankings of rules in relation to the simulated series. This suggests that exchange rates follow a random walk in that there is serial independence in the data series. This was backed up with tests of autocorrelation that found no evidence of serial correlation.

The results of this paper contradict those of previous studies (Dooley & Schafer 1983; Lucac et al., 1988; Sweeney, 1988) and in particular that of Levich and Thomas (1993), who found considerable evidence supporting the profitability of technical trading strategies. The discrepancies between the findings of this article as compared to that of Levich and Thomas (1993) may stem from the differences between the two data sets used. Levich and Thomas (1993) used closing prices as compared to the transaction data used in the present study—and in the application of the trading rules. Levich and Thomas (1993) did note, however, that the profitability of the trading rules they used declined over the period of their study. As this study looks at the two years following their study, it is possible that technical trading systems have ceased to be profitable. This may be the result of increased computerization. As more and more people gain access to computers, the use of technical trading rules becomes more widespread, resulting in the elimination of any profitable opportunities and making the markets efficient.

BIBLIOGRAPHY

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.

Dooley, M. P., & Shafer, J. (1983). Analysis of short-run exchange rate behaviour, March 1973–Nov 1973. In D. Bigman and T. Taya (eds.), Exchange rate and trade instability. Ballinger.

Engel, C., & Hamilton, J. D. (1990). Long swings in the dollar: Are they in the data and do markets know it? American Economic Review, 80(4), 689–713.

Efron, B. (1979). Bootstrap methods: Another look at the jackknife. The Annals of Statistics, 7(1), 1–26.

Glassman, D. (1987). The efficiency of foreign exchange futures markets in turbulent and non-turbulent periods. Journal of Futures Markets, 7, 245–267.

Granger, C. W. J., & Newbold, P. (1986). Forecasting economic time series (2nd ed.). Academic.

Hsieh, D. A. (1988). The statistical properties of daily foreign exchange rates: 1974–1983. Journal of International Economics, 24, 129–145.

Kao, G. W., & Ma, C. K. (1992). Memories, heteroskedasticity, and price limit in currency futures markets. Journal of Futures Markets, 12(6), 679–692.

LePage, R., & Billard, L. (eds.). (1992). Exploring the limits of bootstrap. New York: Wiley.

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, 451–474.

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

Murphy, J. J. (1986). Technical analysis of futures markets. New York: New York Institute of Finance.

Neftci, S. N. (1991). Naive trading rules in financial markets and Wiener–Kolmogorov prediction theory: A study of technical analysis. Journal of Business, 64(4), 549–571.

Ntungo, C., & Boyd, M. (1998). Commodity futures trading performance using neural network models versus ARIMA models. Journal of Futures Markets, 18(8), 965–983.

So, J. (1987, Febraury). The sub-Gaussian distribution of currency futures: Stable paretian or nonstationary. Review of Economics and Statistics, 100–107.

Sweeney, R. J. (1986). Beating the foreign exchange market. Journal of Finance, 41, 163–182.

Sweeney, R. J. (1988). Some new filter rule tests: Methods and results. Journal of Financial and Quantitative Analysis, 23(3), 285–300.

Taylor, S. J. (1994). Trading futures using a channel rule: A study of the predictive power of technical analysis with currency examples. Journal of Futures Markets, 14(2), 215–235.