

Is Technical Analysis in the Foreign Exchange Market Profitable? A Genetic Programming Approach

Christopher Neely, Paul Weller, and Rob Dittmar*

Abstract

Using genetic programming techniques to find technical trading rules, we find strong evidence of economically significant out-of-sample excess returns to those rules for each of six exchange rates over the period 1981–1995. Further, when the dollar/Deutsche mark rules are allowed to determine trades in the other markets, there is significant improvement in performance in all cases, except for the Deutsche mark/yen. Betas calculated for the returns according to various benchmark portfolios provide no evidence that the returns to these rules are compensation for bearing systematic risk. Bootstrapping results on the dollar/Deutsche mark indicate that the trading rules detect patterns in the data that are not captured by standard statistical models.

I. Introduction

In its simplest form, technical analysis uses information about historical price movements, summarized in the form of price charts, to forecast future price trends. This approach to forecasting originated with the work of Charles Dow in the late 1800s, and is now widely used by investment professionals as input for trading decisions. Technical analysts argue that their approach to trading allows them to profit from changes in the psychology of the market. This view is well expressed in Pring's ((1991), pp. 2–3) statement:

*Neely and Dittmar, Research Department, Federal Reserve Bank of St. Louis, St. Louis, MO 63011; Weller, Department of Finance, College of Business Administration, University of Iowa, Iowa City, IA 52242. The authors would like to thank Kent Koch for excellent research assistance and their colleagues at the Federal Reserve Bank of St. Louis for generously sharing their computers at night and during weekends. Jon Faust, Tom George, Michael Melvin, Maureen O'Hara (associate editor and referee), and seminar participants at the Federal Reserve Bank of St. Louis, the University of Iowa, Arizona State University, and the Midwest Finance Association made helpful comments. The views expressed are those of the authors and do not necessarily reflect official positions of the Federal Reserve Bank of St. Louis or the Federal Reserve System.

The technical approach to investment is essentially a reflection of the idea that prices move in trends which are determined by the changing attitudes of investors toward a variety of economic, monetary, political and psychological forces. . . Since the technical approach is based on the theory that the price is a reflection of mass psychology (“the crowd”) in action, it attempts to forecast future price movements on the assumption that crowd psychology moves between panic, fear, and pessimism on one hand and confidence, excessive optimism, and greed on the other.

Although technical analysis was originally developed in the context of the stock market, its advocates argue that it is applicable in one form or another to all asset markets. Since the era of floating exchange rates began in the early 1970s, this approach to trading has been widely adopted by foreign currency traders. In a recent survey of major dealers in the foreign exchange market in London, Taylor and Allen (1992) found that, at short horizons of one week or less, 90% of respondents reported the use of some chartist input, with 60% stating that they regarded such information at least as important as economic fundamentals. Part of the explanation for this state of affairs is to be found in the unsatisfactory predictive performance of models of the exchange rate based upon fundamentals. This assessment is succinctly summarized by Frankel and Rose (1994), p. 29 who state that “the case for macroeconomic determinants of exchange rates is in a sorry state. . . (The) results indicate that no model based on such standard fundamentals like money supplies, real income, interest rates, inflation rates and current-account balances will ever succeed in explaining or predicting a high percentage of the variation in the exchange rate, at least at short- or medium-term frequencies.”

Despite its long history, technical analysis and its claims have traditionally been regarded by academics with a mixture of suspicion and contempt. This attitude was not without justification because its proponents never made serious attempts to test the predictions of the various techniques employed. However, a renewal of academic interest in such forecasting techniques has been sparked by accumulating evidence that financial markets may be less efficient than was originally believed. Foreign exchange markets have proved to be more volatile than was anticipated at the beginning of the floating rate era in the early 1970s, and the “long swings” in the dollar observed in the 1980s have not been satisfactorily explained in terms of movements in economic fundamentals.

Several studies have sought to document the existence of excess returns to various types of trading rules in the foreign exchange market (Dooley and Shafer (1983), Sweeney (1986), Levich and Thomas (1993), Osler and Chang (1995)). These papers find that a class of trading rules makes economically significant excess returns in a variety of currencies over different time periods; however, these results are difficult to interpret. Because the rules considered in these studies are selected for examination *ex post*, there is an inevitable risk of bias. These investigations have deliberately concentrated on the most common and widely used rules, but some doubt remains as to whether the reported excess returns could have been earned by a trader who had to make a choice about what rule or combination of rules to use at the beginning of the sample period.

In this paper, we address this problem by using a genetic program as a search procedure for identifying optimal trading rules. We obtain rules for a variety of currencies in a given sample period (1975–1980), and then examine their performance over the period 1981–1995. The advantage of this approach, and the most important contribution of the paper, is that it enables us to construct a true out-of-sample test of the significance of the excess returns earned by the trading rules. We find strong evidence of economically significant excess returns after transaction costs, not only for currencies against the dollar, but also for Deutsche mark/yen and British pound/Swiss franc. The rules we identify are similar to those commonly used by technical traders. We also document a marked difference in the structure of the best rules we find for dollar markets and for the Deutsche mark/yen market, as indicated by the trading frequency of the rules in the different markets. Such differences in structure have not previously been identified.

To throw light on the possible source of the observed excess returns, we first calculate betas for the returns to the various trading rules using a range of benchmark portfolios, but find no evidence of significant systematic risk associated with use of these rules. We then consider a subsample of the rules obtained for the dollar/Deutsche mark and perform bootstrapping simulations using random walk, ARMA, and ARMA-GARCH models. It is clear that the trading rules detect patterns in the data that are not produced by these models.

The paper is organized as follows. Section II reviews the previous work on trading rules in the foreign exchange market. Section III discusses the implementation of the genetic program. Section IV presents the rules found by the algorithm, and analyzes their characteristics. Section V assesses the significance of the results by means of bootstrapping simulations. Section VI discusses the results and draws conclusions.

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

The trading rules that have been most intensively investigated in the foreign exchange market use filters and moving averages. A simple filter rule takes the form: “Go long when the exchange rate (dollar value of foreign currency) rises by $x\%$ above its previous local low; sell when it falls $x\%$ below its previous local high.” A moving average rule compares a short- to a long-run moving average, producing a buy signal when the short-run moving average cuts the long-run moving average from below. Even these simple rules can take a huge variety of forms. For example, the moving average rules will depend on the time windows chosen for each moving average. The filter rules will depend on filter size and the time window over which the previous high or low is calculated. The two classes of rules have obvious similarities in that they both seek to identify changes in a trend.

In one of the earliest studies to consider technical trading in the foreign exchange market, Dooley and Shafer (1983) focused on filter rules. They reported evidence of substantial profits to all but the largest filters over the period 1973–1981 for the Deutsche mark, yen, and pound sterling. Sweeney (1986) also looked at filter rules. Using daily dollar/Deutsche mark data over the period 1975–1980, he reported excess profits over buy-and-hold of 4% per annum for a 0.5% filter.

Neither study could reliably assess the significance of these profits. The bootstrap technique can be used to address this problem. It was first used by Brock, Lakonishok, and LeBaron (1992) to judge the significance of technical trading rule profits in the context of the stock market. Levich and Thomas (1993) adopted a similar approach in a study in which they used futures prices for a number of foreign currencies to examine the profits earned by various moving average and filter rules over the period 1976–1990. With a random walk as the generating process, they conducted bootstrapping simulations to assess the significance of their results, and concluded that the profit levels generated by the trading rules would have been highly unlikely to have been observed. Kho (1996), however, reports results suggesting that a substantial amount of the profitability of moving average rules in foreign currency markets can be explained by a time-varying risk premium. Osler and Chang (1995) is the first paper we are aware of that aims to evaluate the profitability of the head-and-shoulders pattern, one of the more complex, but very widely used, signals in the technical analyst's toolkit. They construct a computer algorithm to identify this pattern, and look at the returns to this rule in several currencies over the period 1973–1994. With bootstrap methodology, they find evidence of significant profits for the mark and yen, but not for the pound sterling, Canadian dollar, French franc, or Swiss franc. Trading in all six currencies simultaneously would have yielded significant profits even after transactions costs, although these profits were lower than typically reported for the moving average and filter rules.

All the studies described above use a range of rules chosen *ex post*. As Brock, Lakonishok, and LeBaron (1992) acknowledge, the dangers of biasing the results in such a situation are unavoidable. They argue that such dangers are minimized by the deliberate choice of a simple class of rules that has been in common use for a long period of time. But there still remains a significant amount of latitude in choosing the exact form of the rule. For this reason, we investigate the question of whether the data themselves can reveal the form of an "optimal" trading rule. For this, one needs a way of searching efficiently over the space of possible trading rules, a task for which genetic programming is well suited. While we do not claim that use of this procedure completely eliminates all potential bias, because we are forced to limit the search domain to some degree, we do argue that it substantially reduces the risk of such bias. Similar arguments were made by Allen and Karjalainen (1995), who were the first to use genetic programming to identify profitable trading rules in the stock market. After taking account of transactions costs, they found evidence that there were rules that were able to earn economically significant excess returns over a buy-and-hold strategy during the period 1970–1989.

All of the studies cited above that have documented the existence of profitable trading rules in the foreign exchange market have confined attention to exchange rates against the dollar. As a check on the robustness of our techniques, we also examine two cross rates, yen/Deutsche mark and Swiss franc/pound.

III. The Genetic Program

There are many different ways to specify the class of trading rules over which to search. Perhaps the simplest would be to treat the problem as one of forecasting the exchange rate one step ahead. However, it is well established that both linear and non-linear forecasting models perform rather poorly out-of-sample (Diebold and Nason (1990), Meese and Rose (1991)). An alternative would be to use a non-parametric approach such as an artificial neural network. This still requires that the structure of the network be specified in advance. The advantage of the genetic programming approach is that it allows one to be relatively agnostic about the general form of optimal trading rule, and to search efficiently in a non-differentiable space of rules. It also has the attractive feature that one can build in to the search procedure the relevant performance criterion directly in the form of the measure of fitness. This is important because the forecasting problem is not equivalent to finding an optimal trading rule, although the two are clearly linked. A profitable trading rule may forecast rather poorly much of the time, but perform well overall because it is able to position the trader on the right side of the market during large moves.

Genetic algorithms are computer search procedures based on the principles of natural selection originally expounded in Darwin's theory of evolution. These procedures were developed by Holland (1975) and extended by Koza (1992). They have been applied to a variety of problems in a diverse range of fields and are most effectively used in situations where the space of possible solutions to a problem is too large to be handled efficiently by standard procedures, or when the space is in some sense "badly behaved," e.g., non-differentiable, or possessing multiple local extrema. The essential ingredients of genetic programming are: i) a means of representing potential solutions as character strings that can be recombined to form new potential solutions, and ii) a fitness criterion that measures the "quality" of a candidate solution. In Holland's original genetic algorithm, potential solutions were encoded as fixed length character strings. Koza's extension, referred to as genetic programming, permits explicitly hierarchical variable length strings. In this paper, we use genetic programming to search for optimal technical trading rules, and encode these rules in the form of non-recombining trees. The basic steps in constructing a genetic program are as follows:

1. Create an initial randomly generated population of trees.
2. Calculate the fitness of each tree in the initial population according to a suitable criterion.
3. Create a new population by applying the following operations:
 - i) Copy existing individuals to the new population.
 - ii) Randomly select a pair of existing trees and recombine subtrees from them to produce a new tree.

The operations of reproduction and recombination are carried out with the probability of selection for the operations skewed toward selecting individuals with higher levels of fitness.¹

4. Calculate the fitness of each individual in the new population.
5. Repeat these operations, keeping a record of the overall fittest individual.

Holland demonstrated a remarkable theorem for character strings of fixed length. The operation of reproduction and genetic recombination on successive generations causes the candidate solutions to grow in the population at approximately a mathematically optimal rate. Here, optimality is defined by viewing the search strategy as a set of multi-armed bandit problems for which the optimal search strategy is known.² The approximation is closer in the case of problems whose solutions can be built up from subcomponents that are “small” in a suitably defined sense.

A striking feature that contributes to the efficiency of the algorithm is *implicit parallelism*. The algorithm is explicitly parallel in the sense that it simultaneously processes information about a large population of candidate solutions. But, it is possible to show that the operations of reproduction and genetic recombination lead one to expect solutions unrepresented in a given population to be generated in the new population in proportion to their fitness.

In genetic programming, in contrast to genetic algorithms, the individual structures that are subjected to selection and adaptation are explicitly hierarchical character strings of variable rather than fixed length. These structures can be represented as non-recombining decision trees, whose non-terminal nodes are functions and whose terminal nodes are variables or constants that serve as arguments of the functions.

In this paper’s application, we represent trading rules as trees, and make use of the following function set in constructing them:

1. arithmetic operations: plus, minus, times, divide, norm, average, max, min, lag;
2. Boolean operations: and, or, not, greater than, less than;
3. conditional operations: if-then, if-then-else;
4. numerical constants;
5. Boolean constants: true, false.

Norm returns the absolute value of the difference between two numbers, and average, max, min, and lag operate on a time window specified by the numerical argument of the function, rounded to the nearest whole number. Thus, max(3) at time t is equivalent to $\max(p_{t-1}, p_{t-2}, p_{t-3})$, lag(3) at time t is equal to p_{t-3} , and avg(3) is the arithmetic mean of p_{t-1} , p_{t-2} , and p_{t-3} .

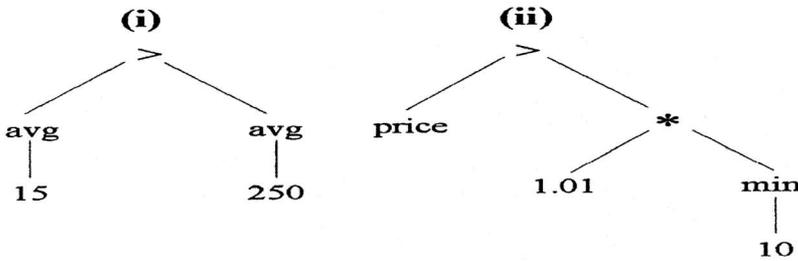
Figure 1 presents examples of two simple trading rules. Rule (i) signals a long position if the 15-day moving average is greater than the 250-day moving average,

¹In some applications, an additional mutation operation is introduced at this stage. Koza (1992) presents evidence to suggest that its impact is insignificant relative to that of recombination, and we choose not to implement it here.

²In a multi-armed bandit problem, the agent has to choose each period between one of n alternatives with constant expected payoffs that are unknown a priori.

otherwise a short position. Rule (ii) signals a long position if the closing exchange rate has risen by more than 1% above its minimum over the previous 10 days. The price function returns the average of bid and ask quotes at a fixed time on the day in question. Figure 2 illustrates the recombination 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. The operation is subject to the restriction that the resulting tree must be a well defined rule. In addition, we place a limit on the maximum possible size (100 nodes) and depth (10 levels) of a rule.

FIGURE 1



Trees representing (i) a moving average rule, (ii) a filter rule. Rule (i) signals long if the 15-day moving average exceeds the 250-day moving average, short otherwise. Rule (ii) signals long if the current exchange rate has risen by at least 1% above its minimum over the last 10 days, short otherwise.

We now turn to the form of the fitness criterion. The rules we examine switch between long and short positions in the foreign currency. We suppose that some amount is held in dollars and is reinvested daily at the domestic overnight interest rate. This can be thought of as the margin held against borrowing an amount equal in value, either in dollars or the foreign currency. If the trading rule signals a long position in the foreign currency at date t , the borrowed dollars are converted to foreign currency at the closing rate for date t and earn the foreign overnight rate. We denote the exchange rate at date t (dollars per unit of foreign currency) by S_t , and the domestic (foreign) overnight interest rate by $i_t(i_t^*)$. Then the gross excess return R_{t+1} (over the overnight rate in the U.S.) is given by

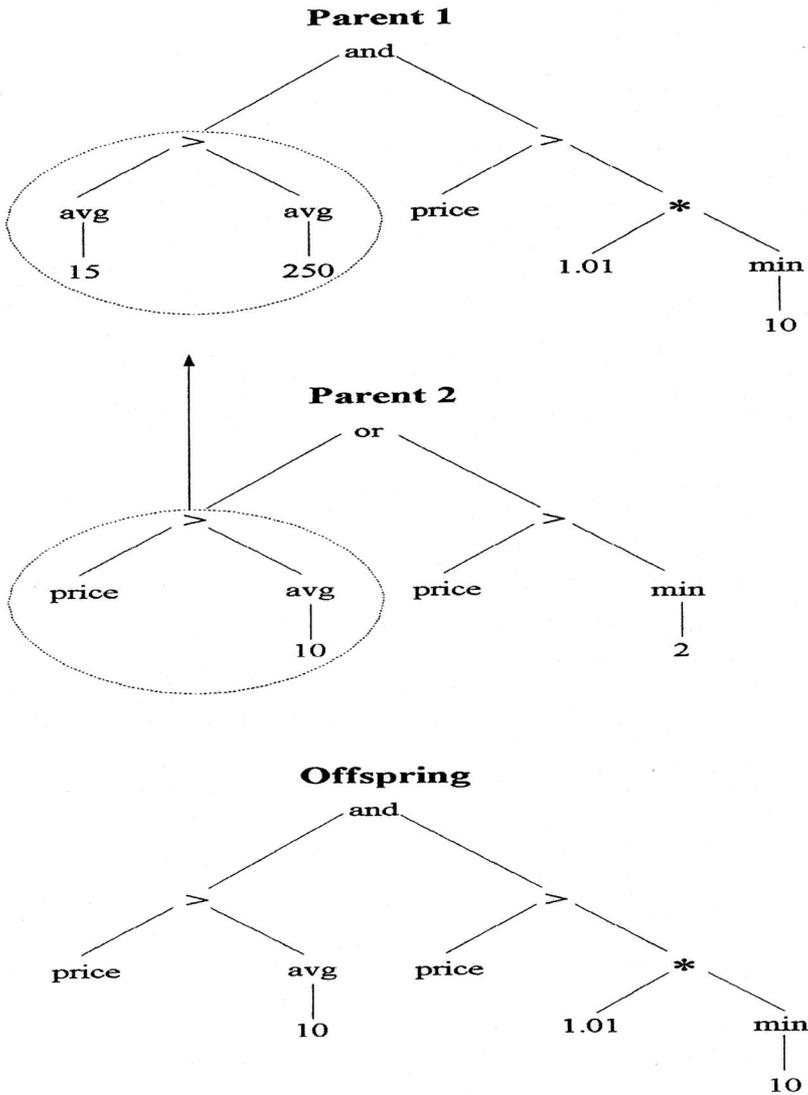
$$(1) \quad R_{t+1} = \frac{S_{t+1}}{S_t} \frac{(1 + i_t^*)}{(1 + i_t)}.$$

For a short position, the excess return is

$$(2) \quad R_{t+1} = 2 - \frac{S_{t+1}}{S_t} \frac{(1 + i_t^*)}{(1 + i_t)}.$$

The form of this expression reflects the fact that the rule now specifies the following transaction: borrow an amount of foreign currency equal in value to the dollar holding and convert to dollars to earn the U.S. overnight interest rate.

FIGURE 2
The Recombination Operation



The recombination operation involves selecting random subtrees from each parent rule and replacing one with the other to create a new rule, labeled offspring.

We denote the continuously compounded (log) excess return by $z_t r_t$, where z_t is an indicator variable taking the value +1 for a long position and -1 for a short position and r_t is defined as³

$$(3) \quad r_t = \ln S_{t+1} - \ln S_t + \ln(1 + i_t^*) - \ln(1 + i_t).$$

³This introduces an approximation error in the case of a short position, but it is very small—less than 0.25% of the total cumulative excess return in all cases.

The cumulative excess return from a single round-trip trade (go long at date t , go short at date $t + k$), with one-way proportional transaction cost c , is

$$(4) \quad r_{t,t+k} = \sum_{i=0}^{k-1} r_{t+i} + \ln(1 - c) - \ln(1 + c).$$

Therefore, the cumulative excess return r for a trading rule over the period from time zero to time T is given by

$$(5) \quad r = \sum_{t=0}^{T-1} z_t r_t + n \ln \left(\frac{1 - c}{1 + c} \right),$$

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

It should be noted that our definition of excess return does not consider the return over a buy-and-hold strategy. We would argue that such a criterion is only appropriate in markets where there is a clearly predictable trend in the asset price, as is the case in the stock market. While it is certainly true that the long-run trend in foreign exchange rates depends on underlying economic fundamentals such as relative inflation rates and levels of productivity, there is no convincing evidence, at least for the currencies we are considering, that these factors are forecastable. In addition, a buy-and-hold strategy is not well defined from the point of view of a global investor. If we consider the dollar/Deutsche mark exchange rate, then the buy-and-hold return for a German investor is the negative of that for a U.S. investor, whereas our measure of excess return is realizable in either currency.

In implementing the genetic program, we follow procedures similar in many respects to those first developed by Allen and Karjalainen (1995) to find trading rules in the stock market. In order to allow the rules to incorporate numerical constants efficiently when the data is non-stationary, we normalize each series by dividing by a 250-day moving average. This means that the data series oscillates around unity and greatly simplifies the task of characterizing the structure of the rules we obtain. The steps involved are detailed below:

1. Create an initial *generation* of 500 random rules.
2. Measure the *fitness* of each rule over the *training period* and rank according to fitness.
3. Select the top-ranked rule and calculate its fitness over the *selection period*. Save it as the initial *best rule*.
4. Randomly select two rules, using weights attaching higher probability to more highly ranked rules. Apply the recombination operator to create a new rule, which then replaces an old rule, chosen using weights attaching higher probability to less highly ranked rules. Repeat this procedure 500 times to create a new generation of rules.
5. Measure the fitness of each rule in the new generation over the training period. Take the best rule in the training period and measure its fitness over the selection period. If it outperforms the previous best rule, save it as the new best rule.
6. Stop if no new best rule appears for 25 generations, or after 50 generations. Otherwise, return to step 4.

The stages above describe one *trial*. Each trial produces one rule. If it produces a negative excess return over the selection period, it is discarded on the grounds that it is dominated by the strategy of never trading and allowing initial margin to accumulate at the overnight interest rate.⁴ If it produces a positive excess return, its performance is then assessed by running it over the *validation period*. We concentrate on this measure of the out-of-sample performance of each rule in the results we report in Section IV.

IV. Results

We obtained trading rules for six exchange rate series.⁵ We use the average of daily U.S. dollar bid and ask quotations for the Deutsche mark, yen, pound sterling, and Swiss franc, obtained from DRI. These rates are collected at 4:00 p.m. local time in London from Natwest Markets and S&P Comstock. Daily overnight interest rates are collected by BIS at 9:00 a.m. London time. All exchange rate data begin on 1/1/74 and end on 10/11/95. All interest rate data begin and end on the same dates, except for Japanese data, for which the start date is 2/1/82. We also created two cross-rate series, Deutsche mark/yen and pound/Swiss franc, assuming the absence of triangular arbitrage opportunities. We will refer to these six series as \$/DM, \$/¥, \$/£, \$/SF, DM/¥, and £/SF. For all exchange rates, we used 1975–1977 as the training period, 1978–1980 as the selection period, and 1981–1995 for validation. The starting date of 1975 for the training period was dictated by the fact that we normalized the series by dividing by a 250-day moving average. Since the interest rate data for the yen were available only from 1982, we measured the excess return of the yen rules in training and selection periods and, for the first year of the validation period, without interest differential. We deliberately chose to standardize training, selection, and validation periods across all currencies to ensure maximum comparability. Training and selection periods were picked to avoid as much as possible the period of dollar appreciation in the first half of the 1980s to minimize the chances of biasing our results.

We distinguish between the level of transaction cost in the training and selection periods and in the validation period. In the latter, we choose a value for the one-way transaction cost, c , which accurately reflects the costs faced by a large institutional trader, in order to assess the economic significance of the profits earned by the trading rules. However, in the former, we simply treat c as a parameter in the search algorithm, and do not impose the restriction that it be realistic. We set c equal to 0.0005 in the training and selection periods, and to 0.0005 in the validation period. We use a higher transaction cost in training and selection periods to bias the search in favor of rules that trade relatively less frequently. If the problem of overfitting in-sample is associated with high trading frequency, which seems intuitively reasonable, using a high transactions cost is one way of guarding

⁴This is equivalent to initializing the search procedure by defining the initial *best rule* under step 3 as the no-trade rule instead of the best rule from the first generation.

⁵This project has been very computer intensive. A back-of-the envelope estimate suggests that it would take at least 81 days of computing time to duplicate the results in this paper on a 120 MHz Pentium. The calculation of the trading rules was the largest element in computation time. Each trial, producing one rule, took approximately two hours.

against this. We report below some evidence in support of these conjectures in the form of an experiment in which c is set to 0.002 in training and selection periods.

Measurement of transactions costs is typically based on an approach proposed by Frenkel and Levich (1977), and uses deviations from exchange rates consistent with the absence of triangular arbitrage. Dooley and Shafer (1983) used a figure of 0.1% for the round-trip transaction cost, and Sweeney (1986) used 0.125%. Consistent with the substantial increase in trading volume and liquidity in foreign exchange markets over the last decade, more recent figures used tend to be somewhat lower. Levich and Thomas (1993) consider a round-trip cost of 0.05% realistic, as do Osler and Chang (1995). Thus, the figure of 0.05% per round trip that we use in our out-of-sample tests can be viewed as reasonable.

In Table 1, we present information on the annual excess return averaged over all 100 rules for each currency.⁶ The picture that emerges is somewhat variable across currencies. For the \$/DM and DM/¥, excess returns are 6.05% and 4.10%, respectively, and for all currencies, they are positive. The average across all currencies is 2.87%. All but four of the \$/DM rules generated positive excess returns. Even the poorest performer in terms of mean return, the £/SF, produced positive excess returns in 89% of cases. This mean return understates the return to a portfolio rule in which each rule receives uniform weight in the portfolio because the mean return overstates transactions costs that would be netted out in the uniform rule. The size of this bias is five to 30 basis points in the annual return.

We also calculated the return to a median portfolio rule, in which a long position was taken if 50 or more of the rules signaled long, and a short position otherwise. That is, a majority vote over all 100 rules determined the position of the trading rule. For \$/¥ and DM/¥, this produced a substantial increase in excess return. Although the performance for two of the currencies deteriorated somewhat, the average excess return over all currencies was significantly improved to 3.67%. We find the performance of the rules surprisingly good, in light of i) the long out-of-sample test period, ii) the well documented difficulty involved in forecasting the exchange rate at short horizons, and iii) the limited information set we provide for the rules. For purposes of comparison, we report in the last row of Table 1 the excess return to a long position in the foreign currency held throughout the validation period.⁷ The mean return to these positions is 0.6%.

There are some striking differences in trading frequency across currencies. The DM/¥ rules trade, on average, about 30 times a year, whereas the £/SF rules trade less than four times a year. In other words, the former sacrifice 1.5% per annum in transactions costs and still achieve a return of 4.1%.

The fact that the exchange rate series are related by triangular arbitrage requires some comment. This does not imply, as it might appear, that some of our exchange rate cases are redundant. A simple example will illustrate this point. Suppose we select the three series \$/DM, DM/¥, and \$/¥, and label the returns r_1 , r_2 , and r_3 , respectively. Then, we have $r_1 + r_2 = r_3$. If rules for the \$/DM and DM/¥ both give the same signal (long or short), then consistency suggests that we

⁶Note that not all rules are necessarily distinct, even though they may appear to have different structures. Classifying two rules as identical if they have the same excess return and the same number of trades, we find that of the 100 \$/DM rules, there are only 87 distinct ones.

⁷For the two cross rates, the yen and SF were treated as the foreign currency.

TABLE 1
 Mean and Median Annual Trading Rule Excess Return for Each Currency over the Period
 1981–1995

	\$/DM	\$/¥	\$/£	\$/SF	DM/¥	£/SF
<i>Mean</i>						
AR*100	6.0485 (28.18)	2.3400 (7.94)	2.2750 (10.55)	1.4154 (5.63)	4.0999 (13.09)	1.0191 (9.48)
MSD*100	3.4897	3.4822	3.6643	3.8794	2.7910	2.9202
% > 0	96	65	85	84	85	89
SR	0.5003	0.1940	0.1792	0.1053	0.4241	0.1007
Trades	106.54	107.98	130.51	156.58	426.61	55.25
% Long	50.5244	78.0106	63.42	81.7283	49.91	93.5729
<i>Median</i>						
AR*100	7.1104	4.5676	1.8466	0.4266	6.5201	1.5633
Trades	35	101	88	14	451	6
<i>Long Rule</i>						
AR*100	1.0342	1.0906	-0.5746	-0.0420	1.4658	0.5257
MSD*100	3.5438	3.5793	3.6062	3.9406	2.9345	2.8910

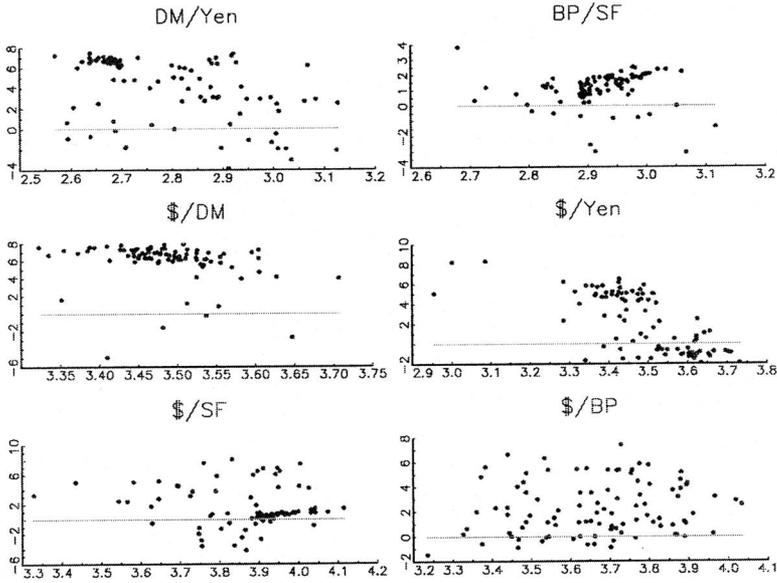
The mean return is calculated over the first 100 rules that produced nonnegative excess returns in the selection period. It is reported as an annual percentage rate in the first row. *t*-statistics are given in parentheses. The third row reports the monthly standard deviation calculated over non-overlapping periods. The fourth row reports the number of rules that generated positive excess returns. The fifth row reports the Sharpe ratio (ratio of mean annual return to standard deviation of annual return). The sixth row reports mean number of trades over all 100 rules. The seventh row relates the mean percentage of long positions taken by the rules. The eighth row shows the annual percentage return to the median portfolio rule. The ninth row provides the number of trades made by the median portfolio rule. The last two rows provide the mean return and monthly standard deviations to a long position in the foreign currency.

see the same signal from the ¥/DM rule. This follows from the fact that we can think of a trading rule as predicting the sign of a future return. But, if the first two rules give different signals, then either a long or a short signal from the ¥/DM rule is consistent and provides additional information about the relative magnitudes of the predicted movements in \$/DM and DM/¥ rates.

In Figure 3, we plot the mean annual excess returns for the individual rules against their monthly standard deviations. There is comparatively little variation in standard deviation across different rules. The explanation for this is that we constrain the rules always to be in the market, either with a long or a short position. Because the return to a short position is the negative of the return to a long position, one would expect the return variances to be similar across rules and close to the variance of the buy-and-hold return. While there is no clear relationship between mean and standard deviation of excess return, there is some tendency for rules that trade very frequently to perform poorly. The effect of transactions costs can only partly explain this. It seems clear that the genetic program occasionally produces a rogue rule that overfits the data, even when the vast majority of rules perform consistently well.

Figure 4 presents a time series of the proportion of all 100 rules giving a long signal over the validation period. This gives a visual representation of the degree of consensus among rules, and of the extent to which their signals are coordinated.

FIGURE 3
 Mean Annual Return vs. Monthly Standard Deviation for All Currencies: 1981–1995



Again, there are sharp differences across currencies. A high proportion of the \$/DM rules identify similar patterns in the data. For most of the time, 90% or more of the rules give the same signal, and they tend to switch from long to short positions and vice versa at much the same time. It is apparent too that the rules identify trends, signaling long positions in the dollar during the period of its rise up to the beginning of 1985, and short positions until the middle of 1988.

Because of the high trading frequency for the DM/¥, it is difficult to see how well coordinated the rules are, but clearly there is a high degree of consensus among them. So, we present the plot over a period of a year (1984) in Figure 5, which reveals that the rules are remarkably well coordinated, with the majority picking switch points at much the same time. But, in the case of the \$/¥, for a substantial amount of time, the proportion hovers around 50%, indicating maximum difference of opinion.

Next, we examine in more detail the performance of a subset of the \$/DM rules. These results are presented in Table 2. We have selected the 10 rules ranked first, eleventh, twenty-first and so on in terms of their performance in the selection period. We find that the rule that did best in the selection period turns out to be one of the poorest performers out-of-sample. It is one of only four rules out of 100 that produced negative excess returns. As noted above, the poor performance is only partly attributable to the transactions costs associated with a much higher frequency of trading. This indicates that the procedure of having separate training and selection periods reduces, but does not eliminate, the problem of overfitting. Rules 2 to 10 all produce excess returns of more than 5% per annum. Trading frequency was quite variable, ranging from 31 to 175. The monthly standard

FIGURE 4
Proportion of Rules Indicating a Long Position for All Currencies: 1981–1995

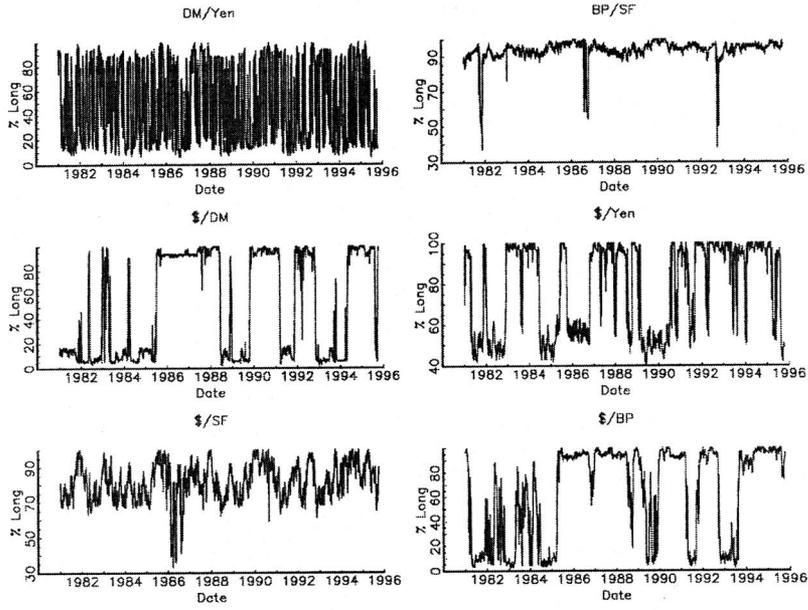
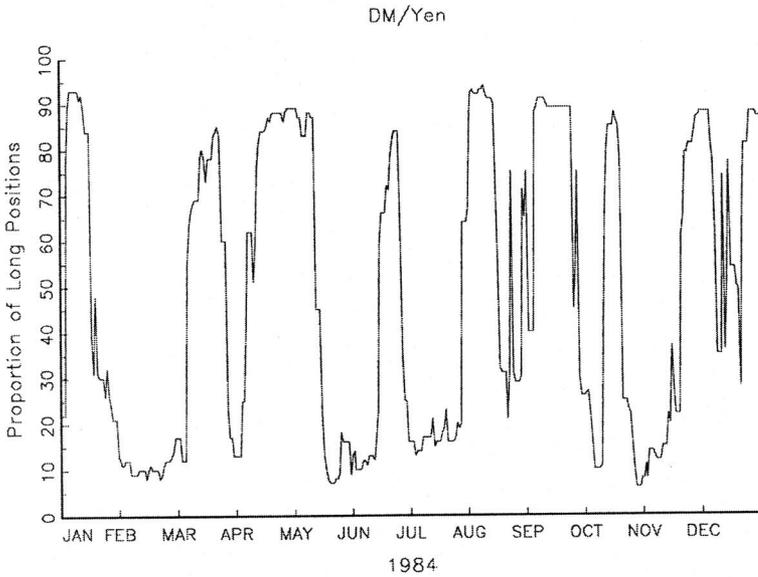


FIGURE 5
Proportion of Rules Indicating a Long Position for DM/¥: 1984



deviations of all rules were rather similar, as was noted above. The rule with the best performance in this subsample, number 7, traded only about twice a year, and the mean number of trades over all 100 rules was seven per year.

TABLE 2
Description of the Performance of 10 Trading Rules for \$/DM from 1981–1995

Rule	Number of Trades	AR	Monthly SD	% Long
1	456.00	-2.7292 (-0.874)	3.6468	58.6
2	119.00	5.7358 (1.904)	3.5559	29.1
3	175.00	6.6309 (2.148)	3.4519	46.1
4	45.00	6.3567 (2.072)	3.4970	50.3
5	77.00	6.1945 (2.005)	3.5375	46.9
6	67.00	5.9612 (1.944)	3.4766	51.0
7	37.00	6.8719 (2.224)	3.5250	47.4
8	37.00	5.9035 (1.916)	3.4943	50.3
9	31.00	6.7446 (2.191)	3.4862	47.7
10	89.00	5.8941 (1.907)	3.5434	50.6
Mean	113.30	5.3564 (5.906)	3.5215	47.8

The rules are ranked according to their performance in the selection period. The first, eleventh, twenty-first, etc., are considered. The third column reports the mean annual percentage excess return, the fourth column the monthly percentage standard deviation, and the fifth column the proportion of the total validation period in which the rule signaled a long position. *t*-statistics are given in parentheses.

We experimented in the case of the \$/DM with a level of transaction cost in training and selection periods four times that used in the validation period ($c = 0.001$) and found that this reduced the average number of trades from 107 to 76 and produced a small improvement in mean excess return of 0.15% per annum. It also increased the number of rules with strictly positive excess returns from 96 to 99. The only rule with negative excess returns earned -0.55% per annum. So, the extremely poor performers were eliminated by this change in procedure, supporting the argument made earlier that the presence of a suitably chosen transaction cost in the training and selection periods introduces a useful friction that provides some protection against overfitting the data.

One question of considerable interest is whether the rules we have identified have a structure that approximates any of those commonly used by technical analysts. This is not always easy to determine, because the rules discovered by the genetic program often have a rather complex nested structure.⁸ However, we

⁸The mean number of nodes for the 100 \$/DM rules was 45.58, and only two rules had fewer than 10 nodes.

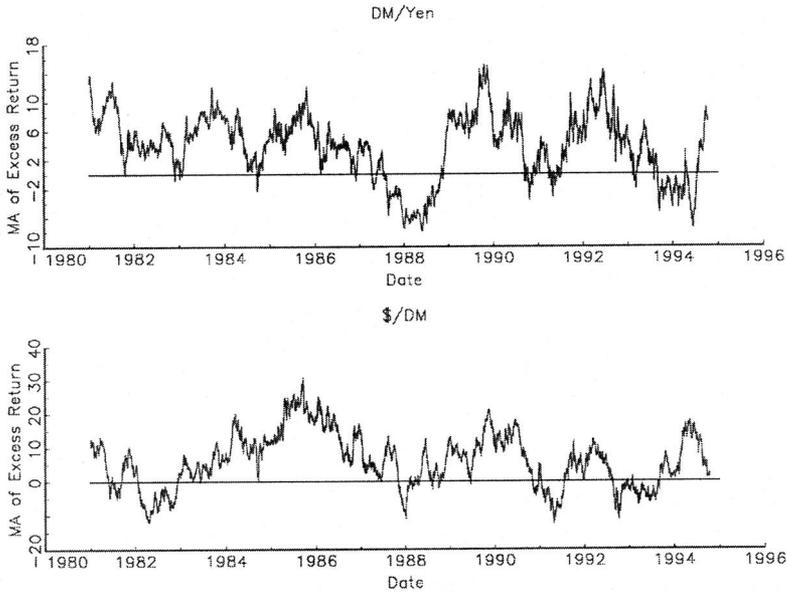
have found that, in a number of cases, rules that appear to be very complicated are highly redundant. The most common form of redundancy arises because a function such as \max that takes an integer argument is followed by a subtree that always evaluates after rounding to the same integer. The most striking illustration is the case of the rule that had the third-best overall performance out-of-sample, with an average excess return of 7.78% and a total of 55 trades. This rule was represented by a tree with 10 levels and 71 nodes, but turned out to be equivalent to the extremely simple rule: "Take a long position at time t if the minimum of the normalized exchange rate over periods $t - 1$ and $t - 2$ is greater than one." Given the normalization we used, the rule is very closely approximated by "Take a long position at time t if the minimum exchange rate over periods $t - 1$ and $t - 2$ is greater than the 250-day moving average." We also analyzed several rules that had a relatively simple structure. One of these, the twenty-fifth best rule over the selection period, had eight nodes and reduced to: "Take a long position at time t if the three-day moving average of the normalized exchange rate is greater than the ratio of the maximum normalized exchange rate over the past six days to yesterday's normalized exchange rate." Another, the fifty-fifth best rule over the selection period, whose daily returns had a correlation of 0.9911 with those of the median rule, prescribed: "Take a long position if the four-day minimum of the normalized exchange rate is greater than one." This is evidently similar to the first rule analyzed above, reflecting the fact that its excess return was 7.34 and its number of trades 37.

Given the high trading frequency of DM/¥ rules, we also attempted to examine several of the simpler rules from that case. We found that at least two of them could be expressed as variants on a 12,250-day moving average rule. Others were much more complex and we were unable to simplify them in any useful way.

In Figure 6, we plot a one-year moving average of the mean annual excess return to the DM/¥ and \$/DM. This provides an informal indication of some of the risk factors associated with trading with an equally-weighted portfolio rule of all 100 individual rules. The moving average for both exchange rates has been positive over a much longer time period than it has been negative. The returns for the \$/DM move between greater extremes: the minimum is around -12% and the maximum over 30% . Comparable figures for the DM/¥ are $(-8\%$ and 14% . Although the recent past has not seen the profits of the 1980s in the \$/DM market, performance since 1990 has still been good. It remains an open question whether rules trained on more recent data will be able to do better than those we have identified.

There is as yet no generally agreed-upon procedure to correct the excess returns earned by a currency trading rule for risk. Existing models do not provide reliable estimates of risk premia associated with individual currencies, and so one is forced to present calculations that, while suggestive, are not firmly grounded in theory. To discover whether the returns to the trading rules could be interpreted as compensation for bearing systematic risk, we calculated betas, presented in Table 3, for the returns to each portfolio of 100 rules over the period 1981–1995. We considered four possible benchmarks: the Morgan Stanley Capital International (MSCI) world equity market index, the S&P 500, the Commerzbank index of German equity, and the Nikkei. Only one beta (\$/¥ on MSCI World Index) was

FIGURE 6
One-Year Moving Average Excess Return for \$/DM and DM/¥



significantly positive, with a value of 0.1689. Interestingly, most of the estimated betas were negative and one (\$/SF on the Commerzbank index) was significantly negative. These figures suggest that the excess returns observed were not compensation for bearing systematic risk. Even if one were to take the extreme position that the associated risk exposure is completely undiversifiable, we find that the Sharpe ratios of the \$/DM and DM/¥ rules are significantly above typical figures for the S&P 500.⁹

Technical analysts persistently claim that the rules they use exploit quite general features of financial markets, and are not specific to any particular market. To investigate this claim, we take the successful \$/DM rules and run them on the data for the other five exchange rates. The results are reported in Table 4. There is a marked improvement in performance over the previous rules in all cases except the DM/¥. The average excess return over the five exchange rates rises from 2.23% to 3.76% and the number of mean returns greater than zero also increases almost uniformly. The most striking change occurs in the \$/SF case, where the excess return rises from 1.41% to 5.40%. One must be somewhat cautious in interpreting these results, because currency movements against the dollar will tend to be positively correlated and, therefore, the reported excess returns are not independent. However, it seems unlikely that all the improvement in performance could be attributable to cross-correlations in the various exchange

⁹The value for the \$/DM is 0.5003, whereas the figure reported for the S&P 500 is generally around 0.3. Osler and Chang (1995) calculate a figure of 0.32, using annual total returns for the S&P 500 from 1973 to 1994.

TABLE 3
 Betas for Returns to Portfolio Rules: 1981–1995

Beta / Standard Error	¥/DM	SF/£	\$/DM	\$/¥	\$/SF	\$/£
World Portfolio Beta (s.e.)	0.0401 0.0448	-0.0599 0.0546	-0.0061 0.0692	0.1689 0.0617	0.0965 0.0608	-0.0029 0.0695
S&P Portfolio Beta (s.e.)	0.0184 0.0304	-0.0480 0.0441	-0.0275 0.0513	-0.0411 0.0398	-0.0406 0.0452	-0.0739 0.0507
Commerzbank Index Beta (s.e.)	0.0125 0.0278	-0.0425 0.0378	-0.0792 0.0428	-0.0624 0.0340	-0.1210 0.0373	-0.0339 0.0433
Nikkei Index Beta (s.e.)	-0.0051 0.0282	-0.0569 0.0382	-0.0679 0.0436	0.0038 0.0348	-0.0133 0.0389	-0.0310 0.0439

The betas are the coefficients from regressing monthly excess returns from a portfolio of 100 rules for each currency on the monthly excess returns for each of the equity indices and no constant. Inference with a constant in the regression was similar. The sample was 1981–1995 for the S&P, Commerzbank, and Nikkei returns and July 1985–1995 for the Morgan Stanley World returns.

TABLE 4
 A Comparison of Results from Running the \$/DM Trading Rules on All Currencies:
 1981–1995

	\$/DM	\$/¥	\$/£	\$/SF	DM/¥	£/SF
AR*100 (1)	6.0485 (28.18)	2.3400 (7.94)	2.2750 (10.55)	1.4154 (5.63)	4.0999 (13.09)	1.0191 (10.55)
(2)	6.0485 (28.18)	4.8726 (25.53)	4.9232 (24.02)	5.4040 (21.83)	2.2822 (19.31)	1.3143 (11.45)
% > 0 (1)	96	65	85	84	85	89
(2)	96	97	93	94	96	89
SR (1)	0.5003	0.1940	0.1792	0.1053	0.4241	0.1007
(2)	0.5003	0.4182	0.3970	0.4040	0.2233	0.1342
Trades (1)	106.54	107.98	130.51	156.58	426.61	55.25
(2)	106.54	103.71	95.60	114.37	116.38	111.69

The rows indexed (1) in the first column reproduce the results of Table 2 for easy comparison. Refer to the notes for that table for explanation of the figures. The rows indexed (2) report the corresponding figures when the 100 \$/DM rules are run on data for other currencies.

rates. We conjecture that one set of rules, trained on data from all four rates against the dollar, would have improved on our baseline case.

Another, perhaps more interesting, result that emerges from this comparison is that the \$/DM rules do substantially *less* well on DM/¥ data than the DM/¥ rules. In other words, there are significant differences between dollar markets and the DM/¥, and the trading rules pick up these differences. This confirms a result that we might have expected from the very large differences in trading frequency. Indeed, we see that the \$/DM rules, when applied to the other currencies, produce rather uniform rates of trade, more than doubling the number of trades for the £/SF and, at the same time, reducing the number of trades for the DM/¥ by 73%.

V. Assessing the Significance of the Results Using Bootstrapping

We carry out a number of bootstrapping simulations to determine whether the rules are simply exploiting known statistical properties of the data. This approach has been used in the context of assessing trading rule performance in the stock market by Brock, Lakonishok, and LeBaron (1992),¹⁰ and by Allen and Karjalainen (1995), and in the foreign exchange market by Levich and Thomas (1993). Bootstrapping permits one to determine whether the observed performance of a trading rule is likely to have been generated under a given model for the data-generating process. The model under consideration is fitted to the data and the residuals are saved and used with the underlying fitted parameters to generate new simulated data sets by resampling with replacement from the distribution of the estimated residuals. The simulated data will possess all the characteristics of the original data captured by the model, but will lose all temporal dependence not captured by the model.

The bootstrapping exercise used three models for the daily log return of the exchange rate: a random walk model, an ARMA model, and an ARMA-GARCH(1,1) model. The order of the ARMA model was chosen by the Akaike information criterion to be (2,2). Four hundred data sets were generated from each of the three models. For each of the 10 rules selected for analysis above, we calculated the cumulative excess return over the simulated data sets according to the measure in equation (5), using observed interest rates. The results of the bootstrapping simulations are presented in Table 5. The pattern that emerges is broadly similar across all models, with p -values generally low except in the case of rule 1. It would appear that the structure captured by the ARMA model can explain about 11% of the excess returns found by the rules. The mean return for the rules run over the simulated ARMA data is 0.57% per annum. The addition of the GARCH term to the ARMA model actually degrades its ability to reproduce the observed returns. Although we report the mean p -value for the 10 selected rules, it cannot be interpreted as the p -value associated with using the 10 rules in combination. An additional set of simulations for the random walk model to examine the performance of the uniform portfolio rule produces a p -value of only 0.0125.

VI. Discussion

Our results present a consistent picture of the potential excess returns that could have been earned in currency markets against the dollar over the period 1981–1995. While there is variation in the size of the profit opportunities, excess returns are significantly positive in all cases. In addition, we show that such profit opportunities existed in the two cross rates ¥/DM and SF/£ and, as far as we are aware, this has not previously been documented. When the \$/DM rules are allowed to determine trades in the other markets, there is a significant improvement in performance in all cases, except for the ¥/DM. The fact that the

¹⁰This paper has a very good discussion of the technique as it has been applied in the trading rule literature.

TABLE 5
 Results of Bootstrapping Using the Random Walk with Drift, the ARMA(2, 2) Model,
 and the GARCH(1, 1)–ARMA(2, 2) Model

	Random Walk		ARMA(2, 2)		GARCH(1, 1)– ARMA(2, 2)	
	AR	% long	AR	% long	AR	% long
1	–1.7202 (0.5975)	66.3	–1.3713 (0.6800)	63.5	–2.0658 (0.6025)	71.8
2	–0.6519 (0.0125)	38.4	–0.6938 (0.0275)	35.6	–0.1815 (0.0125)	41.2
3	–0.9803 (0.0125)	48.1	1.0374 (0.0500)	48.0	–0.6315 (0.0075)	45.9
4	–0.2891 (0.0225)	54.6	1.0277 (0.0575)	53.8	–0.0251 (0.0175)	52.2
5	–0.5890 (0.0200)	49.8	1.2933 (0.0675)	49.7	–0.2198 (0.0125)	46.6
6	–0.3413 (0.0275)	55.3	1.0088 (0.0525)	54.3	–0.1043 (0.0250)	52.9
7	–0.3310 (0.0100)	49.8	0.6571 (0.0275)	49.5	0.0015 (0.0075)	47.5
8	–0.2271 (0.0300)	54.6	0.8898 (0.0625)	53.8	–0.0243 (0.0150)	52.2
9	–0.3033 (0.0125)	49.9	0.7421 (0.0300)	49.6	–0.0780 (0.0050)	47.6
10	–0.4857 (0.0275)	54.6	1.0984 (0.0600)	53.8	–0.2853 (0.0250)	52.2
Mean	–0.5919 (0.0772)	52.1	0.5690 (0.1115)	51.2	–0.3614 (0.0730)	51.0
Portfolio Mean	–0.4677 (0.0125)					

The table reports the mean excess return to each of the 10 rules examined in Table 2, when run on 400 simulated data sets generated by a random walk, ARMA(2, 2), or GARCH(1, 1)–ARMA(2, 2) model. The *p*-values, given in parentheses, report the proportion of simulated excess returns that exceed the value produced when the rule is run on observed data. Columns 3, 5, and 7 give the average percentage of all time periods in which the rule signaled a long position. The bottom row reports mean excess return and *p*-value for a portfolio of all 100 rules, calculated for the random walk model only.

\$/DM rules did substantially less well when run on ¥/DM data than the ¥/DM rules themselves clearly indicates that there are important differences between dollar currency markets and the ¥/DM market, and that these differences have persisted over a considerable period of time.

We find no evidence that the returns to these rules are compensation for bearing systematic risk as measured by betas calculated for various benchmark portfolios (MSCI world index, S&P 500, Nikkei, and Commerzbank equity indices). The bootstrapping results on the \$/DM indicate that the trading rules detect patterns in the data that are not captured by standard statistical models. Our results, therefore, are consistent with and, indeed, strengthen the previous findings on the profitability of technical trading rules in the foreign exchange market.

A number of recent theoretical arguments have demonstrated that using some form of technical analysis might be a rational and, therefore, a profitable strategy (Treynor and Ferguson (1985), Brown and Jennings (1989), Blume, Easley and O'Hara (1994)). However, these studies all hinge on some form of information asymmetry between traders, which we suspect is a less important factor in the foreign exchange market, at least for daily data, than in the stock market.¹¹

A source of information asymmetry, specific to the foreign exchange market, may play a role in generating profitable trading opportunities. Central bank intervention is a relatively common occurrence, and banks have private information about future fundamentals in the form of changes in monetary policy. If the central bank has a targeting objective, this generates a source of speculative profit in the market (Bhattacharya and Weller (1997)). But, this is very unlikely to be the whole story, for several reasons: i) the trading rules for the currencies against the dollar are profitable over the early half of the 1980s, when there was very little intervention by the Federal Reserve; ii) there is evidence from the work of Brock, Lakonishok, and LeBaron (1992) and Allen and Karjalainen (1995) that technical trading is profitable in equity markets, where there is no suggestion that government intervention occurs;¹² and iii) it is implausible that excess returns of the size we report could have been caused by a policy of "leaning against the wind."

Another possibility is that evidence of profitable trading rules signals some form of market inefficiency. We find this plausible for a number of reasons. First, there is other evidence suggesting that the foreign exchange market is inefficient. The forward discount bias is a phenomenon that has received an enormous amount of attention in the literature (see the surveys by Lewis (1994) and Engel (1995)). Thus far, attempts to explain this bias as a consequence of time-varying risk premia have been unsuccessful. Frankel and Froot (1987) have argued that the bias can be accounted for by expectational errors. If this is so, and the errors have some amount of persistence, it suggests that technical analysis may play a role in anticipating the impact of these errors on the market.

We conclude by noting some of the limitations of our analysis, which point to potentially fruitful areas of future research. We have made no attempt to optimize the various features of the genetic programming approach that are under the control of the researcher, such as length of training and selection periods, size of initial generation, choice of function set, limits on the number of nodes or depth of the tree representing a rule, or level of transactions cost in training and selection periods. We have deliberately concentrated here on a restricted information set consisting only of past prices, although we are currently exploring the effect of allowing the rules to make use of information about intervention by the Federal Reserve (Neely and Weller (1997)).

¹¹There is recent evidence from the work of Lyons (1995) that information asymmetry relating to dealer inventory imbalances is an important factor in explaining very high frequency (intraday) data.

¹²This statement needs some minor qualification. There have been occasions when government intervention in the stock market has occurred in Japan in the recent past, and Mark Rubinstein has suggested that there may have been intervention by the Federal Reserve in the immediate aftermath of the stock market crash in 1987 (personal communication). But such occurrences are exceptional.

References

- Allen, F., and R. Karjalainen. "Using Genetic Algorithms to Find Technical Trading Rules." Working Paper, The Wharton School, Univ. of Pennsylvania (Oct. 1995).
- Bhattacharya, U., and P. A. Weller. "The Advantage to Hiding One's Hand: Speculation and Central Bank Intervention in the Foreign Exchange Market." *Journal of Monetary Economics*, 39 (July 1997), 251–277.
- Blume, L.; D. Easley; and M. O'Hara. "Market Statistics and Technical Analysis: The Role of Volume." *Journal of Finance*, 49 (March 1994), 153–181.
- Brock, W.; J. Lakonishok; and B. LeBaron. "Simple Technical Trading Rules and the Stochastic Properties of Stock Returns." *Journal of Finance*, 47 (Dec. 1992), 1731–1764.
- Brown, D. P., and R. H. Jennings. "On Technical Analysis." *The Review of Financial Studies*, 2 (1989), 527–551.
- Diebold, F., and J. Nason. "Nonparametric Exchange Rate Prediction." *Journal of International Economics*, 28 (May 1990), 315–332.
- Dooley, M. P., and J. R. Shafer. "Analysis of Short-Run Exchange Rate Behavior: March 1973 to November 1981." In *Exchange Rate and Trade Instability: Causes, Consequences, and Remedies*, D. Bigman and T. Taya, eds. Cambridge, MA: Ballinger (1983).
- Engel, C. "Why is the Forward Exchange Rate Forecast Biased? A Survey of Recent Evidence." Working Paper No. 95-06, Federal Reserve Bank of Kansas City (Sept. 1995).
- Frankel, J. A., and K. Froot. "Using Survey Data to Test Standard Propositions Regarding Exchange Rate Expectations." *American Economic Review*, 77 (March 1987), 133–153.
- Frankel, J. A., and A. K. Rose. "A Survey of Empirical Research on Nominal Exchange Rates." Working Paper No. 4865, National Bureau of Economic Research (Sept. 1994).
- Frenkel, J. A., and R. M. Levich. "Transaction Costs and Interest Arbitrage: Tranquil versus Turbulent Periods." *Journal of Political Economy*, 85 (Dec. 1977), 1209–1226.
- Holland, J. *Adaptation in Natural and Artificial Systems*. Ann Arbor, MI: Univ. of Michigan Press (1975).
- Kho, B. C. "Time-Varying Risk Premia, Volatility, and Technical Trading Rule Profits: Evidence from Foreign Currency Futures Markets." *Journal of Financial Economics*, 41 (June 1996), 249–290.
- Koza, J. R. *Genetic Programming: On the Programming of Computers by Means of Natural Selection*. Cambridge, MA: MIT Press (1992).
- Levich R., and L. Thomas. "The Significance of Technical Trading Rule Profits in the Foreign Exchange Market: A Bootstrap Approach." *Journal of International Money and Finance*, 12 (Oct. 1993), 451–474.
- Lewis, K. "Puzzles in International Financial Markets." Working Paper No. 4951, National Bureau of Economic Research (Dec. 1994).
- Lyons, R. K. "Is Profit Taking Rational? A Foreign Exchange Dealer's Response." Working Paper, Univ. of California-Berkeley (Dec. 1995).
- Meese, R. A., and A. Rose. "An Empirical Assessment of Non-Linearities in Models of Exchange Rate Determination." *Review of Economic Studies*, 58 (May 1991), 603–619.
- Neely, C., and P. Weller. "Technical Analysis and Central Bank Intervention." Working Paper No. 97-002A, Federal Reserve Bank of St. Louis (Jan. 1997).
- Osler, C. L., and P. H. K. Chang. "Head and Shoulders: Not Just a Flaky Pattern." Staff Papers No. 4, Federal Reserve Bank of New York (Aug. 1995).
- Pring, M. J. *Technical Analysis Explained*, Second Ed. New York, NY: McGraw-Hill (1991).
- Sweeney, R. J. "Beating the Foreign Exchange Market." *Journal of Finance*, 41 (March 1986), 163–182.
- Taylor, M. P., and H. Allen. "The Use of Technical Analysis in the Foreign Exchange Market." *Journal of International Money and Finance*, 11 (June 1992), 304–314.
- Treynor, J. L., and R. Ferguson. "In Defense of Technical Analysis." *Journal of Finance*, 40 (April 1985), 757–773.