ELSEVIER

Journal of International Money and Finance
20 (2001) 949–970

Journal of International Money and Finance

www.elsevier.com/locate/econbase

Technical analysis and central bank intervention

Christopher J. Neely [a,*], Paul A. Weller [b]

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

Abstract

This paper extends genetic programming techniques to show that US foreign exchange intervention information improves technical trading rules' profitability for two of four exchange rates over part of the out-of-sample period. Rules trade contrary to intervention and are unusually profitable on days prior to intervention, indicating that intervention is intended to halt predictable trends. Intervention seems to be more successful in checking such trends in the out-of-sample (1981–98) period than in the in-sample (1975–80) period. Any improvement in performance results from more precise estimation of the relationship between current and past exchange rates, rather than from information about contemporaneous intervention.

Keywords: Technical analysis; Trading rule; Genetic programming; Exchange rate; Central bank; Intervention

1. Introduction

There is now a considerable amount of evidence to suggest that technical trading rules can earn economically significant excess returns in the foreign exchange market (Dooley and Shafer, 1984; Levich and Thomas, 1993; Neely et al., 1997; Neely and Weller, 1999; Sweeney, 1986). However, the reasons for the existence of these excess returns are still not well understood. One possible explanation is that the intervention activities of central banks in the market may account for at least part of the profitability of technical trading rules (Dooley and Shafer, 1984; LeBaron,

* Corresponding author. Tel.: +1-314-444-8568; fax: +1-314-444-8731.
E-mail addresses: neely@stls.frb.org (C.J. Neely); paul-weller@uiowa.edu (P.A. Weller).

PII: S0261-5606(01)00033-X

1999; Szakmary and Mathur, 1997; Neely, 1998). The arguments advanced in favor of this hypothesis focus on the fact that central banks are not profit maximizers, but have other objectives that may make them willing to take losses on their trading. Thus, the stated goal of intervention by the Federal Reserve is to maintain orderly market conditions, and the unstated goals may include the achievement of macroeconomic objectives such as price stability or full employment.[1] If the target for the exchange rate implied by these goals is inconsistent with the market's expectations of future movements in the exchange rate, there may be an opportunity for speculators to profit from the short-run fluctuations introduced (Bhattacharya and Weller, 1997).

LeBaron (1999) investigated the relationship between intervention by the Federal Reserve and returns to a simple moving average trading rule. He used daily intervention data to show that most excess returns were generated on the day before intervention occurred. He found that removing returns on the days prior to US intervention reduced the trading rule excess returns to insignificance.[2] Szakmary and Mathur (1997) examined the link between monthly trading rule returns and monthly changes in the foreign exchange reserves — a proxy for intervention — of five central banks. They also found evidence of an association between intervention activity and trading rule returns.

The fact that trading rule returns were abnormally high on the day *before* intervention tends to support the hypothesis that strong and predictable trends in the foreign exchange market cause intervention, rather than that intervention generates profits for technical traders. But it still leaves open the possibility that a sophisticated technical trader might be able to respond to the fact that intervention had occurred to modify his position and increase his profits. If this is the case, then observing intervention carries additional useful information about the future path of the exchange rate that is not contained in current and past rates.

Although intervention by the Federal Reserve is not publicly announced at the time it occurs, there is evidence that foreign exchange traders quickly become aware of it.[3] Thus we are interested in determining whether knowledge of central bank intervention can increase excess returns to trading rules in dollar exchange rate markets. We investigate this question using the methodology developed in Neely et al. (1997). This allows us to identify optimal ex ante trading rules that use information about whether intervention has occurred, and to compare their profitability to that

[1] The goal of maintaining "orderly market conditions" is stated in the "Foreign Currency Directive", published annually in the Federal Reserve Bulletin with the minutes of the first Federal Open Market Committee meeting of the year.

[2] The timing of the data used by LeBaron (1999) — exchange rates observed at 9:00 am and 11:00 am New York time —left it unclear whether the high returns preceded or were coincident with the high exchange rate returns. Experimentation with data collected before the opening of the New York market makes it clear that the high returns precede the intervention activity. Those results are not reported for brevity.

[3] Klein (1993) finds that for US intervention from 1985 to 1989, 72% of interventions were reported and that 88% of reports were correct. In addition, practitioners with whom we have spoken express confidence that they are aware when the Federal Reserve is intervening.

of rules obtained without the use of such information. We find substantial differences between different time periods, suggesting that either the policies determining intervention or its effects on the market have not been stable over time. We also find some evidence that the use of in-sample intervention data improves the out-of-sample profitability of the trading rules for two currencies, the British pound and Swiss franc, over the period 1981–92. However, we show that this is a consequence of more precise estimation of the relationship between past and future exchange rates. We find no evidence for any currency to suggest that trading profits can be improved out of sample by using rules that condition on contemporaneous intervention information.

2. Methodology

We use genetic programming as a search procedure to identify trading rules that use information both on the past exchange rate series and on intervention activity. We have previously used this technique to find profitable rules that use data on exchange rates alone (Neely et al., 1997) and exchange rates and interest rates (Neely and Weller, 1999). It has also been applied in the equity market (Allen and Karjalainen, 1999). The method is particularly useful for our purposes as it permits flexible incorporation of additional information on central bank intervention into the trading rule.

The genetic program searches for optimal trading rules over a very large population of possible rules using the principles of natural selection. The program creates successive populations of rules according to certain well-defined procedures. Profitable rules are more likely to have their components reproduced in subsequent populations. The basic features of the genetic program are: (a) a means of encoding trading rules so that they can be built up from separate subcomponents; (b) a measure of profitability or “fitness”; (c) an operation which splits and recombines existing rules in order to create new rules.

Before we describe these features, let us first introduce some notation. The exchange rate at date t (USD per unit of foreign currency) is given by S_t. Intervention at date t is given by the indicator variable, I_t, which can take on values 1, 2, or 3, according to whether the US authorities buy dollars, do not intervene, or sell dollars respectively at date t. A trading rule can be thought of as a mapping from past exchange rates and intervention data to a binary variable, z_t, which takes the value +1 for a long position in foreign exchange at time t, and -1 for a short position. Trading rules may be represented as trees, whose nodes consist of various mathematical functions, logical operators and constants, described in Table 1. The functions are distinguished by the data series on which they operate. Thus $\max_S(k)$ is equivalent to $\max(S_{t-1}, S_{t-2}, \ldots, S_{t-k})$, and $\mathrm{lag}_I(k)$ is equal to I_{t-k}.

Fig. 1 presents an example of a simple trading rule that makes use of both exchange rate and intervention data. It signals a long position in foreign currency at date t if the 15-day moving average is greater than the 250-day moving average of the normalized exchange rate, or if the US authorities intervened to buy dollars in the last two days, otherwise a short position.

Table 1
Genetic programming parameters of interest[a]

Size of a generation	500
Termination criterion	50 generations or no improvement for 25 generations
Probability of selection for reproduction with rules ranked from 1 (best) to 500 (worst)	1/(5.7948(1+rank in population))
Arithmetic functions	+, −, ∗, /, norm, constant between (0, 6)
Boolean operators	“if-then”, “and”, “or”, “<”, “>”, “not”, “true”, “false”
Functions of the data	“moving average”, “local maximum”, “local minimum”, “lag of data”, “current data”

[a] In rules trained using intervention data, the functions could be applied to either the normalized exchange rate series or to the intervention indicator series. In rules trained without the intervention data, the functions could only be applied to the normalized exchange rate.

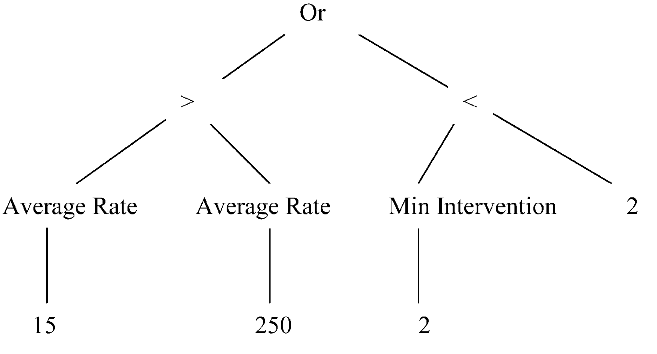

Fig. 1. An example of a trading rule.

The fitness criterion we use in the genetic program is the excess return to a fully margined long or short position in the foreign currency. The continuously compounded (log) excess overnight return is given by $z_t r_t$ where z_t is the indicator variable described above, and r_t is defined as:

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

The domestic (foreign) overnight interest rate is i_t (i_t^*). The cumulative excess return from two round-trip trades[4] (go long at date t, go short at date $t+k$), with round-trip proportional transaction cost c, is

$$r_{t,t+k} = \sum_{i=0}^{k-1} r_{t+i} + \ln(1-c) - \ln(1+c). \tag{2}$$

[4] Each trade incurs a round-trip transaction cost because it involves closing a long (short) position and opening a short (long) one.

Therefore, the cumulative excess return r for a trading rule giving signal z_t at time t over the period from time zero to time T is:

$$r = \sum_{t=0}^{T-1} z_t r_t + \frac{n}{2}\ln\left(\frac{1-c}{1+c}\right). \tag{3}$$

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

To implement the genetic programming procedures we define three separate subsamples, the training, selection and validation periods. The first two periods are equivalent to an in-sample estimation period. The third, the validation period, is used to test the rules trained and selected in the first two periods. The results from this period therefore constitute a true out-of-sample test of the performance of the rules.

The distinct time periods for all currencies were chosen as follows: training period, 1975–77; selection period, 1978–80; validation period, 1981–98. These training and selection periods coincide with those used in Neely et al. (1997) and provide the longest possible out-of-sample period with floating exchange rates.

To examine the stability of the results obtained with this data sample, we repeated the analysis with a more recent subsample of the data: training period, 1987–89; selection period, 1990–92; validation period 1993–98. We chose to begin the second estimation period in 1987 for the following reasons. The period 1981–84 would not be useful in training rules using intervention information because it coincided with a conscious policy decision on the part of the first Reagan administration to avoid intervention.[5] This paucity of intervention is clearly illustrated in Fig. 3. Also, the years 1985–86 coincided with an enormous decline in the value of the dollar; the DEM price of the dollar fell by 39% during those 2 years. Any in-sample period using those years would have introduced a very substantial bias in favor of uninteresting rules that were always long in the foreign currency. Therefore, a second set of rules was constructed using 1987–92 data.

The separate steps involved in implementing the genetic program are described below.

Step 1. Create an initial *generation* of 500 randomly generated rules.

Step 2. Measure the *excess return* of each rule over the *training period* and rank according to excess return.

Step 3. Select the highest ranked rule and calculate its excess return over the *selection period*. If this rule generates a positive excess return, save it as the initial *best rule*. Otherwise, designate the no-trade rule as the initial best rule, with zero excess return.

Step 4. Select two rules at random from the initial generation, 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

[5] There was some — but not much — intervention in 1981–84. See the introduction to Edison (1993).

weights attaching higher probability to less highly-ranked rules. Repeat this procedure 500 times to create a new generation of rules.

Step 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 this best-of-generation rule outperforms the previous best rule, save it as the new best rule.

Step 6. Return to step 4 and repeat until we have produced 50 generations or until no new best rule appears for 25 generations.

The stages above describe one *trial*. Each trial produces one rule whose performance is assessed by running it over the validation period. The validation period for the rules derived over the 1975–80 training/selection period was 1981–98. The validation period for the rules derived over the 1987–92 training/selection period was 1993–98.

Fig. 2 illustrates the splitting and recombination operation referred to in step 4. A pair of rules is selected at random from a 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.[6]

The round-trip transaction cost c was set to 0.0005 (five basis points) in the validation period to reflect accurately the costs to a large institutional trader.[7] In the training and selection periods, however, we treat c as a parameter in the search algorithm and set it equal to 0.001 to bias the search in favor of rules that trade less frequently. We have shown in Neely et al. (1997) that this is an effective way of reducing the chances of overfitting the data.

3. The data

We use the noon (New York time) buying rates for the German mark, yen, pound sterling and Swiss franc (DEM, JPY, GBP, and CHF) from the H.10 Federal Reserve Statistical Release. Daily interest rate data are from the Bank for International Settlements (BIS), collected at 9:00 am GMT (4:00 am, New York time).

As in Neely et al. (1997), we normalize the exchange rate data by dividing by a 250-day moving average. The intervention data we use is the “in market” series from the Federal Reserve Board aggregated across all currencies. The “in market” transactions are explicitly conducted to influence the exchange rate.[8] We construct a variable that can take on one of three values, 1, 2 or 3, depending on whether the US authorities bought dollars, did not transact, or sold dollars on a particular day.

[6] The operation is carried out subject to the requirement that the resulting rule must be well defined. We also impose a restriction that a rule may not exceed a specified size (10 levels and 100 nodes).

[7] Neely et al. (1997) discuss estimates of transaction costs.

[8] Daily US intervention data are released by the Board of Governors of the Federal Reserve with a 1-year lag. Thus 1998 data became available in January 2000.

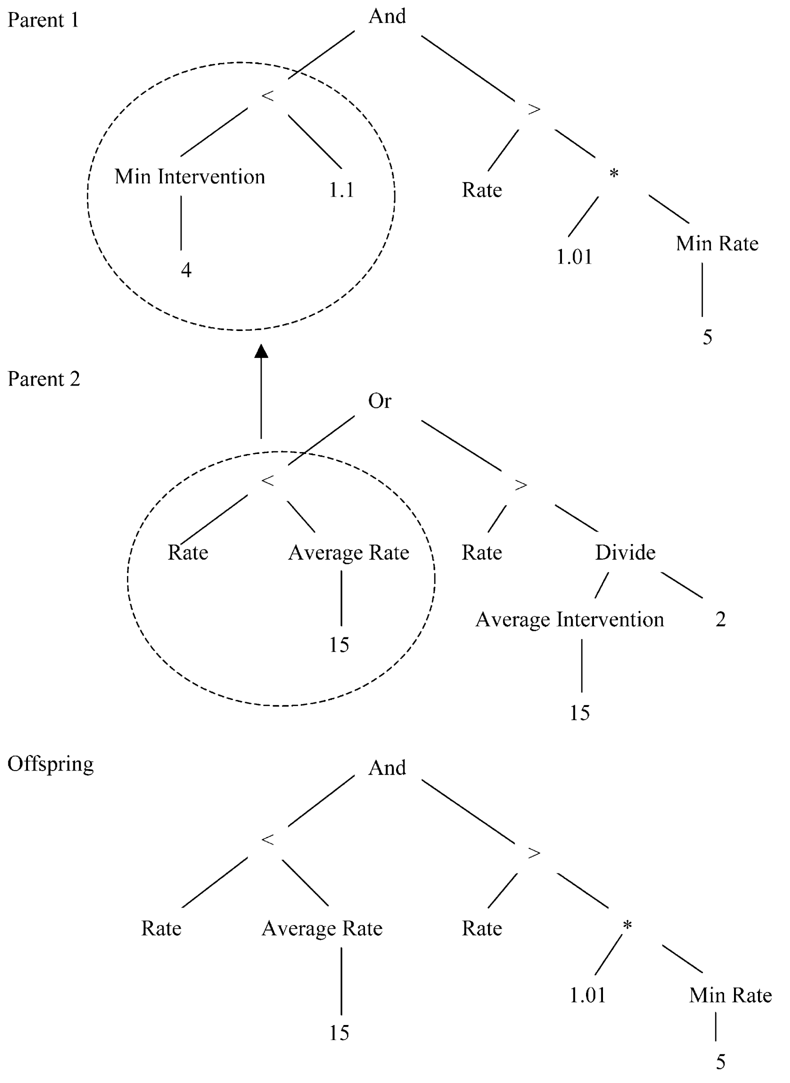

Fig. 2. The recombination operation.

We deliberately avoid the use of quantitative intervention data since this would not have been observable by traders at the time.

Table 2 presents some summary statistics for the various exchange rate returns, including the interest differential but excluding interest accruing over weekends and other missing observations. There is little evidence of significant skewness, and all return series are strongly leptokurtic. However, kurtosis declines in the second half of the sample period, in some cases quite sharply.

Table 3 provides summary statistics on US intervention. We see that the frequency

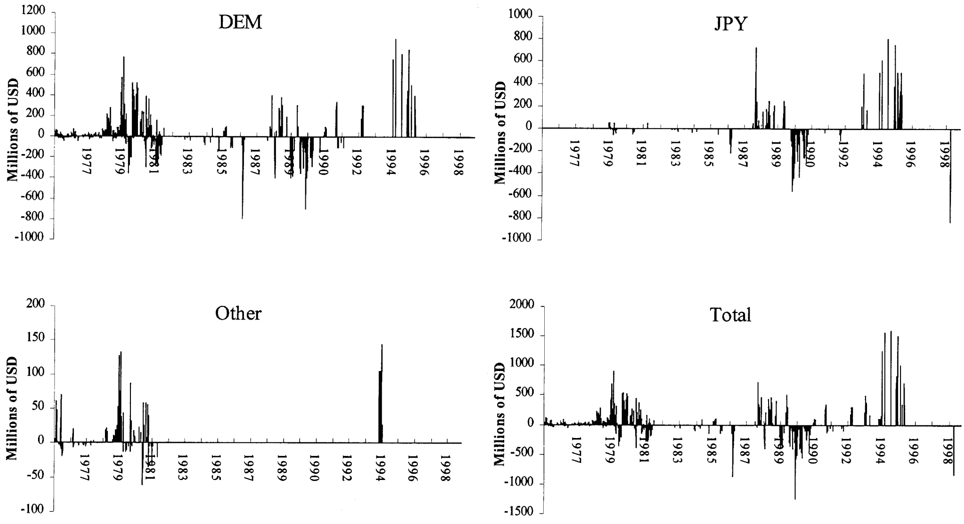

Fig. 3. US official intervention by currency, from 1975 through 1998. The panels of the figure display US official intervention in the DEM, JPY, other currencies and total intervention in millions of US dollars. Purchases (sales) of dollars are displayed as positive (negative) numbers.

Table 2
Summary statistics: daily exchange rate returns including interest differential but excluding weekends and missing observations: 1975–98 and 1987–98[a]

	1975–1998				1987–1998			
	DEM	JPY	GBP	CHF	DEM	JPY	GBP	CHF
Observations	5909	5931	5878	5910	2964	2966	2965	2966
Mean∗100	0.0014	0.0115	0.0005	−0.0002	0.0042	0.0048	0.0111	0.0013
SD∗100	0.66	0.67	0.64	0.76	0.68	0.72	0.64	0.75
Skewness	−0.05	0.56	−0.11	0.01	−0.02	0.52	−0.29	0.13
Kurtosis	3.42	4.85	3.71	3.12	1.77	4.35	2.43	1.60
Min∗100	−5.89	−3.56	−3.86	−5.85	−3.09	−3.37	−3.27	−3.10
Max∗100	4.13	5.62	4.60	4.39	2.90	5.62	2.90	3.89

[a] The kurtosis and skewness statistics are marginally distributed as standard normals under the null hypothesis that the distribution of the exchange rate returns is normal. See Kendall and Stuart (1958) for a derivation of these statistics.

of intervention has declined dramatically over time. Dollar purchases were eight times more frequent during the selection period 1978–80 than the validation period 1981–98. This partly reflects the fact that from 1981 to 1985 there was very little intervention by the USA. There has also been relatively little intervention during the Clinton administration. Fig. 3 illustrates the pattern of intervention over time. We see also that the mean size of intervention has increased substantially over time. The

Table 3
Summary statistics on US intervention data: "in market" series: 1975–98 and 1987–98[a]

	1975–80 rules				1987–92 rules			
	Training 1975–77	Selection 1978–80	Validation 1981–98	Overall 1975–98	Training 1987–89	Selection 1990–92	Validation 1993–98	Overall 1987–98
Observations	735	741	4456	5932	742	742	1483	2967
%>0	18.1	26.5	3.2	7.9	10.0	3.9	1.9	4.4
%<0	18.4	22.7	4.8	8.7	16.8	3.2	0.1	5.1
Mean>0	20.4	136.6	228.4	131.5	180.8	124.3	529.5	242.8
Mean<0	−9.6	−57.8	−172.7	−92.6	−215.9	−126.3	−833.0	−205.6
SD>0	20.4	146.6	297.2	204.6	154.2	96.5	513.0	305.0
SD<0	7.7	77.4	168.5	136.1	175.6	91.2	NA	175.2
Min	−45.5	−379.1	−1250.0	−1250.0	−1250.0	−450.0	−833.0	−1250.0
Max	112.3	904.6	1600.0	1600.0	720.2	336.0	1600.0	1600.0

[a] The data subsamples are: training period 1975:1–1977:12; selection period 1978:1–1980:12; validation period 1981:1–1998:12 for the 1975–80 rules and training period 1987:1–1989:12; selection period 1990:1–1992:12; validation period 1983:1–1998:12 for the 1987–92 rules. Positive intervention corresponds to purchases of dollars in millions. These figures are for the series matched to the USD/JPY exchange rate series. Because each exchange rate series has different missing values, there will be small differences for the intervention series matched to other exchange rate data. %>0 reports the percentage of all trading days on which purchases of dollars were made. Mean>0 and SD>0 give the mean and standard deviation of those purchases. The figures for sales are recorded similarly.

average dollar purchase was 11 times greater in the period 1981–98 than in the period 1975–77.

Table 4 breaks down intervention into different currencies. The DEM was the dominant intervention currency early in the sample period, but the JPY has been used

Table 4
Proportion of intervention in different currencies: 1975–98 and 1987–98[a]

			DEM	JPY	Other
1975–98					
Training	1975	1977	86.17	0.00	13.83
Selection	1978	1980	92.19	1.92	5.89
Validation	1981	1998	53.96	44.95	1.08
Overall	1975	1998	67.90	28.94	3.16
1987–98					
Training	1987	1989	51.98	48.02	0.00
Selection	1990	1992	64.22	35.78	0.00
Validation	1993	1998	45.53	49.53	4.94
Overall	1987	1998	51.77	47.06	1.17

[a] Each column gives the percentage of absolute intervention in different currencies from the "in market" series provided by the Federal Reserve. The first panel reports figures from the subsamples of the 1975–98 period while the second panel reports figures for the 1987–98 period.

almost as often in the 1980s and 1990s. Although there were no JPY interventions at all during the training period 1975–77, the currency accounted for 45% of intervention volume during the validation period 1981–98. Fig. 3 illustrates how the volume of intervention in different currencies has changed over time.

4. Results

4.1. Performance comparisons

Neely et al. (1997) showed that trading rules identified by genetic programming and based only on past observations of the exchange rates earn significant excess returns in the out-of-sample period 1981–95. Here we compare the performance of trading rules trained only on exchange rate data with rules trained on both exchange rate and intervention data. We permit trade conditioned on intervention to occur at 12 noon on the day that it is recorded. Since intervention is generally timed to take place before the London close at 11 am New York time (Goodhart and Hesse, 1993; Humpage, 1998), this makes us confident that we are not allowing the trading rules to use information that would not have been known to the market at the time of trade.[9] Alternatively, the delay in trading on the intervention information means that the rule will miss the immediate response of the price to intervention. This delay is necessary to guard against using future information in trading decisions. For each of the in-sample periods — 1975–80 and 1987–92 — we run 200 trials for each currency, 100 with intervention data and 100 without. This generates a set of 100 rules for each currency under each informational scenario and each in-sample period. We generate sets of rules because the output from a genetic program is inherently stochastic. Although successful rules should detect similar predictive patterns in the data, there is generally some variation in the structure of rules generated from distinct trials. A large sample reduces variation caused by the stochastic nature of the genetic program and produces a more reliable estimate of the average difference in excess return.

We aggregate the individual rules into two portfolio trading rules: the *uniform portfolio rule* and the *median portfolio rule*. The uniform portfolio rule allocates a

[9] In Neely et al. (1997) we used exchange rate data from DRI. In an earlier version of this paper we used DRI data, but later discovered that the time of collection of the data had been incorrectly documented by DRI. In fact, the time at which the data were collected changes in mid-sample. Prior to 8 October 1986 the time of collection was 9:00 am New York time (the New York open), and after that 11:00 am New York time (the London close). Since the vast majority of intervention is timed to occur within the window bracketed by these two times, the information set for traders at date t changes in a crucial way. This is important since, as Peiers (1997) has shown, there are significant information asymmetries around the time of intervention, which in her study of Bundesbank interventions, did not get resolved until shortly before a Reuters report. Not surprisingly, the results with the DRI data set exaggerated the impact of intervention information on trading rule profitability.

fraction 1/100 of the value of the portfolio to each rule.[10] The median portfolio rule generates a long signal at date *t* if 50% or more of the rules give a long signal at date *t*. Otherwise it gives a short signal.

Let us first consider the out-of-sample performance of the portfolio rules generated over the 1975–80 training/selection periods, with and without intervention information. In order to be able to compare the performance of the 1975–80 set of rules with that generated from the 1987–92 in-sample period, we divide the 1981–98 validation period into two subperiods, 1981–92 and 1993–98. The latter subperiod coincides with the validation period for the second (1987–92) set of rules. Because the usual statistical procedures would have little power to discriminate between the portfolio rule returns with and without intervention information, we report Bayesian posterior probabilities to summarize the weight of the evidence in favor of the hypothesis that intervention information increases excess returns. A probability greater than 0.5 favors the hypothesis.

The upper left-hand panel of Table 5 shows the results from the uniform portfolio rule over the 12-year out-of-sample period from 1981 to 1992. We find some evidence that intervention information improved uniform-rule profitability for two currencies. In the case of the GBP the posterior probability is 96.9%. Despite a larger increase in excess return, the evidence in the case of the CHF is weaker. This is a consequence of higher variability in the difference between the two returns. However, the number of profitable rules nearly doubles, lending some additional support to the hypothesis. The figures for the median portfolio rule for the GBP and CHF over the same period provide stronger support for the hypothesis that intervention information improved performance. Excess returns increased from 0.52 to 7.19% (GBP) and from −0.57 to 6.22% (CHF). The associated posterior probabilities were 99.5 and 92.1, respectively.

Over the period 1993–98 the average profitability of the rules declines sharply. The effect of training with intervention information is reversed for the GBP, and is much weaker for the CHF. In the case of the GBP the decline in performance is also reflected in the number of rules which earn a positive excess return over the period. It is important to emphasize the uncertainty associated with our estimates over such a short out-of-sample period. The usual tests — not shown for brevity — would almost always fail to reject the hypothesis that the returns were the same over the 1981–92 period as they were for the 1993–98 period. The only exception is for the GBP with intervention information. However, this decline in performance may be an indication that the forces of competition have reduced profit opportunities as traders have learned about them. When we turn to consider the effect of training with intervention data, the strongest evidence is now for an adverse impact, in the case of the JPY and GBP. To throw further light on this evidence of changing performance over time, we compare these results with those in Table 6 from the second

[10] The excess return to the uniform rule coincides exactly with the average excess return over all 100 rules when transaction costs are zero. Because a simple averaging procedure results in some double counting of transaction costs, the uniform portfolio rule return will always be at least as great as the mean return.

Table 5
Annual portfolio trading rule excess return for each currency over the periods 1981–92 and 1993–98; rules obtained from 1975–80 data using intervention information vs. rules not using intervention information[a]

		1981–92				1993–98			
		DEM	JPY	GBP	CHF	DEM	JPY	GBP	CHF
Panel A: uniform portfolio rule									
AR∗100	CBI	7.12	3.08	5.46	4.64	3.71	3.41	−3.10	1.86
	No CBI	8.29	3.46	3.58	1.72	2.06	7.51	−0.73	−0.67
t-statistic	CBI	2.22	1.37	2.06	1.40	1.03	0.96	1.28	0.46
	No CBI	2.78	1.72	1.27	0.74	0.58	2.20	0.26	0.23
Post prob.		10.00	37.90	96.90	81.70	97.00	1.90	2.20	75.80
Sharpe	CBI	0.63	0.36	0.53	0.39	0.42	0.38	−0.54	0.20
	No CBI	0.77	0.45	0.33	0.22	0.23	0.87	−0.12	−0.10
No. of rules>0	CBI	96	92	96	89	97	52	28	92
	No CBI	94	67	90	45	86	89	57	23
Trades/year	CBI	9.04	4.25	6.58	10.47	7.28	2.92	6.91	6.51
	No CBI	4.15	6.80	7.12	8.04	4.33	7.08	4.89	7.75
% Long	CBI	48.82	81.04	53.55	50.80	37.08	75.18	70.01	50.54
	No CBI	50.10	64.93	61.37	76.28	36.58	58.03	82.85	77.95
Long return		−0.23	2.14	−1.33	−2.07	−0.68	−1.76	2.85	−0.88
Panel B: median portfolio rule									
AR∗100	CBI	7.80	1.77	7.19	6.22	3.17	−1.74	−5.81	1.53
	No CBI	9.73	4.48	0.52	−0.57	1.91	11.87	2.04	−0.65
t-statistic	CBI	2.25	0.57	2.02	1.64	0.81	0.33	1.70	0.33
	No CBI	2.80	1.45	0.14	0.15	0.49	2.36	0.60	0.14
Post prob.		2.80	22.60	99.50	92.10	85.70	1.90	0.00	63.50
Sharpe	CBI	0.64	0.15	0.50	0.46	0.33	−0.13	−0.68	0.15
	No CBI	0.79	0.39	0.04	−0.04	0.19	1.04	0.27	−0.06
Trades/year	CBI	8.49	2.83	3.58	8.49	6.85	0.00	9.85	6.18
	No CBI	2.33	5.00	7.41	4.33	2.84	2.67	0.17	2.84
% Long	CBI	47.69	99.02	46.23	46.68	34.66	100.00	70.05	45.45
	No CBI	48.05	62.21	64.47	98.70	33.20	49.84	89.67	99.09

[a] The rows denoted CBI show results for the rules that use central bank intervention data. The rows denoted No CBI show results for the rules identified only from exchange rate data. AR∗100 is the mean annual percent excess return over the validation period for the 100 rules. The table shows the Newey-West corrected *t*-statistic for the null hypothesis that each portfolio rule has a return equal to zero. Posterior prob. is the Bayesian posterior probability that the excess return of the portfolio using intervention information is greater than that of the rule that does not use such information. The Sharpe ratio is the annual mean excess return divided by the annual standard deviation of the excess return. No. of rules>0 in panel A gives the total number of the 100 rules for each currency which earned a positive excess return over the given period. Trades per year for the uniform portfolio are normalized by the fraction of the portfolio traded. % Long is the percentage of the time the rule was long in the foreign (non-dollar) currency. For the uniform portfolio this represents an average over all individual rules. Long return in panel A is the return to a long position in the foreign currency (buy-and-hold return).

Table 6
Annual portfolio trading rule excess return for each currency over the period 1993–98; rules obtained from 1987–92 data using intervention information vs. rules not using intervention information[a]

		DEM	JPY	GBP	CHF
Panel A: Uniform portfolio					
AR∗100	CBI	−1.73	0.69	−1.84	0.06
	No CBI	−0.84	3.98	−3.80	0.32
t-statistic	CBI	0.75	0.17	1.25	0.02
	No CBI	0.38	1.43	2.56	0.12
Posterior prob.		22.20	12.00	94.00	46.10
Sharpe ratio	CBI	−0.33	0.08	−0.50	0.01
	No CBI	−0.16	0.58	−0.98	0.06
No. of rules>0	CBI	36	39	45	60
	No CBI	42	84	15	52
Trades per year	CBI	39.41	70.22	9.55	8.23
	No CBI	29.60	23.77	24.24	18.66
% Long	CBI	47.12	46.54	68.09	29.94
	No CBI	49.22	45.99	54.59	46.79
Long return		−0.68	−1.76	2.85	−0.88
Panel B: median portfolio					
AR∗100	CBI	−5.47	−0.71	−2.77	−2.46
	No CBI	0.04	2.65	−5.34	0.94
t-statistic	CBI	1.38	0.13	0.82	0.55
	No CBI	0.01	0.48	1.54	0.21
Posterior prob.		5.20	23.10	73.40	24.50
Sharpe ratio	CBI	−0.62	−0.06	−0.35	−0.23
	No CBI	0.00	0.18	−0.62	0.11
Trades per year	CBI	67.96	92.00	16.70	17.87
	No CBI	36.57	40.24	48.76	31.72
% Long	CBI	50.66	45.36	91.59	37.22
	No CBI	48.70	42.57	56.65	48.01

[a] See the notes to Table 5 for a description of each row in the table.

set of rules generated with a training/selection period of 1987–92. Here too, we find a general decline in the profitability of the rules trained without intervention information, compared to the results from 1981–92.[11] Furthermore, there is no evidence that training with intervention information improves performance over this more recent period. We see also that trading frequency is much higher in a number of cases, a clear indication that the structure of the rules identified from this period differs in an important way from that of the earlier set.

[11] Again, we must read the results with some caution because this decline may be due to sampling variation.

4.2. How is the intervention data used?

We conduct two experiments in order to illuminate the way in which the information on intervention influences the performance of the rules. First, to see whether observing intervention during the out-of-sample period contributed to profitability, we compute returns to the rules that are trained with intervention data but are then supplied with a fictitious series out-of-sample indicating that intervention is always zero (null intervention). This is intended to represent the situation in which intervention is not observed. For this exercise, we concentrate on the 1981–92 sample in which the Bayesian posterior probabilities indicate that intervention information improved the performance of GBP and CHF rules. Comparing the null intervention figures for the uniform rule in Table 7 with the results in Panel A of Table 5 we see that performance actually improves for the DEM, GBP and CHF, and is modestly reduced for the JPY. Thus it appears that not observing intervention was, if anything, an advantage for the trading rules during the period 1981–92. This suggests that the intervention variable no longer had the predictive power of the in-sample period. The fact that, nevertheless, out-of-sample performance improved for the GBP and CHF is an indication that during the in-sample period intervention was a significant correlated omitted variable. Its inclusion produced a more precise estimate of the predictive relationship between the exchange rate next period and the past exchange rate series.

One possible reason why the relationship between intervention and the exchange rate might have altered is the change in the nature of intervention that is reported in Tables 3 and 4. In particular, the substantial changes in the volume of intervention in different currencies may have had a significant impact.[12] We report further evidence below to suggest that the response of exchange rates to intervention did change between 1975–80 and later periods.

Next we perform the following simulation experiment. We assume that a simple Markov switching model generates the intervention series independently of the return series. We generate 100 simulated intervention series using the transition probabilities estimated from the validation period and run each set of 100 rules on the observed exchange rate data and the simulated intervention series. This procedure eliminates any predictive power that intervention might have had for future exchange rate returns. In Table 7 we report the performance of each set of rules. The uniform returns are broadly comparable to those in Table 5 but the median returns fall off somewhat. There is also a substantial increase in trading frequency for the DEM and CHF rules. The changes in rule performance suggest that the simulation procedure has eliminated features of the joint distribution of intervention and exchange rate returns that had been incorporated into the trading rules.

The most direct method for determining how the genetic programming rules use the central bank intervention data is to analyze the structure of individual rules.

[12] LeBaron (1999) reports a significant association between trading rule returns and the currency of intervention.

Table 7
Mean annual excess returns over the period 1981–92 for the trading rules run on actual exchange rate data and fictitious intervention data[a]

	DEM	JPY	GBP	CHF
Uniform rule, null intervention				
AR∗100	9.26	2.11	7.08	7.66
t-statistic	2.67	0.68	1.99	2.02
Sharpe ratio	0.74	0.18	0.49	0.57
Trades per year	2.50	0.00	3.75	3.83
% Long	48.28	100.00	46.11	47.96
Uniform rule, Markov intervention				
AR∗100	6.63	2.14	7.03	4.91
t-statistic	1.92	0.69	1.98	1.31
Sharpe ratio	0.54	0.18	0.49	0.37
Trades per year	20.39	0.65	3.81	20.83
% Long	48.39	99.86	46.18	46.50
Median rule, null intervention				
AR∗100	8.10	2.92	5.26	5.76
t-statistic	2.53	1.28	2.03	1.76
Sharpe ratio	0.71	0.33	0.52	0.49
Trades per year	3.05	2.73	6.44	4.75
% Long	49.88	81.10	53.38	52.29
Median rule, Markov intervention				
AR∗100	5.89	2.67	5.32	3.40
t-statistic	1.84	1.16	2.01	1.06
Sharpe ratio	0.52	0.30	0.52	0.29
Trades per year	20.63	6.11	6.88	21.89
% Long	49.52	83.02	53.08	50.64

[a] The panels display portfolio-rule results from the rules generated on 1975–80 data but using either null or simulated intervention data. Panels labeled "null intervention" show the results, comparable to those in Table 5, for the case in which the rules are provided with fictitious intervention data during the period 1981–92 in which all intervention data are set to zero. Panels labeled "Markov intervention" display the mean results from providing the rules with fictitious data generated by drawing 100 sets of intervention data from a calibrated Markov switching process. See the notes to Table 5 for a description of each row in the table.

However, this approach is generally informative only when the structure of the rule to be analyzed is fairly simple. Although such rules may not be representative of the total population, it is interesting to examine an example of a simple rule produced for the CHF. The rule, illustrated in Fig. 4, had a mean annual excess return of 3.72% per annum over the out-of-sample period 1981–98, and a correlation of 98.2% with the median portfolio rule. It provides a clear illustration of the way in which intervention information influences the signal. The rule instructs "Take a long position in foreign currency if the normalized exchange rate is greater than the norm (absolute value of the difference) of the maximum value of the intervention variable

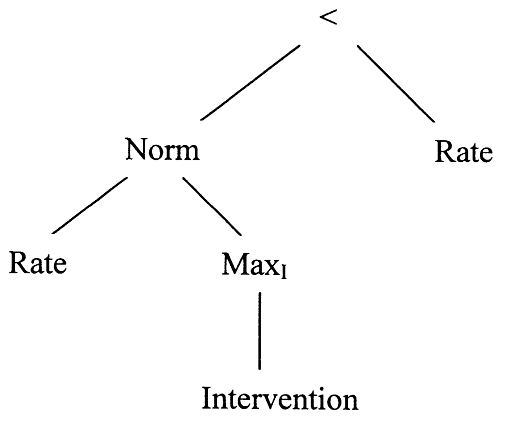

Fig. 4. A trading rule for the CHF found by the genetic program.

(over a time window determined by current intervention) and the normalized exchange rate". The price normalization (division by a 250-day moving average) means that the exchange rate series moves fairly closely around unity. So on a date when the Federal Reserve buys dollars ($I=1$) the rule will always signal a long position in foreign currency, and conversely on a day when it sells dollars ($I=3$) the rule will always signal a short position. Otherwise the rule takes a form that is essentially equivalent to "Take a long position in foreign currency if the current value of the exchange rate is greater than its 250-day moving average". Thus on the day of intervention the trading rule takes a position on the *opposite* side of the market from the Federal Reserve.

4.3. Returns around intervention

We next investigate the behavior of returns around days when intervention took place. Table 8 presents the raw exchange rate returns — returns to a long position in the foreign currency — conditional on intervention at date t. Recall that our exchange rate data are collected at midday New York time, so that the return at $t-1$ is the return to a long position in the foreign currency from midday on the day before intervention to midday on the day of intervention.[13] Since the majority of interventions are timed to occur before the London close (11:00 am New York time), the $t-1$ return will include the immediate response to the intervention (Goodhart and Hesse, 1993; Humpage, 1998).

A strikingly consistent pattern emerges over the period 1975–80. Returns to a long position in the foreign currency were extremely high at $t-1$ when the Federal Reserve bought dollars at t. Similarly, returns to a short position in the foreign cur-

[13] Conditional returns are measured over business days. Thus if intervention occurs on a Monday the return at $t-1$ is measured from midday on the previous Friday to midday on Monday, inclusive of the interest differential.

Table 8
Exchange rate returns conditional on intervention/no intervention by the USA[a]

	DEM			JPY			GBP			CHF		
	$t-1$	t	t+1	$t-1$	t	t+1	$t-1$	t	t+1	$t-1$	t	t+1
1975–80												
AR∗100 (Fed buys USD)	76.2	26.5	5.0	44.7	19.2	−2.9	43.8	22.4	7.8	93.0	32.2	4.2
MSD∗100	4.2	3.4	3.5	6.2	3.1	3.3	4.4	3.6	3.3	5.8	4.8	5.9
AR∗100 (Fed sells USD)	−71.0	−22.8	−17.4	−29.8	0.7	5.0	−23.8	−5.1	−9.5	−75.2	−24.7	−17.0
MSD∗100	4.0	3.0	3.3	5.3	2.9	3.1	2.4	2.6	3.1	4.3	3.4	3.8
AR∗100 (Fed out)	−3.7	−1.8	4.4	4.8	3.9	11.0	−2.5	−1.0	6.2	−9.9	−4.5	3.2
MSD∗100	1.8	2.4	2.8	3.6	3.7	4.1	2.6	2.8	3.1	3.7	4.3	3.3
1981–92												
AR∗100 (Fed buys USD)	67.6	6.7	−15.2	85.4	25.7	−4.5	72.3	26.9	2.5	77.6	5.3	−25.8
MSD∗100	5.3	4.8	3.2	6.7	7.1	4.3	5.5	1.2	1.6	5.3	4.6	2.8
AR∗100 (Fed sells USD)	−63.1	−6.1	6.1	−50.3	3.3	8.0	−59.3	−20.4	0.0	−63.6	−6.4	10.4
MSD∗100	5.8	5.6	4.0	5.9	5.6	4.8	4.8	5.9	4.3	6.1	6.3	5.0
AR∗100 (Fed out)	1.9	−0.1	−0.4	2.8	1.0	1.8	0.2	−1.1	−1.8	−0.5	−2.1	−2.3
MSD∗100	4.2	4.2	4.4	3.7	3.9	4.0	4.8	5.0	5.0	4.7	4.7	4.9
1993–98												
AR∗100 (Fed buys USD)	−36.3	64.2	44.1	−80.2	15.5	72.3	−9.8	18.5	−4.1	−15.0	67.9	60.2
MSD∗100	5.4	4.3	3.9	6.6	4.2	4.0	3.8	3.2	2.4	5.9	5.7	5.2
AR∗100 (Fed sells USD)	144.8	−76.0	98.6	1136.4	9.4	628.9	145.6	129.6	46.1	90.2	−78.4	39.7
MSD∗100	n/a	n/a	n/a	n/a	n/a	n/a	n/a	n/a	n/a	n/a	n/a	n/a
AR∗100 (Fed out)	−0.2	−1.8	−1.7	−1.3	−2.1	−3.6	3.1	2.5	2.5	−0.7	−2.1	−2.1
MSD∗100	3.3	3.0	2.9	4.6	4.7	4.6	2.3	2.3	2.4	4.0	3.7	3.6

[a] Exchange rate returns and standard deviations of returns conditional on intervention for each of the four exchange rates are displayed moving left to right. Each of the three horizontal panels displays results from a different subsample: 1975–80, 1981–92, and 1993–98. Row 1 of each panel gives the annualized percentage return conditional on intervention at date t to buy dollars. The first figure in the column headed $t-1$ gives the return from the rate collected at 12:00 noon New York time on date $t-1$ to the rate collected at the same time on date t. The figures for the succeeding columns are interpreted similarly. Business days only are included so that if date t is a Monday, date $t-1$ is a Friday and the return is calculated from midday Friday to midday Monday. Row 2 reports the monthly standard deviation of the exchange rate return on the date specified. Rows 3–4 and 5–6 report figures conditional on intervention to sell dollars at date t, and on no intervention at date t, respectively.

rency were very high at $t-1$ when the Federal Reserve sold dollars at t.[14] This is consistent with evidence obtained from intraday data — not presented here — that shows that US interventions were preceded by sharp moves in the exchange rate (Neely, 2000). However, we see that returns usually continue to have the same sign on date t, i.e. from midday on the day of intervention to midday on the day after, but they are much reduced. Thus interventions, although they might have checked the appreciation or depreciation of the currency on the day they occurred, on average did not reverse it. Even on the day after intervention in most cases — except for the JPY — we see only a further slowing of the existing trend. Therefore, from 1975 to 1980 it would clearly have been a profitable strategy to trade on the opposite side of the market from the Federal Reserve on a day when it intervened. From 1981 to 1992, intervention at date t is again associated with abnormally high returns from $t-1$ to t, but the pattern of return continuation is less pronounced for the DEM and CHF. Indeed there are usually reversals of trend at $t+1$. Consistent with results found by Humpage (1998) these figures indicate that, at least over a very short horizon, the intervention appears to be successful in checking or reversing a trend. We note that the change in response to intervention between the two periods is coincident with a substantial increase in the size of intervention and a decrease in the frequency of intervention (see Table 2). Both these facts could plausibly be associated with the apparent increase in effectiveness.

We include the information on returns for 1993–98 for completeness, although there were only 29 interventions during this period — including only one dollar sale. In a complete reversal of the earlier pattern, the dollar appreciated on average from $t-1$ to t when the Fed bought dollars. Given the timing of our observations, which means that the $t-1$ return includes the immediate response to intervention, a possible explanation for these results is that intervention had a much larger immediate impact on the exchange rate, but that it was very short-lived.

In Table 9 we show the median portfolio returns — from the 1975–80 rules, using intervention information — for the DEM, GBP and CHF over the periods 1976–80, 1981–92, and 1993–98.[15] The results for the JPY are not informative because the median portfolio rule, trained with intervention information, performed poorly out of sample and took long positions over 99% of the time. We therefore omit them. If we consider the in-sample results for the DEM (top-left panel of Table 9) we see that the median portfolio rule took a long position in foreign currency at $t-1$ in 89% of cases when support intervention occurred the following day. This is strong evidence in favor of the hypothesis that a persistent depreciating trend in the dollar tended to precede intervention. Similar results, although not quite as pronounced, are found for dollar sales. It is also clear that the rule has detected the profitability

[14] These results are consistent with LeBaron (1999), who found that removing returns on days prior to non-zero intervention reduced the profitability of a simple moving average trading rule to insignificance. Note that LeBaron's study used data from DRI and that the time at which these data were collected changed in mid-sample (see footnote 8).

[15] The first sample starts in 1976 because the year 1975 was used as a window for lagged variables by the rules.

Table 9
Median portfolio returns and positions conditional on intervention by the Federal Reserve[a]

1975–80 rules	DEM			GBP			CHF		
	$t-1$	t	t+1	$t-1$	t	t+1	$t-1$	t	t+1
1976–80 data (in-sample)									
AR∗100 (buy)	41.5	29.9	−0.6	52.2	32.2	14.4	45.2	34.0	3.2
MSD∗100	4.6	3.3	3.8	3.8	3.0	3.3	6.1	4.5	5.5
% Long	88.6	100.0	87.9	88.2	88.2	88.2	82.1	100.0	85.4
AR∗100 (sell)	22.1	16.9	16.3	−9.6	4.3	4.4	16.7	20.0	22.6
MSD∗100	2.9	2.9	3.0	2.7	3.0	2.5	3.6	3.3	3.3
% Long	36.3	0.0	38.1	81.9	81.9	81.4	25.2	0.0	8.0
AR∗100 (out)	−7.8	−1.5	10.2	−0.5	2.7	9.4	−13.8	−10.2	0.6
MSD∗100	1.8	2.7	2.4	2.4	2.8	2.8	3.0	4.0	3.6
% Long	72.1	78.7	71.8	67.7	67.7	68.0	67.7	68.5	71.7
1981–92 data (out-of-sample)									
AR∗100 (buy)	46.9	5.1	−15.0	84.7	31.3	−7.0	54.2	3.5	−21.9
MSD∗100	4.5	3.7	4.0	3.6	3.5	3.5	5.1	4.2	4.3
% Long	87.7	100.0	88.6	92.1	92.1	92.1	86.0	100.0	86.8
AR∗100 (sell)	61.8	4.7	2.7	22.1	−5.0	−4.7	64.0	5.5	−11.6
MSD∗100	4.6	4.1	4.1	4.9	4.2	4.0	5.0	4.5	4.7
% Long	10.7	0.0	10.7	38.1	37.2	36.3	7.0	0.0	0.5
AR∗100 (out)	1.8	8.2	9.4	2.7	7.2	8.7	−0.5	6.4	9.0
MSD∗100	3.9	4.0	4.0	3.9	4.0	4.1	4.3	4.4	4.3
% Long	48.9	49.3	48.9	44.8	44.9	44.9	48.2	48.1	48.6
1993–98 data (out-of-sample)									
AR∗100 (buy)	−25.2	58.3	29.4	−12.7	9.8	−5.3	−59.8	65.7	54.3
MSD∗100	5.4	4.3	3.9	3.7	3.2	2.4	5.8	5.7	5.2
% Long	53.6	100.0	50.0	64.3	64.3	60.7	82.1	100.0	82.1
AR∗100 (sell)	−144.8	76.0	−98.6	133.0	129.6	46.1	−90.2	78.4	−39.7
MSD∗100									
% Long	0.0	0.0	0.0	100.0	100.0	100.0	0.0	0.0	0.0
AR∗100 (out)	3.7	2.1	2.9	−5.8	−6.2	−5.4	2.7	0.3	0.7
MSD∗100	3.2	3.3	3.3	2.9	2.9	2.9	3.8	3.8	3.8
% Long	34.2	33.4	34.4	69.9	69.9	70.0	44.4	44.1	44.5

[a] Median rule returns and their monthly standard deviations conditional on intervention are displayed. JPY results are omitted because the median rule took a long position over 99% of the time. The top panel displays results from 1976–80 — rather than 1975–80 as in Table 8 — because 1975 was used as a data window for lags in constructing the trading rules. The rules were obtained from training and selection periods 1975–80. Row headings “buy”, “sell” and “out” give figures conditional on the Fed buying USD, selling USD, and not transacting respectively. See the notes to Table 8 for additional details.

of trading on the opposite side of the market from the Fed at t, since in every instance when the Fed bought dollars, the rule took a long position. This trading pattern is reproduced in the 1981–92 period, but its profitability is much reduced. The results for the CHF are very similar. For the GBP, although the rule is predominantly on the right side of the market for dollar purchases, it does less well during both the in-sample and out-of-sample periods around sales.

5. Discussion and conclusion

The profitability of a trading rule is closely related to the predictability of the exchange rate one period ahead. However, it is important to recognize the differences between this investigation and one that uses standard statistical procedures to address the issue of predictability. The application of Granger causality tests to the data (not reported) provides strong evidence for all currencies except the JPY that returns and squared returns help predict intervention and also lends support to the hypothesis that intervention causes returns. These conclusions are based on the results from running two-variable vector autoregressions, including past exchange rates and magnitude of intervention. However, this evidence of Granger causality does not necessarily imply that a trading strategy that conditions on intervention will be more profitable than one that does not. There are several reasons for this. First, the linear predictive power attributed to intervention by the Granger causality tests may also be present as a non-linear component in the past exchange rate return series. The genetic program may already have incorporated the information into the trading rules trained only on exchange rate data. Second, the Granger causality tests use the magnitude of intervention, while we provide the genetic program only with information about the sign of intervention. Third, the linear relationship may not be economically significant; transactions costs may eliminate any excess returns.

The volatility associated with exchange rate returns cautions us to be circumspect in drawing conclusions about the value of intervention information as an input to trading rules. Given this caveat, however, the weight of the evidence suggests that training with intervention information improved the performance of the GBP and CHF rules over the period 1981–92, although there is no such evidence for the major intervention currencies, DEM and JPY. For both DEM and JPY, providing intervention information led to some deterioration in performance for the median portfolio rule during the out-of-sample period 1981–92. This may be explained by the fact that intervention policy changed in some significant ways between in-sample and out-of-sample periods. For example, the DEM was by far the most used intervention currency in the 1975–80 period (Table 4) but the JPY was used nearly as often in the 1981–98 period. In addition, interventions were much more frequent, but much smaller in the former period.

Experiments with null and simulated intervention (Table 7) show that the improved performance of the GBP and CHF rules comes about not because the intervention signal itself has predictive power out of sample, but because training rules with intervention data better identify the predictive component in the past exchange rate series. This suggests that the predictive relationship between past and future exchange rates has been more stable than the relationship between intervention and the exchange rate.

Because we find no evidence for any currency that contemporaneous information about the occurrence of intervention improves trading rule performance, our findings do not support the view that intervention activity is a source of profit for technical traders in the foreign exchange market. On the contrary, our results indicate that the profitability of technical trading rules is a consequence of strong and persistent trends in exchange rates, which intervention is intended to reverse.

Acknowledgements

The authors would like to thank Robert Dittmar for excellent programming assistance, Kent Koch for excellent research assistance, Michael Melvin, Owen Humpage and two referees for helpful comments and their colleagues at the Federal Reserve Bank of St Louis for generously sharing their computers at night and over weekends. Paul Weller would like to thank the Research Department of the Federal Reserve Bank of St Louis for its hospitality while he was a Visiting Scholar, when this work was initiated. The views expressed are those of the authors and do not necessarily reflect official positions of the Federal Reserve Bank of St Louis or the Federal Reserve System.

References

Allen, F., Karjalainen, R., 1999. Using genetic algorithms to find technical trading rules. Journal of Financial Economics 51, 245–271.

Bhattacharya, U., Weller, P.A., 1997. The advantage to hiding one’s hand: speculation and central bank intervention in the foreign exchange market. Journal of Monetary Economics 39, 251–277.

Dooley, M.P., Shafer, J.R., 1984. In: Bigman, D., Taya, T. (Eds.), Floating Exchange Rates and the State of World Trade Payments. Ballinger Press, Cambridge, MA, pp. 43–69.

Edison, H.J., 1993. The effectiveness of central-bank intervention: a survey of the literature after 1982. Special Papers in International Economics No. 18, Department of Economics, Princeton University.

Goodhart, C.A.E., Hesse, T., 1993. Central bank intervention assessed in continuous time. Journal of International Money and Finance 12, 368–389.

Humpage, O., 1998. U.S. Intervention: assessing the probability of success. Federal Reserve Bank of Cleveland Working Paper 9608.

Kendall, M.G., Stuart, A., 1958. The Advanced Theory of Statistics. Hafner Press, New York.

Klein, M.W., 1993. The accuracy of reports of foreign exchange intervention. Journal of International Money and Finance 12, 644–653.

LeBaron, B., 1999. Technical trading rule profitability and foreign exchange intervention. Journal of International Economics 49, 125–143.

Levich, R., Thomas, L., 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.

Neely, C.J., 1998. Technical analysis and the profitability of U.S. foreign exchange intervention. Federal Reserve Bank of St Louis Review 80, 3–17.

Neely, C.J., 2000. The temporal pattern of trading rule returns and central bank intervention: Intervention does not generate technical trading profits. Federal Reserve Bank of St Louis Working Paper 00-18B. Forthcoming in the Journal of International Economics.

Neely, C.J., Weller, P.A., 1999. Technical trading rules in the European Monetary System. Journal of International Money and Finance 18, 429–458.

Neely, C.J., Weller, P.A., Dittmar, R., 1997. Is technical analysis in the foreign exchange market profitable? A genetic programming approach. Journal of Financial and Quantitative Analysis 32, 405–426.

Peiers, B., 1997. Informed traders, intervention and price leadership: a deeper view of the microstructure of the foreign exchange market. Journal of Finance 52, 1589–1614.

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

Szakmary, A.C., Mathur, I., 1997. Central bank intervention and trading rule profits in foreign exchange markets. Journal of International Money and Finance 16, 513–535.