FORECASTING FOREIGN EXCHANGE RATES USING RECURRENT NEURAL NETWORKS

PAOLO TENTI

Tenti Financial Management s.a., Lugano, Switzerland

This article proposes the use of recurrent neural networks in order to forecast foreign exchange rates. Artificial neural networks have proven to be efficient and profitable in forecasting financial time series. In particular, recurrent networks, in which activity patterns pass through the network more than once before they generate an output pattern, can learn extremely complex temporal sequences. Three recurrent architectures are compared in terms of prediction accuracy of futures forecast for Deutsche mark currency. A trading strategy is then devised and optimized. The profitability of the trading strategy, taking into account transaction costs, is shown for the different architectures. The methods described here, which have obtained promising results in real-time trading, are applicable to other markets.

For years, opposing views have existed between the trading and academic communities about the statistical properties of foreign exchange rates. Traders considered exchange rates to have persistent trends that permitted mechanical trading systems (systematic methods for repeatedly buying and selling on the basis of past prices and technical indicators) to consistently generate profits with relatively low risk. Researchers, on the other hand, presented evidence supporting the *random walkhypothesis*, which implies that exchange rate changes are independent and have identical statistical distributions. When prices follow a random walk, the only relevant information in the historical series of prices, for traders, is the most recent price. The presence of a random walk in a currency market is a sufficient, but not necessary, condition to the existence of a weak form of the *efficient market hypothesis*, i.e., that past movements in exchange rates could not be used to foretell future movements.

While there is no final word on the diatribe between practitioners and academicians about the efficiency of currency markets, the prevalent view in economic literature that exchange rates follow a random walk has been dismissed by recent empirical work. There is now strong evidence that exchange rate returns are not independent of past changes. Before the advent of nonlinear dynamics, statistical tests for the random walk were usually conducted by verifying that there was no linear dependence, or that autocorrelation coefficients were not statistically different from zero. However, the lack of linear dependence did not rule out nonlinear

Address correspondence to Paolo Tenti, Tenti Financial Management s.a., Via Nassa 7 - 6900 Lugano, Switzerland. dependence, the presence of which would negate the random walk hypothesis. Therefore many tests were often inappropriate, and some conclusions were questionable. Recent evidence has clearly shown that while there is little linear dependence, the null hypothesis of independence can be strongly rejected, demonstrating the existence of nonlinearities in exchange rates (Brock et al., 1991; De Grauwe et al., 1993; Fang et al., 1994; Taylor, 1986).

It would seem to be very difficult to predict exchange rates, characterized by nonlinearities and high noise, using only high-frequency (weekly, daily, or even intraday) past prices. Surprisingly though, there are anomalies in the behavior of the foreign exchange markets that cannot be explained under the existing paradigm of market efficiency. Sweeney (1986) applied the academic filter rule to several spot exchange rate series from 1973 to 1980, with successful results for several filter sizes. Lukac et al. (1988) simulated trading of 12 technical systems for British pound and Deutsche mark currency futures for the period 1978-1984, obtaining significant risk-adjusted returns for 4 of them. These early tests of mechanical trading systems often received unflattering criticism, presumably because they used less than rigorous methodology. While it is true that mechanical trading systems need to be designed and optimized with care in order to avoid the risk of overfitting (Pardo, 1992), new evidence has emerged, which reinforces previous tests, on the profitability and statistical significance of mechanical trading systems in currency markets (Bilson, 1992; LeBaron, 1993; Levich & Thomas, 1993; Taylor, 1994). It seems that technical trading rules are able to pick up some of the hidden patterns in the inherently nonlinear price series, contradicting the conclusions reached by many earlier studies that found technical analysis to be useless.

Some technical indicators were used as inputs to three recurrent neural network architectures, in order to forecast foreign exchange rates. The importance of the trading strategy (when to enter, when to exit, number of contracts per trade, etc.) can hardly be underestimated. Research shows that identifying the appropriate trading strategy for each forecasting problem is vital to each system's trading performance. It is also important to emphasize that prediction accuracy is not the goal in itself, and it should not be used as the guiding selection criteria in the tests. (While this simple concept is part of the wealth of knowledge of mechanical traders (Schwager, 1984), it is rarely considered in tests undertaken by academicians.) The three recurrent architectures are compared in terms of prediction accuracy of futures forecast of Deutsche mark currency and its consequential profitability and riskiness of the trading strategy. Building a trading system and testing it on the evaluation criteria that are pertinent to the application is the only practical and relevant approach to evaluate the forecasting method.

RECURRENT NEURAL NETWORK AS A FORECASTING TOOL

The potential advantages and limitations of an artificial neural network (ANN), and, in particular, of a multilayer feedforward neural network (Werbos, 1974;

Rumelhart et al., 1986) over other statistical methods or expert systems are well known. They are universal function approximators and, being inherently nonlinear, are notoriously good at detecting nonlinearities, but suffer from long training time and a very high number of alternatives as far as architectures and parameters go. They are also prone to overfitting data. Another common critique that is made about ANNs is that they are "black boxes": knowledge of the value of the weights and biases in the network gives, at best, only a rough idea of the functional relationships. Thus, even if ANNs are based on causally related data, the resulting model may not give a great amount of insight into the strength and nature of the relationships within it. This elusiveness of ANNs is the price to be paid in return for their being model-free estimators.

Recurrent neural networks (RNNs), in which the input layer's activity patterns pass through the network more than once before generating a new output pattern, can learn extremely complex temporal patterns. Several researchers have confirmed the superiority of RNNs over feedforward networks when performing nonlineartime series prediction (Connor & Atlas, 1993; Logar et al., 1993; Adam et al., 1994). (See also (Kamijo & Tanigawa, 1990) for an application of RNNs to recognition of stock price patterns.) The main disadvantage of RNNs is that they require substantially more connections, and more memory in simulation, than standard backpropagation networks. RNNs can yield good results because of the rough repetition of similar patterns present in exchange rate time series. These regular but subtle sequences can provide beneficial forecastability.

Recurrent architecture proves to be superior to the windowing technique of overlapping snapshots of data, which is used with standard backpropagation. In fact, by introducing time-lagged model components, RNNs may respond to the same input pattern in a different way at different times, depending on the sequence of inputs. Prediction using an RNN involves the construction of two separate components: one or more recurrent layers that provide the temporal context, usually referred to as short-term memory, and a predictor, usually the feedforward part of the network. The short-term memory retains features of the input series relevant to the prediction task and captures the network's prior activation history. Therefore the appropriate response at a particular point in time could depend not only on the current input, but potentially on all previous inputs.

The tests were performed with three variations of RNNs. They belong to the RNN family known as local feedback networks, where only local connections are activated. The rationale is that instead of learning with complex, fully connected recurrent architectures, redundant connections should be eliminated in order to significantly increase the network's generalization capability. The first architecture used (RNN1) is similar to that developed by Jordan (1986), known as sequential network and used to solve some sequential tasks in cognition. The network has one hidden and one recurrent layer. The output layer is fed back into the hidden layer, by means of the recurrent layer, showing resulting outputs of previous patterns (Figure 1; self-loops of recurrent neurodes are not shown in this and subsequent



Figure 1. Recurrent backpropagation with output layer feedback link (memory: outputexponential).

figures). The recurrent neurode allows the network's hidden neurodes to see their own previous output, so that their subsequent behavior can be shaped by previous responses. The recurrent layer is what gives the network its memory. It follows the taxonomy proposed by Mozer (1993), which distinguishes between the short-term memory's content and form. The version I used was characterized by output-exponential memory.

With respect to the form of the memory, the use of an exponential trace memory acts on the series of inputs $x(1), \ldots, x(t)$, creating a state representation [$\overline{x_1}(t), \overline{x_2}(t), \ldots, \overline{x_t}(t)$], where each $\overline{x_t}(t)$ is related to the input sequence by the function e_i :

$$\overline{X}_i(t) = \sum_{\tau=1}^{\infty} e_i(t-\tau) X(\tau)$$

where $e_i(t) = (1 - \mu_i)\mu_i^t$ with $0 < \mu_i < 1$.

An important property of exponential trace memories is that $\overline{x}_i(t)$ can be calculated incrementally:

$$\overline{x}_i(t) = (1 - \mu_i)x_i(t) + \mu_i \overline{x}_i(t - 1)$$

These memories can then be seen as exponentially weighted moving averages of past inputs. The exponential memory, used also for the other two versions, makes the strength of more distant inputs decay exponentially. The rate of decay is governed by μ^{j} .

In the second version (RRN2, shown in Figure 2), similar to that of Frasconi et al. (1992), the hidden layer is fed back into itself through an extra layer of recurrent neurodes. Both the input layer and recurrent layer feed forward to activate the hidden layer, which then feeds forward to activate the output layer. Therefore the features



Figure 2. Recurrent backpropagation with hidden layer feedback link (memory: transformed input-exponential).

detected in all previous patterns are fed back into the network with each new pattern. These recurrent neurodes remember the previous internal state.

In the third version (RNN3, shown in Figure 3), patterns are processed from the input layer through a recurrent layer of neurodes, which holds the input layer's contents as they existed when previous patterns were trained, and then are fed back into the input layer.

The memory's content is the dimension that differentiates the three versions of RRN. It refers to the fact that although it must hold information about the input sequence, it does not have to be a memory of the raw input series. In the three versions used here there were one-for-one linear connections between each recurrent neurode and, respectively, each output, hidden, or input neurode.

Issues such as learning parameters, number of hidden neurodes, and activation functions are also important in determining the chances of success of different configurations of RNNs. Several alternatives regarding the parameters were tested. The best results were provided, contrary to the tests performed by Refenes et al.



Figure 3. Recurrent backpropagation with input layer feedback link (memory: input-exponential).

(1993) with standard backpropagation, by the symmetric sigmoid logistic activation function:

$$f(x) = \frac{e^{x} - e^{-x}}{e^{x} + e^{-x}}$$

where f(x) has the same shape as the standard sigmoid function, except that its range is [-1, 1] rather than [0, 1]. The learning rate was initially set at 0.05 and decreased gradually to 0.0005 during the first 500 passes, while the momentum term was fixed at 0.1. The rate of decay μ_i was set at 0.6. Choosing the appropriate number of hidden neurodes was extremely important. The configuration that obtained the best results in terms of generalization had 18 inputs, 5 hidden neurodes, and 1 output.

EMPIRICAL DESIGN OF RNN

The experiments, in order to be useful and applicable to real-time trading, must create conditions that are as close as possible to reality. Therefore one must take into account, when using spot exchange rates, the interest rate differential among the currencies in question. Unfortunately, this crucial criterion is often overlooked. There will always be a difference between "paper" profits and real profits: the objective is to minimize that difference. By using forward rates or currency futures, it is possible to overcome this problem because they already include a premium or a discount due to the differences in interest rates. Any trading system based solely on spot exchange rates is just an approximation because it falls short of dealing with the problem of interest rate differentials. (It would be possible to use spot exchange rates by taking into account overnight interest rates on spot interbank deposits for returns calculations, but at the price of more approximations. In addition, the spread between bids and offers will tend to be greater for forward than for spot rates, increasing as the maturities grow longer.) Currency futures are traded primarily on the International Monetary Market of the Chicago Mercantile Exchange. They offer the advantage of being standardized by size, highly liquid, and for our purposes, providing a reliable source of prices. In the futures market, a speculator would buy a currency futures if he expects the spot rate at maturity to exceed the current futures price, and he would sell currency futures if he expected the spot rate at maturity to be less than the futures price. The set of data used in the experiments consisted of opening and closing daily prices of Deutsche mark currency futures from January 1990 to December 1994.

Choosing the appropriate price series for currency futures is hardly a trivial matter and is the first step in building a trading system. Currency futures contracts are traded for delivery at a fixed maturity, namely, the third Wednesday of March, June, September, and December. Using individual contracts complicates the task of

training and testing the system. The training and testing usually require a price data history that is much longer than the typical liquid trading period for an individual contract. Furthermore, the simultaneous use of individual contracts is difficult because it is necessary to combine a large number of individual results for each contract as well as to deal with possible divergence of trading signals when switching from the expiring contract to the next one. The commonly proposed solution is to create a single continuous price series by using the nearest futures prices, with a jump to the prices of the successive contract made at the beginning of the month or at a specified number of trading days before expiration. The fatal distortion of this system is that there could be significant price gaps created in the series at the rollover dates, between the expiring and the subsequent contracts. The nearest futures series will create illusory price moves at the transition points, distorting both training and testing activities. In addition, the nearest futures series does not allow direct calculation of the profitability of a trading system. The solution adopted here is to use spread-adjusted continuous price series, by which, except for the most recent contract in the series, prices are adjusted by a constant that compensates for price differences that exist at rollover dates (Schwager, 1984). This method alters the prices of the future contracts prior to the most recent contract but maintains identical price relationships, thereby avoiding the distortions mentioned above.

The transition between contracts was performed 7 days before expiration. Several sets of data were prepared: each contained a training set of 424 consecutive trading days, a test set of 100 consecutive trading days (which begins the day after the training set ends), and a validation set of 100 consecutive trading days (which begins the day after the test set ends).

RNNs are predisposed, as are standard backpropagation networks, to overfit training data. Rather than learning the fundamental structure of the training set, which would enable them to generalize adequately, RNNs learn insignificant details of individual cases. This problem is generated by two conflicting purposes of ANNs: they have to be as general as possible, so that they learn a broad range of problems, and yet they need to perform well in out-of-sample tests, on examples not previously seen. There are two approaches to the overfitting problem. The first is to train the model on the training set and to evaluate the model's performance on the test set. The second approach is to use one of the many network pruning algorithms (Weigend et al., 1992) to reduce the network size, thereby limiting the number of hidden neurodes and hence the number of parameters to be estimated. The solution I adopted is based on a parsimonious choice of the number of hidden neurodes as suggested by the generalization capability of the network on the test set. In this procedure, I trained the network until convergence, observed the point at which the test set error began to rise, and then restored the network weights at the iteration cycle where the test set error was minimum. How well the network generalized was deduced by analyzing its performance on the validation set, and not on the test set, as this was used to decide when to stop training and therefore introduced a dangerous bias in the evaluation.

The ultimate goal of the experiments is to create a trading system, a set of interrelated rules to enter and exit the market, that produces profits. While accuracy is related to profitability, the trading system should not be evaluated using only standard statistical error measures (mean square error and the like). As an example, a trading system might consistently miss a large number of small moves but correctly forecast a small number of large moves. Therefore the researcher must take into account the out-of-sample profitability of the system, as well as its forecasting accuracy, when choosing the neural architecture, activation functions, data sets, and forecast horizon.

Choosing the kind of outputs to be forecast is an important decision. The most common options are actual price values, first differences of prices, returns, and binary signals, such as -1 short (sell 1 future contract), 1 long (buy 1 future contract). One of the problems in forecasting actual prices is that activation functions tend to emphasize the importance of intermediate output values, so that the range of predicted values is compressed with respect to target values. Solutions range from using special forms of normalization to linear activation functions. The experiments were conducted by using price changes as the output. As currency futures are nonstationary, it is better to analyze price changes in terms of compound return: $r_{ct} = \log (f_t) - \log (f_{t-1})$.

Another critical point is to identify the appropriate set of inputs relevant for the RNN architecture and for the chosen output. In particular, the inputs should be adapted to the "needs" of RNNs: they should have a temporal structure and should not be too numerous. An analysis was performed of alternative sets of inputs based on transformations of the set of data, drawing from the vast base of technical indicators. This selection was based on previous work performed on the optimal choice of parameters of different technical indicators and on their combined use in trading systems (Tenti, 1991). Inputs included the compound returns of the last *n* periods (where n = 1, 2, 3, 5, 8), the running standard deviation of the *k* last periods (where k = 13, 21, 34), and technical indicators such as the average directional movement index (ADX) (Wilder, 1978), trend movement index (TMI) (Bookstaber, 1985), rate of change (ROC) (Kaufman, 1980), and Ehlers leading indicator (ELI) (Elhers, 1992). Inputs were normalized to zero mean and two standard deviations.

TRADING STRATEGY

The forecast formulated by the three versions of RNNs is just the initial part of a trading strategy. The transformation from predictions into market actions is obtained by specifying a set of rules to buy and sell currency futures. In particular, according to the uses of the Chicago Mercantile Exchange, two types of orders were used:

- Market Opening Only Order, where the order is filled only during the opening range at the first available offer (sell order) or bid (buy order); and
- Market on Close Order, where the order is filled at any time during the closing range.

I account for two types of costs in executing a trade—transaction costs and slippage costs. Transaction costs are the charges levied by a brokerage firm to buy or to sell a future contract. Slippage costs reflect the reality of a moving market and are the difference between the theoretical execution price and the actual fill price. It was assumed that one contract should be traded at a time and also that transaction costs should equal \$80 per trade (\$25 for commission and \$55 for slippage). Transaction costs are very important in short-term trading systems because they can have a dramatic impact on performance.

The network's objective was to forecast the compound return of the following day's Open (O) and Close (C). Using prices up to the Open at time $t(O_t)$, the forecast was made two steps ahead for the Open at time t + 2 (O_{t+2}). Similarly, using prices up to C_{t+1} the forecast was made for C_{t+3} . For the sake of generating trading signals, it is possible to build a continuous price series F_t, \ldots, F_{t+n} by alternating O_t and C_t prices.

Trading Strategy 1

Entry rule at time t:

- 1. If $f(F_{t+2}) > x$, then Long F_{t+1}
- 2. If $f(F_{t+2}) < -x$, then Short F_{t+1}
- 3. If $-x < f(F_{t+2}) < x$, then Flat

Exit rule at time t + 1:

1. If $f(F_{t+3}) > x$, then stay Long (if 1 at *t*), stop and reverse to Long (if 2), go Long (if 3)

2. If $f(F_{t+3}) < -x$ then stay Short (if 2 at *t*), stop and reverse to Short (if 1), go short (if 3)

3. If $-x < f(F_{t+3}) < x$ then cover Short (if 2), cover Long (if 1), stay Flat (if 3)

where f() stands for the network's compound return forecast and x is a numerical filter.

Trading Strategy 2

Entry rule at time t + 1:

- 1. If $f(F_{t+2}) > [r(F_{t+1}) + x]$, then Long F_{t+1}
- 2. If $f(F_{t+2}) \le [r(F_t+1) x]$, then Short F_{t+1}
- 3. If $[r(F_{t+1}) x] \le f(F_{t+2}) \le r(F_{t+1}) + x]$, then Flat

Exit rule at time t + 2:

1. If $f(F_{t+3}) > [r(F_{t+2}) + x]$, then stay Long (if 1 at t + 1), stop and reverse to Long (if 2), go Long (if 3)

2. If $f(F_{t+3}) < [r(F_{t+2}) - x]$, then stay Short (if 2 at t + 1), stop and reverse to Short (if 1), go Short (if 3)

3. If $[r(F_{t+2}) - x] \le f(F_{t+3}) \le [r(F_{t+2}) + x]$, then cover short (if 2 at t + 1), cover long (if 1), stay flat (if 3)

where r() is the compound return.

Trading strategy 1 is more realistic than trading strategy 2 because it forecasts and decides at time t to go Long (predicting a rise in price), Short (predicting a fall in price), or Flat (inactive). The purchase or sale of the future is then done at the next time step t + 1. Trading strategy 2, on the other hand, forecasts at time t but waits until time t + 1 to compare the subsequent market Open or Close with the forecast and then decides whether to buy, sell, or do nothing. For this second strategy the slippage is likely to be significantly larger than for the first because the fill is not made at the Open but immediately after, and a decision must be reached within a few seconds of the Close.

Any sensible trading strategy should somehow restrict the number of trading signals because of the incidence of transaction costs. The filter x was used to provide a way to avoid false signals as much as possible. Its size was optimized using genetic algorithms based on the profitability of the trading strategy across different RNN versions and periods. The fitness function was the cumulative profit obtained over the training, test, and validation data sets for each trading strategy. The filter was subject to the constraint of being a positive number. An additional constraint was that each one of the three data sets had to show a profit. (Filter values that were used in the experiments are shown in parentheses in Tables 1–3.)

EVALUATION

The different versions of RNN were compared by focusing on the accuracy and reliability of the forecasts on training, test, and validation data. Differences in the architecture yield significantly different results. A standard error measure to evaluate the quality of predictions is the normalized mean squared error (NMSE):

NMSE =
$$\frac{\sum_{f \in \tau} (\text{observation}_{t} - \text{prediction}_{t})^{2}}{\sum_{f \in \tau} (\text{observation}_{t} - \text{mean}_{t})^{2}}$$

where t = 1, ..., N enumerates the patterns in each data set (τ) used. The above is the ratio between the mean squared errors of both the prediction method and the method that forecasts by using the mean at every step. A value of NMSE = 1 thus corresponds to the value obtained by simply predicting the average. Yet, prediction accuracy statistics such as NMSE by themselves are of little use. The purpose is rather to build trading systems that would provide a consistent profitability on a risk-adjusted basis, with a high degree of confidence.

The results in terms of profitability of the trading strategy, net of trading commissions, and slippage are shown for the different versions. Margin requirements are usually satisfied by posting Treasury Bills. The interest income earned is not accounted for in the following performances. Therefore reported net profits are based on trading profits only and represent the return earned in excess of the Treasury Bill rate. In itself, mere profitability is not enough to evaluate the relative value of a trading system. Profit has to be computed across several different periods, as it could be an expression of one isolated period of extraordinary performance. In addition, other measures of relative performance are needed, such as

$$ROE = \left(1 + \frac{\text{net profit}}{\text{MaxValueFuture}}\right)^{(\text{days/255})} - 1$$

Return on equity (ROE) measures the relative unlevered profitability, being the annualized ratio between the net profit and the maximum value of the Deutsche mark future contract in the period analyzed.

$$ROC = \left(1 + \frac{\text{net profit}}{2(\text{MaxDr} + \text{InMar})}\right)^{(\text{days/255})} - 1$$

Return on capital (ROC) expresses the dollar net profit relative to the funds required for trading by an individual trader who is subject to double margins. It is the annualized ratio of the net profit over twice the sum of the maximum drawdown, the largest cumulative loss on which the particular trading system would incur, plus the initial margin required. It is a measure of the efficiency in the use of capital of a trading system. Tables 1–3 show the net profit, the percentage of correct trading signals, ROE, ROC, and NMSE for each strategy. The values shown are for the last of the five different time periods used (ending in December 1994) of the best network configurations. Each RRN version and trading strategy has its own filter, optimized using genetic algorithms.

Combining the use of the validation set (in addition to the test set) with the use of different periods adds greater reliability to the trading systems. Results here are reported only for the Deutsche mark, though similar performances were obtained on other currencies. Although it is probably unrealistic to expect any single system

	Training	Test	Validation
# of trading days	424	100	100
TrSt1 (0.13)	\$37,488	\$2,925	\$2,725
TrSt2 (0.17)	\$20,275	\$1,950	\$2,800
%TrSig1	69.1%	69.4%	55.6%
%TrSig2	45.2%	46.9%	45.7%
ROE TrSt1	24.9%	9.1%	8.6%
ROE TrSt2	13.9%	6.0%	8.7%
ROC TrSt1	375.7%	124.3%	241.1%
ROC TrSt2	203.2%	58.2%	234.3%
NMSE	0.8949	0.9622	0.9699

Table 1. Performance measures for RNN1

TrSt1, trading strategy 1; TrSt2, trading strategy 2; TrSig1, trading signal 1; TrSig2, trading signal 2; ROE, return on equity; ROC, return on capital; NMSE, normalized mean squared error.

to work in all markets, a good system should demonstrate profitability at least in related markets, such as is the case of currency futures. In addition, if a trading strategy is devised on currency futures, very similar results should be expected by using forward contracts.

RNN2 provides the best overall profitability. It is the best of the three versions even if it is judged in terms of smaller decay of performance going from the training set to the test set, and then from the test set to the validation set. RNN1's results were not quite as good as those of RNN2, while RNN3 did not show a good generalization capability. However, these results should not be taken as the final verdict on the relative merits of the three versions of RNN. As far as trading strategies are concerned, strategy 2 had greater accuracy in forecasting large price movements than strategy 1, even though the percentage of correct trading signals was signifi-

	Training	Test	Validation
# of trading days	424	100	100
TrSt1 (0.13)	\$40,675	\$3,925	\$8,437
TrSt2 (0.17)	\$35,713	\$3,788	\$8,262
%TrSig1	64.4%	63.3%	63.5%
%TrSig2	54.4%	52.1%	51.0%
ROE TrSt1	26.9%	12.4%	27.7%
ROE TrSt2	23.8%	11.9%	27.1%
ROC TrSt1	217.8%	256.9%	353.7%
ROC TrSt2	191.2%	247.9%	383.5%
NMSE	0.9519	0.9683	0.9795

Table 2. Performance measures for RNN2

See Table 1 footnote for definitions of abbreviations.

	Training	Test	Validation
# of trading days	424	100	100
TrSt1 (0.05)	\$28,413	\$200	\$1,250
TrSt2 (0.04)	\$36,063	\$1,625	-\$625
%TrSig1	48.9%	46.0%	48.50%
%TrSig2	43.8%	44.1%	43.00%
ROE TrSt1	24.0%	0.6%	3.9%
ROE TrSt2	28.8%	5.0%	-1.9%
ROC TrSt1	185.1%	3.2%	8.7%
ROC TrSt2	232.4%	15.3%	-9.3%
NMSE	0.9654	0.9984	1.018

Table 3. Performance measures for RNN3

See Table 1 footnote for definitions of abbreviations.

cantly smaller. This resulted in absolute and relative profitabilities that were slightly less than those obtained by trading strategy 1.

Comparisons with standard backpropagation methods have shown that RNN1 and RNN2 have better profitability and generalization capacities. However, as there are an almost infinite number of configurations and parameters in standard backpropagation, I cannot say that it would be impossible to find one that could yield better results than RNN1 and RNN2. However, in repeated tests it has been the case that RNN1 and RNN2 outperformed standard backpropagation.

CONCLUSIONS

RNNs, often avoided because of fears of time-consuming training sessions, are particularly useful for financial forecasting applications. The methods described here are equally applicable to other markets. They are particularly well suited to forecasting foreign exchange markets because of the network's adherence to nonlinearities as well as the subtle regularities found in these markets.

The above findings can be considered preliminary, as I am in the process of expanding my research to the following areas:

- comparisons of ANNs with standard statistical techniques;
- · comparisons of ANNs with mechanical trading systems;
- application of the Modern Portfolio Theory framework to ANN financial forecasting;
- · diversification of trading systems through the use of regime-switching models; and
- development of criteria to be used in the evaluation phase.

The evaluation of forecasting techniques is a very complex task involving many factors; the most important one is the application of the technique to real-world situations. By applying RNNs successfully to trading in the currency futures markets, it appears that RNNs are not a "passing fad," as critics would have us believe.

REFERENCES

- Adam, O., J. L. Zarader, and M. Milgram. 1994. Identification and prediction of non-linear models with recurrent neural network. Laboratoire de Robotique de Paris.
- Bilson, J. F. O. 1992. Technical currency trading. Chicago, Ill.: The Chicago Corporation.
- Bookstaber, R. 1985. The complete investment book. Glenview, Ill.: Scott, Foresman and Company.
- Brock, W. A., D. A., Hsieh, and B. LeBaron. 1991. Nonlinear dynamics, chaos, and instability. Cambridge, Mass.: MIT Press.
- Connor, J., and L. Atlas. 1993. Recurrent neural networks and time series prediction. *Proceedings of the International Joint Conference on Neural Networks* 1:301–306.
- De Grauwe, P., H. Dewachter, and M. Embrechts. 1993. Exchange rate theory. Chaotic models of foreign exchange markets. London: Blackwell.
- Ehlers, J. F. 1992. MESA and trading market cycles. New York: Wiley.
- Fang, H., K. S. Lai, and M. Lai. 1994. Fractal structure in currency futures price dynamics. Journal of Futures Markets 14:169–181.
- Frasconi, P., M. Gori, and G. Soda. 1992. Local feedback multilayered networks. Neural Computation4:120–130.
- Jordan, M. I. 1986. Serial order: A parallel distributed processing approach. Institute of Cognitive Science Report 86-104, University of California–San Diego.
- Kamijo, K., and T. Tanigawa. 1990. Stock price pattern recognition: A recurrent neural network approach. In Proceedings of the International Joint Conference on Neural Networks, 1215–1221.
- Kaufmann, P. J. 1980. The new commodity trading systems and methods. New York: Wiley.
- LeBaron, B. 1993. Practical comparisons of foreign exchange forecasts. *Neural Network World* (November):779–791.
- Levich, R. M., and L. R. Thomas. 1993. The significance of technical trading-rule profits in the foreign exchange market: A bootstrap approach. *Journal of International Money and Finance* 12:451–465.
- Logar, A. M., E. M., Corwin, and W. J. B. Oldham. 1993. A comparison of recurrent neural network learning algorithms. In Proceedings of the International Joint Conference on Neural Networks, 1129–1134.
- Lukac, L. P., B. W. Brorsen, and S. H. Irwin. 1988. A test of futures market disequilibrium using twelve different technical trading systems. *Applied Economics* 20:623–639.
- Mozer, M. C. 1993. Neural net architectures for temporal sequence processing. In *Time series prediction: Forecasting the future and understanding the past*, eds. A. S. Weigend and N. A. Gershenfeld, 243–264. Reading, Mass.: Addison-Wesley.
- Pardo, R. 1992. Design, testing, and optimization of trading systems. New York: Wiley.
- Refenes, A. N., M. Azema-Barac, L. Chen, and S. A. Karoussos. 1993. Currency exchange rate prediction and neural network design strategies. *Neural Computing and Applications* 1:46–58.
- Rumelhart, D. E., G. H. Hinton, and R. J. Williams. 1986. Learning internal representations by error propagation. In *Parallel distributed processing*, eds. D. E. Rumelhart, McClelland, and the PDP Research Group, 318–362. Cambridge, Mass.: MIT Press.
- Schwager, J. D. 1984. A complete guide to the futures markets. New York: Wiley.
- Sweeney, R. J. 1986. Beating the foreign exchange market. Journal of Finance 41:163–182.
- Taylor, S. J. 1986. Modelling financial time series. New York: Wiley.
- Taylor, S. J. 1994. Trading futures using a channel rule: A study of the predictive power of technical analysis with currency examples. *Journal of Futures Markets* 14:215–235.
- Tenti, P. 1991. Optimal selection of parameters in technical trading systems. Boston University-Rome.

- Weigend, A. S., B. A. Huberman, and D. E. Rumelhart. 1992. Predicting sunspots and exchange rates with connectionist networks. In *Nonlinear modeling and forecasting*, eds. M. Casdagli and S. Eubank, 395–432. Reading, Mass.: Addison-Wesley.
- Werbos, P. 1974. Beyond regression: New tools for prediction and analysis in the behavioral sciences. Ph.D. thesis, Harvard University, Cambridge, Mass.
- Wilder, W. J. 1978. New concepts in technical trading systems. McLeansville, N. C.: Trend Research.

PAGE 582 BLANK