Studies in Nonlinear Dynamics and Econometrics

Quarterly Journal
January 1998, Volume 2, Number 4
The MIT Press

Studies in Nonlinear Dynamics and Econometrics (ISSN 1081-1826) is a quarterly journal published electronically on the Internet by The MIT Press, Cambridge, Massachusetts, 02142. Subscriptions and address changes should be addressed to MIT Press Journals, Five Cambridge Center, Cambridge, MA 02142; tel.: (617) 253-2889; fax: (617) 577-1545; e-mail: journals-orders@mit.edu Subscription rates are: Individuals $40.00, Institutions $130.00. Canadians add 7% GST. Prices subject to change without notice.

Subscribers are licensed to use journal articles in a variety of ways, limited only as required to insure fair attribution to authors and the Journal, and to prohibit use in a competing commercial product. See the Journal's World Wide Web site for further details. Address inquiries to the Subsidiary Rights Manager, MIT Press Journals, Five Cambridge Center, Cambridge, MA 02142; tel.: (617) 253-2864; fax: (617) 258-5028; e-mail: journals-rights@mit.edu.

©1998 by the Massachusetts Institute of Technology

Forecasting Exchange Rates Using Neural Networks for Technical Trading Rules

Philip Hans Franses
Econometric Institute and
Rotterdam Institute for Business Economic Studies
Erasmus University, Rotterdam, The Netherlands
franses@few.eur.nl

Kasper van Griensven
ABN-AMRO Bank, Amsterdam, The Netherlands

Abstract. *We examine the performance of artificial neural networks (ANNs) for technical trading rules for forecasting daily exchange rates. The main conclusion of our attempt is that ANNs perform well, and that they are often better than linear models. Furthermore, the precise number of hidden layer units in ANNs appears less important for forecasting performance than is the choice of explanatory variables.*

Acknowledgments. We thank an anonymous referee and the Special Issue editor for helpful comments.

Keywords. technical analysis, neural networks

1 Introduction

In this short paper, we aim to provide a modest contribution to the literature on forecasting daily exchange rates using nonlinear time series models. When various diagnostic tests for linearity are applied to exchange-rate data, one almost invariably finds evidence of nonlinearity; see, e.g., Brooks (1996) for a recent survey of the literature, and Hsieh (1988, 1989) and Scheinkman and LeBaron (1989) for earlier accounts. Given this evidence, the natural question often is whether any nonlinearity can be exploited for improved forecasting. The most commonly considered nonlinear models are the threshold model and the bilinear model [see Chappell et al. (1996) and Brooks (1997) among others], and the artificial neural network (ANN) model [see, for example, Kuan and Liu (1995)]. The nonlinear models reported in the literature almost always include lagged returns on the exchange rates themselves as explanatory variables. The general conclusion from the current empirical literature appears to be that nonlinear time series models do not seem to forecast well [see Brooks (1997) for a recent survey], although the results in Kuan and Liu (1995) and Swanson and White (1995) suggest that ANNs sometimes are favorable models for successful out-of-sample forecasting.

Given the disappointing findings so far, it seems important to examine the possible causes of the failure of commonly applied nonlinear models to yield proper forecasts. One cause may be that current tests and

© 1998 by the Massachusetts Institute of Technology

models for linearity are misleading, and that these mistakes neglected GARCH for nonlinearity; see Brooks (1996). Extensive simulation results in Franses and Van Homelen (forthcoming), however, suggest that neglected GARCH does not appear to lead to spuriously successful ANNs in terms of forecasting. Additionally, the simulations in that article show that *if* there is some form of nonlinearity (such as bilinearity and threshold nonlinearity), ANNs do exploit this for improved forecasting, thereby excluding a possible second cause of nonlinear model failure, which is that the models may not capture the nonlinear features of the data. In this paper, we therefore investigate a third possible cause, which is that the input variables for the nonlinear models may be inappropriate. Following Gencay (1996), we examine the potential usefulness of moving averages of past exchange rates as explanatory variables. To save space, we restrict our attention to ANNs, especially since these models can approximate nonlinear functions of the data arbitrarily close. Notice that by considering ANNs, there is an advantage that one can sidestep complicated model-selection issues for often non-nested nonlinear models.

The consideration of moving averages of past data as predictors for future patterns is often referred to as *technical analysis* (TA) or momentum strategy. An example of a very simple technical trading rule (TTR) is that one buys today if yesterday's exchange rate exceeds a long-term moving average of the exchange rates over the last, say, 100 days. A more detailed account of this TTR and much more advanced TTRs can be found in Plummer (1990). Surveys among dealers show that many use TTRs to make decisions on buying and selling. For example, interviews reported in Taylor and Allen (1992) show that at least 90% of the respondents (among 353 chief foreign-exchange-rate dealers in London) say that they put weight on technical analysis when forming views for one or more time horizons. On the theoretical side, Sweeny (1986) reports on the profitability of some rules for the German mark exchange rate, while Brock, Lakonishok, and LeBaron (1992) use simulations to demonstrate that TTRs can outperform parametric models with respect to out-of-sample forecasting of stock returns. Finally, Neftci (1984) shows the predictability of TTRs for specific futures prices.

In an important article, Neftci (1991) demonstrates that technical trading rules can only be exploited usefully if the underlying process is nonlinear. The intuition for this result is that the descriptive models corresponding to TTRs are, so to say, piecewise linear models for the (returns on) exchange rates, where there typically are thresholds between the different regimes. This implies that data generated from TTRs show nonlinear patterns, and hence that a model for observed returns on exchange rates, when these are "explained" by certain moving averages of past exchange rates, should best be nonlinear. Since an ANN is a very general and flexible nonlinear model, we examine in this paper the use of ANNs for forecasting returns on exchange rates when the explanatory variables are functions of past averages of the exchange rates.

Our approach to link ANNs with TA is not new. To our knowledge, Gencay (1996) is the first study that uses TA as input for ANNs, in an attempt to forecast stock-market returns. It is found that such ANNs can yield a successful forecasting record. In this paper we focus on daily exchange rates. Other differences between our approach and that in Gencay (1996) is (1) that we do not a priori select a model using some model-selection criterion like Schwarz's or Akaike's [also because Swanson and White (1995) suggest that these criteria seem unreliable for selecting the best forecasting model], and (2) that we evaluate the directional accuracy of the forecasts instead of mean-squared prediction error (MSPE). This measure corresponds with the outcome of technical analysis, which is to buy or to sell. For simplicity, we abstain from other measures, for example, the probability measure proposed in Leitch and Tanner (1991). To our knowledge, there is yet no consensus on which criterion is best, so we opt for our measure simply because it is simple to calculate and easy to understand. It is our experience that the MSPE is seldom used in financial practice.

The outline of our short paper is as follows. In Section 2 we review the ANNs we consider, and the linear model. In Section 3, we apply these models to forecasting five daily exchange-rate series, and we summarize our findings in a single table. In Section 4, we conclude with some remarks.

2 Forecasting Models

In this section we provide some information on the TTRs and ANNs we use in the empirical analysis. Our choice for the TTRs is inspired by Gencay (1996), and our selection for the ANN by Kuan and Liu (1995) and Swanson and White (1995). We denote the exchange rate as y_t, where t is a daily index, and we define the return on the exchange rate as $r_t = \log y_t - \log y_{t-1}$. Consider the moving average $m_t(n)$ when it is defined as:

$$m_t(n) = n^{-1} \sum_{i=0}^{n-1} y_{t-i}. \tag{1}$$

Clearly, for $n = 1$, $m_t(n)$ equals y_t. Very simple technical trading rules consider the signal $s_t(n_1, n_2)$ defined by:

$$s_t(n_1, n_2) = m_t(n_1) - m_t(n_2), \tag{2}$$

where $n_1 < n_2$. When $s_t(n_1, n_2)$ exceeds zero (or some other preset value), the short-term moving average exceeds the long-term moving average to a certain extent, and a "buy" signal is generated. Conversely, when $s_t(n_1, n_2)$ is negative (or below a certain threshold), a "sell" signal is given. For daily data, typical choices in practice are $n_1 = 1$ or 5, and $n_2 = 50$, 100, or 150.

If a moving average rule like Equation (2) is an explanatory variable for the returns r_t, one may want to consider the linear model

$$r_t = \alpha s_{t-1}(n_1, n_2) + \beta + \varepsilon_t, \tag{3}$$

where α and β are unknown parameters and ε_t is some kind of error process. A motivation for Equation (3) is that when a buy signal is given at $t - 1$, and $s_t(n_1, n_2)$ is a relevant explanatory variable, one would expect that r_t is indeed positive; and when a sell signal is given, r_t should be negative. This already suggests that one may be interested more in the sign of the forecasted r_t in conjunction with the sign of the observed r_t than in the specific values of the forecast errors. We will use Equation (3) as the benchmark model. Other studies on forecasting exchange rates have used, for example, linear ARMA models. As these ARMA models appear not to work well, we choose to use Equation (3), especially since it is nested in our ANN, as we will show below.

Theoretically, according to the arguments of Neftci (1991), Equation (3) cannot be an adequate descriptive model. In this paper we therefore consider Equation (3) only as the benchmark model that should be beaten by a nonlinear model. Given the potential usefulness of ANNs, we choose to modify Equation (3) as follows:

$$r_t = \alpha s_{t-1}(n_1, n_2) + \sum_{j=1}^{q} \beta_j G[\gamma_j s_{t-1}(n_1, n_2)] + \varepsilon_t, \tag{4}$$

with

$$G(a) = (1 + \exp(-a))^{-1},$$

where α and γ_j are unknown parameters for $j = 1, 2, \ldots, q$. This ANN model contains $2q + 1$ parameters. Kuan and White (1994) call it a feedforward single-hidden-layer neural network with q units in the hidden layer. The function $G(a)$ is often taken as the logistic activation function, which connects the q input components with the hidden layer. The hidden layer is connected with the output variable r_t through the β_j parameters. Using a Taylor expansion of $G(a)$ around $a = 0$, it is seen that Equation (3) is nested within Equation (4). The expression in Equation (4) indicates that it allows the intercept in Equation (3) to be time varying and to depend on the input variable. This is one of the main features of an ANN. Finally, given the simulation results in Franses and Van Homelen (forthcoming), where GARCH effects are shown not to lead to the spurious forecasting success of ANNs, we do not modify Equation (4) by assuming specific GARCH

properties of the error process ε_t. In other words, when an ANN works for forecasting, it is unlikely owing to neglected GARCH.

The model shown by Equation (4) is often used in the empirical analysis of financial time series; see, for example, Kuan and Liu (1995) and Swanson and White (1995). One possible motivation for its frequent use is that Kuan and White (1994), among others, show that with large enough q, an ANN can approximate (with arbitrary precision) any function f that connects r_t with explanatory variables. This property of ANNs makes them useful for rather precise pattern recognition in time-series data; see, for example, Bishop (1995), Ripley (1994), and Franses and Draisma (1997). Whether they are also of use for out-of-sample forecasting is a matter of recent research in several of the aforementioned studies.

We estimate the parameters in the ANN by minimizing the residual sum of squares $\Sigma\varepsilon_t^2$. In the first round of optimization, we impose a penalty on the parameter values as advocated in Ripley (1994). In the last estimation stage, this penalty is removed. Our estimation procedure uses the simplex method of Nelder and Mead and the Broyden-Fletcher-Goldfarb-Shanno algorithm, of which the latter is available in the Gauss Optimization Toolbox (version 3.1.1). The Gauss programs are available from the authors on request.

We estimate the parameters in Equation (4) for $q \in \{1, 2, 3, 4, 5\}$. The linear and ANN models are used to generate one-day-ahead forecasts, both for within-sample and for out-of-sample data. To mimic a practically realistic situation, we investigate whether the sign of the fitted return $\hat{r}_{i+1}$ equals the sign of the truly observed return r_{i+1}. For this purpose, we define the success ratio (SR) as $\mathrm{SR} = k^{-1} \sum_{i=1}^{k} I_i[\hat{r}_{i+1}.r_{i+1} > 0]$, where $I_i[.]$ is an indicator function that takes a value of 1 when its argument is true and a value of 0 otherwise, and where k is the number of forecasts. We test whether the SR value differs significantly from an SR value that would be obtained in the case where the processes that generate $\hat{r}_{i+1}$ and r_{i+1} are independent. For this purpose, we use the test proposed in Pesaran and Timmermann (1992). When this PT test value is positive and significant, the model yields good forecasts, and when it is negative and significant, the model should not be used. In this paper we evaluate the null hypothesis in a two-sided test procedure because we also want to allow for significantly poor forecasts, if there are any. Our decisions are based on the assumption of a 10% significance level.

3 Empirical Results

The data we use in this paper are daily exchange-rate data on the German mark (GM), U.S. dollar (USD), Dutch guilder (DG), Japanese yen (JY), and the Swiss franc (SF), all measured in British pounds. We have daily observations for 1981, to and including 1995. Missing values (owing to holidays and religious festivals) are replaced by immediately preceding observations.

We estimate the model parameters for the sample 1981 to and including 1993. One-day-ahead forecasts for 1994 and 1995 are generated, conditional on the available information. Hence, to forecast for day i in 1994, we assume the availability of information for days $i-1$, $i-2$ and so on. The parameters are only estimated once, and they are kept fixed for 1994 and 1995. In what follows, we set $n_1 = 1$ and $n_2 = 50$ or 150.

Before we turn to a discussion of our results, we must stress that we are aware of the facts that our empirical approach can be refined in various ways: for example, we can use smaller samples, reestimate parameters for moving-windows samples, and forecast several steps ahead and make n_1 and n_2 time varying and depend on past fluctuations. In the present paper, we abstain from all these possible modifications, simply because we only aim to obtain a first impression of the usefulness of linking ANNs with TA.

We summarize our main findings in Table 1. More detailed results can be obtained from the authors upon request. In this table we report the number of times (out of the five exchange rates considered) that the PT test is positive (or negative) and significant at the 10% level. Several conclusions can be drawn from this table. The first is that within-sample forecasts (where the return to be forecasted is also included when estimating parameters) are always better than genuine out-of-sample forecasts. This proves again the importance of

Table 1
Number of Times (out of Five Exchange Rates) that the Test Statistic for Directional Accuracy is Positive and Significant at the 10% Level

Trading Rule	Hidden Layer Units, q	Estimation Sample: 1981–1993	Forecasting Sample: 1994	Forecasting Sample: 1995
$n_1 = 1$ $n_2 = 50$	0 (linear)	2	0	0[a]
	1	5	1 (SF)	1 (USD)
	2	4	0	1 (USD)
	3	5	1 (SF)	1[a] (USD)
	4	5	1 (SF)	1 (USD)
	5	5	0	1 (USD)
$n_1 = 1$ $n_2 = 150$	0 (linear)	5	0	0[a]
	1	5	2 (JY,SF)	2 (GM,DG)
	2	5	2 (JY,SF)	2 (GM,DG)
	3	4	2 (JY,SF)	1 (GM)
	4	5	1 (SF)	2 (GM,DG)
	5	5	2 (JY,SF)	2 (GM,DG)

[a]For $q = 0$ and forecasting sample 1995, the PT test is three times significantly negative for $n_1 = 1, n_2 = 50$, and once significantly negative for $n_1 = 1, n_2 = 150$. For $q = 3$ and forecasting sample 1995, the PT is significant and negative for the JY. In parentheses, we give the exchange rates for which useful forecasts can be obtained: SF denotes Swiss franc, USD is U.S. dollar, JY is Japanese yen, GM is German mark, and DG denotes Dutch guilder

evaluating models on hold-out data, and it reinforces the statements in Brooks (1997). The second conclusion is that the linear model never yields good out-of-sample forecasts, even though the within-sample results tend to suggest otherwise. In fact, sometimes the linear model even predicts poorly (that is, obtains negative and significant PT test values). This outcome corresponds with the theoretical results where it is argued that TA fits in a nonlinear environment. Notice, by the way, that successful forecasting performance applies to both 1994 and 1995, while also all five exchange rates appear forecastable, at least once. The (1,150) TTR seems best, though. The third main conclusion is that the number of hidden-layer units in the ANNs does not seem to matter much with regard to forecasting. Hence, perhaps one needs to put less effort in designing selecting criteria for q for the within-sample data. In fact, it seems that more research is needed as to which variables are to be included as explanatory variables. In sum, our results add to those in Gencay (1996), and suggest that ANNs for technical trading rules can yield good forecasts.

4 Final Remarks

This short paper showed the potential usefulness of artificial neural networks for technical trading rules to forecast daily exchange-rate data. Hence, the recently documented poor performance of ANNs may be owing to the inclusion of inappropriate explanatory variables.

Of course, our first attempt in this paper can and should be refined in various ways. Other more sophisticated trading rules, different sample sizes, and various exchange rates can be considered; and one may want to reestimate the parameters for every new sample. Also, alternative versions of ANNs, which can pick up even more complicated relations between returns and explanatory variables, can be considered. To us, the results in the present paper only suggest that further consideration of ANNs for technical trading rules is a fruitful enterprise.

References

Bishop, C. M. (1995). *Neural Networks for Pattern Recognition*. Oxford: Oxford University Press.

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

Brooks, C. (1996). "Testing for non-linearity in daily sterling exchange rates." *Applied Financial Economics*, 6:307–317.

Brooks, C. (1997). "Linear and non-linear (non-)forecastability of high-frequency exchange rates." *Journal of Forecasting*, 16:125–145.

Chappell, D., J. Padmore, P. Mistry, and C. Ellis (1996). "A threshold model for the French franc-deutschmark exchange rate." *Journal of Forecasting*, 15:155–164.

Franses, P. H., and G. Draisma (1997). "Recognizing changing seasonal patterns using artificial neural networks." *Journal of Econometrics*, 81:273–280.

Franses, P. H., and P. Van Homelen (forthcoming). "On forecasting exchange rates using neural networks." *Applied Financial Economics*.

Gencay, R. (1996). "Non-linear prediction of security returns with moving average rules." *Journal of Forecasting*, 15:165–174.

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

Hsieh, D. A. (1989). "Testing for nonlinear dependence in daily foreign exchange rates." *Journal of Business*, 62:339–368.

Kuan, C. M., and T. Liu (1995). "Forecasting exchange rates using feedforward and recurrent neural networks." *Journal of Applied Econometrics*, 10:347–364.

Kuan, C. M., and H. White (1994). "Artificial neural networks: An econometric perspective (with discussion)." *Econometric Reviews*, 13:1–91.

Leitch, G., and G. E. Tanner (1991). "Economic forecast evaluation: Profit versus the conventional error measures." *American Economic Review*, 81:580–590.

Neftci, S. N. (1984). "Can chartists outperform the market? Market efficiency tests for 'technical analysis.'" *Journal of Futures Markets*, 4:465–478.

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

Pesaran, M. H., and A. Timmermann (1992). "A simple nonparametric test of predictive performance." *Journal of Business and Economic Statistics*, 10:461–465.

Plummer, T. (1990). *Forecasting Financial Markets; Technical Analysis and the Dynamics of Price*. New York: Wiley.

Ripley, B. D. (1994). "Neural networks and related methods for classification." *Journal of the Royal Statistical Society B*, 56:409–456.

Scheinkman, J. A., and B. LeBaron (1989). "Nonlinear dynamics and stock returns." *Journal of Business*, 62:311–337.

Swanson, N. R., and H. White (1995). "A model-selection approach to assessing the information in the term structure using linear models and artificial neural networks." *Journal of Business and Economic Statistics*, 13:265–275.

Sweeny, R. J. (1986). "Beating the foreign exchange market." *Journal of Finance*, 41:163–182.

Taylor, M. P., and H. Allen (1992). "The use of technical analysis in the foreign exchange market." *Journal of International Money and Finance*, 11:304–314.