ELSEVIER
Journal of Empirical Finance 5 (1998) 347–359
Journal of EMPIRICAL FINANCE

The predictability of security returns with simple technical trading rules

Ramazan Gençay
Department of Economics, University of Windsor, 401 Sunset, Windsor, Ont., Canada N9B 3P4

Accepted 29 July 1997

Abstract

Technical traders base their analysis on the premise that the patterns in market prices are assumed to recur in the future, and thus, these patterns can be used for predictive purposes. This paper uses the daily Dow Jones Industrial Average Index from 1897 to 1988 to examine the linear and nonlinear predictability of stock market returns with simple technical trading rules. The nonlinear specification of returns are modelled by single layer feedforward networks. The results indicate strong evidence of nonlinear predictability in the stock market returns by using the past buy and sell signals of the moving average rules. © 1998 Elsevier Science B.V. All rights reserved.

JEL classification: C45; C52; C53; G14

Keywords: Market efficiency; Technical trading rules; Feedforward networks

1. Introduction

Technical traders base their analysis on the premise that the patterns in market prices are assumed to recur in the future, and thus, these patterns can be used for predictive purposes. The motivation behind the technical analysis is to be able to identify changes in trends at an early stage and to maintain an investment strategy until the weight of the evidence indicates that the trend has reversed. One popular rule for deciding when to buy and sell in a market is the moving average rule. This rule basically involves the calculation of a moving average of the raw price data. The simplest version of this rule indicates a buy signal whenever the price climbs

PII: S0927-5398(97)00022-4

above its moving average and a sell signal when it drops below. The underlying notion behind this rule is that it provides a means of determining the general direction or trend of a market by examining the recent history. For instance, an *n*-period moving average is computed by adding together the *n* most recent periods of data, then dividing by *n*. This average is recalculated each period by dropping the oldest data and adding the most recent, so the average moves with its data but does not fluctuate as much. The larger the *n*, the smoother is the moving average rule and measures a longer-term trend.

The literature on financial forecasting documented evidence on the predictability of current returns from past returns as well as the predictability of current returns from other variables such as dividend yields and various term structure variables. This literature also documented significant relationships between expected returns and fundamental variables such as the price earnings ratio, the market-to-book ratio and evidence for systematic patterns in stock returns related to various calendar periods such as the weekend effect, the turn-of-the-month effect, the holiday effect, the January effect and the predictability originating from bid–ask bounces. For instance, Lo and MacKinlay (1988) find that weekly returns on portfolios of NYSE stocks grouped according to size show positive autocorrelation. Conrad and Kaul (1988) examine the autocorrelations of Wednesday-to-Wednesday returns (to mitigate the nonsynchronous trading problem) for size-grouped portfolios of stocks that trade on both Wednesdays. Similar to the findings of Lo and MacKinlay (1988) they find that weekly returns are positively autocorrelated. Cutler et al. (1991) present results from many different asset markets generally supporting the hypothesis that returns are positively correlated at the horizon of several months and negatively correlated at the 3–5 year horizon. Lo and MacKinlay (1990) report positive serial correlation in weekly returns for indices and portfolios and negative serial correlation for individual stocks. Chopra et al. (1992), De Bondt and Thaler (1985), Fama and French (1986) and Poterba and Summers (1988) find negative serial correlation in returns of individual stocks and various portfolios over three to ten year intervals [1]. Jegadeesh (1990) finds negative serial correlation for lags up to two months and positive correlation for longer lags. Lehmann (1990) and French and Roll (1986) report negative serial correlation at the level of individual securities for weekly and daily returns. Overall, the findings of recent literature confirm the findings of earlier literature that the daily and weekly returns are predictable from past returns and other economic and financial variables.

[1] One interpretation of negative correlation in longer horizons is that there can be substantial mean-reversion in stock market prices at longer horizons. On the other hand, the obvious concern is that inferences based on long horizon returns are based on extremely small sample size which makes the mean-reversion hypothesis fragile.

The predictability of stock market returns led the researchers to investigate the sources of this predictability. In Brock et al. (1992), two of the simplest and most popular trading rules, moving average and the trading range brake rules, are tested through the use of bootstrap techniques. They compare the returns conditional on buy (sell) signals from the actual Dow Jones Industrial Average Index to returns from simulated series generated from four popular null models: the random walk, the AR(1), the GARCH-M and the exponential GARCH (EGARCH). Brock et al. (1992) find that returns obtained from buy (sell) signals are not likely to be generated by these four popular null models. They document that buy signals generate higher returns than sell signals and the returns following buy signals are less volatile than returns on sell signals. In addition, they find that returns following sell signals are negative which is not easily explained by any of the currently existing equilibrium models. Their findings indicate that the GARCH-M model fails not only in predicting returns, but also in predicting volatility. They also document that the EGARCH model performs better than the GARCH-M in predicting volatility, although it also fails in matching the volatility during sell periods.

Overall, the findings in Brock et al. (1992) show that buy signals consistently generate higher returns than sell signals and the second moments of the distribution of the buy and sell signals behave quite differently because the returns following buy signals are less volatile than returns following sell signals. The asymmetric nature of the returns and the volatility of the Dow series over the periods of buy and sell signals indicate that the linear conditional mean estimators fail to characterize the temporal dynamics of the security returns and suggest the existence of nonlinearities as the data generation mechanism.

This paper investigates the nonlinear predictability of security returns from past returns as well as from the simplest forms of technical trading rules, namely moving average rules. The single layer feedforward networks are used as the nonlinear conditional mean estimator. To measure the performance of the nonlinear conditional mean estimator against linear specifications, the popular linear null models such as the simple AR and GARCH-M models are also studied. The evidence of predictability should contain a diligent search for out-of-sample confirmation and subsample analysis in a reasonably lengthy data set so that the problems or interpretations relating to the sample specific conditions are avoided. This paper provides a detailed out-of-sample analysis of both the linear and nonlinear conditional mean estimators over 90 years of daily data and in 22 subsamples. The ten most recent observations of each subsample are kept out for out-of-sample prediction purposes, for a total of 220 observations. The advantage of constructing the forecast horizon from 22 different subsamples is that it enables us to analyze the performance of the technical trading rules under different market conditions. This is especially important in observing the performance of these rules in trendy versus sluggish market conditions in which there is no clear trend in either direction. As a measure of performance the out-of-sample mean square prediction error (MSPE) is used. The data set is the daily Dow Jones Industrial Average Index from 1897 to 1988.

The results of this paper indicate that nonlinear conditional mean specifications provide significant forecast predictions over the linear models. The improvement in the forecast accuracy is more pronounced in the nonlinear models which use past buy–sell signals, relative to the nonlinear models which use past returns. The OLS and GARCH-M (1,1) models with past buy–sell signals provide an average of 1.65% and 2.95% forecast improvements over the benchmark model with past returns. In nonlinear conditional mean specifications, the models with past returns provide an average of 4.7% forecast improvement over the benchmark linear model with past returns. This forecast improvement is an average of 10.8% for the nonlinear model with past buy–sell signals.

In section two a brief description of the data is presented. Estimation techniques are described in section three and empirical results in section four. Conclusions follow thereafter.

2. Data summary

The data series includes the first trading day in 1897 of the Dow Jones Industrial Average (DJIA) Index to June 30, 1988, a total of 90 years of daily data. All of the stocks are actively traded and problems associated with nonsynchronous trading should be of little concern with the DJIA.

The data set is studied in subsample periods 1817–1914, 1915–1938, 1939–1962 and 1963–1988. These subsamples are chosen for two reasons. The first subsample ends with the closure of the stock exchange during World War I. The second includes both the rise of the twenties and the turbulent times of the depression. The third includes the period of World War II and ends in June 1962, the date at which the Center for Research in Securities Prices (CRSP) begins its daily price series. The last covers the period that was extensively researched because of data availability. Secondly, Brock et al. (1992) studied the same data set with the same subsamples, which makes our results comparable to theirs.

To study the sensitivity of our results to sample variation, each of these four subsamples are further split into four year subsamples. The volatile periods are the 1980–1988, 1955–1962, 1927–1930 and 1931–1938. The 1927–1930 and the 1931–1938 periods exhibit the rise of the 1920's and the times of the great depression. The 1923–1926 and 1951–1954 periods exhibit an upward trend towards the end of 1922 and 1950, respectively. The other remaining periods do not demonstrate a strong upward or downward trend when the reference point is the level of the DJIA at the starting date of that particular period.

The daily returns are calculated as the log differences of the Dow level. All of the 22 subperiods except one show signs of skewness and all periods show evidence of kurtosis. The subsamples in the 1897–1914 and 1963–1988 periods are highly skewed and have unusually high kurtosis. The source of this leptokurtic structure in the 1897–1914 period originates from the 1910–1914 period which contains the World War I period. All periods show some evidence of autocorrelation in the first lag. The Ljung-Box-Pierce statistics are also calculated for the first 10 lags and are distributed $\chi^2(10)$ under the null of identical and independent observations. Twenty-one series out of 26 give strong rejection of the null hypothesis of identical and independent observations.

3. Methodology

Let p_t, $t = 1, 2, \ldots, T$ be the daily Dow series. The return series are calculated by $r_t = \log(p_t) - \log(p_{t-1})$. Let m_t^n denote the time t value of a moving average rule of length n. m_t^n is calculated by $m_t^n = (1/n)\sum_{i=0}^{n-1} p_{t-i}$. The buy and sell signals [2] are calculated by $s_t^{n_1,n_2} = m_t^{n_1} - m_t^{n_2}$ where n_1 and n_2 are the short and the long moving averages, respectively. The rules used in this paper are $(n_1, n_2) = [(1, 50), (1, 200)]$ where n_1 and n_2 are in days. To compare the performance of the test regressions the linear regression

$$r_t = \alpha + \sum_{i=1}^{p} \beta_i r_{t-i} + \epsilon_t \quad \epsilon_t \sim ID\left(0, \sigma_t^2\right) \tag{1}$$

with lagged returns is used as the benchmark model. The test regressions are constructed such that each test regression embeds the benchmark model as a special case. The linear test regression is

$$r_t = \alpha + \sum_{i=1}^{p} \beta_i r_{t-i} + \sum_{i=1}^{p} \eta_i s_{t-i}^{n_1,n_2} + \epsilon_t \tag{2}$$

where $\epsilon_t \sim \mathrm{ID}(0, \sigma_t^2)$. In case of the GARCH-M (1,1) process the test model is written as

$$r_t = \alpha + \sum_{i=1}^{p} \beta_i r_{t-i} + \sum_{i=1}^{p} \eta_i s_{t-i}^{n_1,n_2} + \gamma h_t^{1/2} + \epsilon_t \tag{3}$$

where $\epsilon_t \sim N(0, h_t)$ and $h_t = \delta_0 + \delta_1 h_{t-1} + \delta_2 \epsilon_{t-1}^2$. The test regression for the single layer feedforward network model with d hidden units is written as

$$r_t = \alpha_0 + \sum_{i=1}^{p} \beta_{ij} r_{t-i} + \sum_{j=1}^{d} \eta_j G\left(\alpha_j + \sum_{i=1}^{p} \gamma_{ij} s_{t-i}^{n_1,n_2} \right) + \epsilon_t \quad \epsilon_t \sim ID\left(0, \sigma_t^2\right) \tag{4}$$

where G is the activation function which is chosen to

[2] The buy and sell signal is not a discrete indicator but a continuous variable in this paper.

$$G(x) = \frac{1}{1 + e^{-\alpha x}}$$

be the logistic function. It has the property of being a sigmoidal [3] function. In this paper, the network architecture is a single layer feedforward network. Many authors have investigated the universal approximation properties of neural networks (Gallant and White (1988, 1992); Cybenko (1989); Funahashi (1989); Hecht-Nielsen (1989); Hornik et al. (1989, 1990)). Using a wide variety of proof strategies, all have demonstrated that under general regularity conditions, a sufficiently complex single hidden layer feedforward network can approximate any member of a class of functions to any desired degree of accuracy where the complexity of a single hidden layer feedforward network is measured by the number of hidden units in the hidden layer. For an excellent survey of the feedforward and recurrent network models, the reader may refer to Kuan and White (1994).

The out-of-sample forecast performance of Eqs. (2)–(4) are measured by the ratio of their mean square prediction error's (MSPE) to that of the linear benchmark model in Eq. (1). To differentiate the predictive power of the past returns, the MSPE of the GARCH-M (1,1) model with lagged returns

$$r_t = \alpha + \sum_{i=1}^{p} \beta_i r_{t-i} + \gamma h_t^{1/2} + \epsilon_t \quad \epsilon_t \sim N(0, h_t),\ h_t = \delta_0 + \delta_1 h_{t-1} + \delta_2 \epsilon_{t-1}^2 \tag{5}$$

is also compared to that of the benchmark model in Eq. (1). In addition, the out-of-sample forecast performance of the single layer feedforward network model with lagged returns

$$r_t = \alpha_0 + \sum_{j=1}^{d} \beta_j G\left(\alpha_j + \sum_{i=1}^{p} \gamma_{ij} r_{t-i} \right) + \epsilon_t \quad \epsilon_t \sim ID(0, \sigma_t^2) \tag{6}$$

is also studied and compared to the benchmark model in Eq. (1). The choice for the number of hidden units in a feedforward network is determined by cross-validation.

4. Out-of-sample evidence

For each subsample the out-of-sample predictive performance of the benchmark and test models is examined. The forecast horizon is chosen to be 10 days. There

[3] G is a sigmoidal function if $G: R \rightarrow [0, 1]$, $G(a) \rightarrow 0$ as $a \rightarrow -\infty$, $G(a) \rightarrow 1$ as $a \rightarrow \infty$ and G is monotonic.

are 22 subsamples in equally spaced intervals for the entire data set, a total of 220 observations for the out-of-sample forecasts. Each forecast horizon corresponds to either a downward, upward or a non-trending market behavior. Therefore, the results across the subsamples reveal evidence for the performance of technical trading rules across various market conditions. Out-of-sample forecasts are completely ex ante by using only the information actually available.

Let $MSPE^{t}$ and $MPSE^{b}$ be the mean square prediction errors of the test and benchmark models, respectively. To measure the out-of-sample performance between the test and benchmark models, the ratio of the mean square prediction errors, $MSPE^{t}/MSPE^{b}$ is used. $MSPE^{t}/MSPE^{b}$ is less than one if the test model provides more accurate predictions. Similarly, the ratio is greater than one if the predictions of the test model are less accurate relative to the benchmark model. In this study, a 10% forecast improvement is adopted as a benchmark measure of performance. The ratios of mean square prediction errors which provide 10% or more forecast improvement are in bold in each table.

The MSPE's of the benchmark model (Eq. (1)), the GARCH-M (1,1) (Eq. (5)) and the feedforward network models (Eq. (6)) with past returns are presented in Table 1. The MSPE's of the benchmark model are reported in levels. MSPE's of the GARCH-M (1,1) and the feedforward network models are reported as a ratio to the MSPE's of the benchmark model. All three specifications are estimated with four lags, $p=4$, of the past returns. The results in Table 1 indicate that the out-of-sample forecast performance of the GARCH-M (1,1) model does not outperform the benchmark model. The ratio of the MSPE's of both models are close to the unitary value except for the 1935–1938 period. The 1935–1938 period provides a 9.3 forecast improvement over the benchmark model.

The findings in Brock et al. (1992) show that the linear conditional mean estimators fail to characterize the temporal dynamics of the security returns and suggest the existence of nonlinearities as the data generation mechanism. Here, the single layer feedforward network regression is used as a nonparametric method to model the conditional mean returns of the stock data. The results of feedforward network regression (Eq. (6)) are presented in the last column of Table 1. Those subperiods which provide a 10% forecast improvement are shown in bold. In the majority of the subperiods, the feedforward network model provides smaller MSPE's in comparison to the benchmark and the GARCH-M (1,1) (Eq. (5)) models. The largest forecast improvement occurs at the 1976–1979 period which is a 12.3 forecast improvement over the benchmark model. Overall, the results of the feedforward network regression with past returns indicate forecast improvement over the benchmark model and the GARCH-M (1,1) (Eq. (5)) specification.

The predictability of the current returns with the past buy–sell signals of the moving average rules are investigated in three different regression models. These test models are OLS (Eq. (2)), GARCH-M (1,1) (Eq. (3)) and feedforward networks (Eq. (4)). The test models are specified such that each test model is nested with the benchmark model. Each of these specifications are estimated for

Table 1
MSPE's of the models with past returns

Date	Obs.	Benchmark	GARCH-M (1,1)	Feedforward
2/1/1963–29/12/1967	1259	[0.1653]	1.013	**0.884**,7
2/1/1968–31/12/1971	983	[0.3117]	1.053	0.988,8
3/1/1972–31/12/1975	1009	[0.3805]	0.991	0.936,8
2/1/1976–31/12/1979	1010	[0.1272]	0.995	**0.877**,8
2/1/1980–30/12/1983	1012	[0.2687]	0.996	1.037,7
3/1/1984–30/06/1988	1137	[1.1276]	0.999	0.920,6
3/1/1939–31/12/1942	1204	[0.2649]	0.998	1.012,6
2/1/1943–31/12/1946	1166	[0.3920]	0.997	0.947,7
2/1/1947–30/12/1950	1129	[0.6372]	1.021	1.069,7
2/1/1951–31/12/1954	1057	[0.3735]	1.016	**0.888**,7
3/1/1955–31/12/1958	1007	[0.3897]	1.005	0.914,8
2/1/1959–31/12/1962	1007	[0.2573]	1.028	1.041,7
2/1/1915–31/12/1918	1198	[0.7973]	1.019	0.970,7
2/1/1919–30/12/1922	1190	[0.3700]	1.008	1.024,8
2/1/1923–31/12/1926	1129	[0.3632]	1.018	0.921,6
3/1/1927–31/12/1930	1185	[3.8557]	1.011	0.903,6
2/1/1931–31/12/1934	1188	[1.1855]	0.983	0.965,7
2/1/1935–31/12/1938	1202	[0.3877]	0.907	0.961,7
2/1/1897–31/12/1900	1191	[1.4790]	0.948	**0.897**,7
2/1/1901–31/12/1904	1189	[1.0559]	0.979	0.911,8
3/1/1905–31/12/1909	1503	[0.0882]	1.006	1.004,8
3/1/1910–31/12/1914	1385	[1.0887]	0.987	0.905,8
Average			0.999	0.953

The linear benchmark model is $r_t = \alpha + \sum_{i=1}^{p} \beta_i r_{t-i} + \epsilon_t$. The GARCH-M (1,1) model is $r_t = \alpha + \sum_{i=1}^{p} \beta_i r_{t-i} + \gamma h_t^{1/2} + \epsilon_t$, $\epsilon_t \sim N(0, h_t)$ and $h_t = \delta_0 + \delta_1 h_{t-1} + \delta_2 \epsilon_{t-1}^2$. The feedforward network model is $r_t = \alpha_0 + \sum_{j=1}^{d} \beta_j G(\alpha_j + \sum_{i=1}^{p} \gamma_{ij} r_{t-i}) + \epsilon_t$. MSPE's are $\times 10^{-4}$. The MSPE's of the benchmark model (Eq. (1)) are reported in levels. The MSPE's of the GARCH-M (1,1) (Eq. (5)) and the feedforward network models (Eq. (6)) are reported as a ratio to the MSPE's of the benchmark model.

two different moving average rules. These rules are $(n_1, n_2) = [(1, 50), (1, 200)]$ where n_1 and n_2 are in days. For convenience, these rules are referred to as $A \equiv (1, 50)$ and $B \equiv (1, 200)$ in the following tables. All three specifications are estimated with four lags, $p = 4$, of the buy–sell signals.

In Table 2, the OLS regression (Eq. (2)) results are presented. For each rule, only the 1927–1930 period provides more than 10% forecast improvement. These are 17.7 and 18.3% for rules A and B, respectively. Also, there is evidence of more predictability for the time period between 1897–1914. The MSPE's of this time period are consistently lower than the benchmark model. The results in Table 2 indicate that the past buy–sell signals in a linear conditional mean specification

Table 2
The ratio of the MSPE's of the OLS model with past buy–sell signals to the MSPE's of the benchmark model

Date	Obs.	Rule *A*	Rule *B*
2/1/1963–29/12/1967	1259	0.981	1.021
2/1/1968–31/12/1971	983	0.952	0.987
3/1/1972–31/12/1975	1009	0.995	1.011
2/1/1976–31/12/1979	1010	0.997	1.012
2/1/1980–30/12/1983	1012	1.001	1.060
3/1/1984–30/06/1988	1137	1.042	0.993
3/1/1939–31/12/1942	1204	0.993	1.004
2/1/1943–31/12/1946	1166	1.003	1.013
2/1/1947–30/12/1950	1129	0.993	1.004
2/1/1951–31/12/1954	1057	1.021	1.002
3/1/1955–31/12/1958	1007	0.965	0.915
2/1/1959–31/12/1962	1007	1.003	0.987
2/1/1915–31/12/1918	1198	0.983	0.989
2/1/1919–30/12/1922	1190	0.995	1.012
2/1/1923–31/12/1926	1129	1.043	1.034
3/1/1927–31/12/1930	1185	**0.823**	**0.817**
2/1/1931–31/12/1934	1188	0.976	0.973
2/1/1935–31/12/1938	1202	1.003	1.011
2/1/1897–31/12/1900	1191	0.969	0.983
2/1/1901–31/12/1904	1189	0.941	0.932
3/1/1905–31/12/1909	1503	0.927	0.935
3/1/1910–31/12/1914	1385	0.981	0.987
Average		0.981	0.986

The linear benchmark model is $r_t = \alpha + \sum_{i=1}^{p} \beta_i r_{t-i} + \epsilon_t$. The linear model with buy–sell signals is $r_t = \alpha + \sum_{i=1}^{p} \beta_i r_{t-i} + \infty_{i=1}^{p} \eta_i s_{t-i}^{n_1,n_2} + \epsilon_t$. Rules A and B refer to the $(n_1, n_2) \equiv [(1, 50), (1, 200)]$ where n_1 and n_2 are in days.

provide a slight forecast improvement in comparison to the benchmark parametric model with past returns.

In Table 3, the results of the GARCH-M (1,1) model (Eq. (3)) are presented. The GARCH-M (1,1) model makes an improvement over the OLS model. In general, the MSPE's of the GARCH-M (1,1) model are smaller than the OLS model in Eq. (2). The average MSPE's of the GARCH-M (1,1) model for rules *A* and *B* are 3.2 and 2.7% smaller than the MSPE's of the benchmark model. For the 1897–1914 period, the GARCH-M (1,1) model provides 18.9 and 18.3% forecast improvement over the benchmark model for rules *A* and *B*, respectively.

In Table 4, the feedforward network regression (Eq. (4)) results with past returns and past buy–sell signals are presented. The results indicate that 14 of the subperiods provide at least 10% forecast improvement in comparison to the

Table 3
The ratio of the MSPE's of the GARCH-M (1,1) model with past buy–sell signals to the MSPE's of the benchmark model

Date	Obs.	Rule *A*	Rule *B*
2/1/1963–29/12/1967	1259	0.943	0.989
2/1/1968–31/12/1971	983	0.976	0.987
3/1/1972–31/12/1975	1009	0.998	0.999
2/1/1976–31/12/1979	1010	0.979	0.984
2/1/1980–30/12/1983	1012	1.004	1.010
3/1/1984–30/06/1988	1137	0.991	0.993
3/1/1939–31/12/1942	1204	0.995	0.996
2/1/1943–31/12/1946	1166	0.987	0.997
2/1/1947–30/12/1950	1129	0.988	0.997
2/1/1951–31/12/1954	1057	1.003	1.011
3/1/1955–31/12/1958	1007	0.939	0.896
2/1/1959–31/12/1962	1007	0.976	0.991
2/1/1915–31/12/1918	1198	0.981	0.989
2/1/1919–30/12/1922	1190	0.995	0.996
2/1/1923–31/12/1926	1129	1.001	1.008
3/1/1927–31/12/1930	1185	**0.811**	**0.817**
2/1/1931–31/12/1934	1188	0.964	0.948
2/1/1935–31/12/1938	1202	1.004	1.012
2/1/1897–31/12/1900	1191	0.921	0.933
2/1/1901–31/12/1904	1189	0.935	0.936
3/1/1905–31/12/1909	1503	0.925	0.932
3/1/1910–31/12/1914	1385	0.974	0.987
Average		0.968	0.973

The linear benchmark model is $r_t = \alpha + \sum_{i=1}^{p} \beta_i r_{t-i} + \epsilon_t$. The GARCH-M (1,1) model with buy–sell signals is $r_t = \alpha + \sum_{i=1}^{p} \beta_i r_{t-i} + \sum_{i=1}^{p} \eta_i s_{t-i}^{n_1,n_2} + \gamma h_t^{1/2} + \epsilon_t$ where $\epsilon_t \sim N(0,\ h_t)$ and $h_t = \delta_0 + \delta_1 h_{t-1} + \delta_2 \epsilon_{t-1}^2$. Rules A and B refer to the $(n_1,\ n_2) \equiv [(1, 50), (1, 200)]$ where n_1 and n_2 are in days.

benchmark model with past returns. Especially for the 1915–1918 subperiod, the improvement in the MSPE of the feedforward network model is as large as 27.9%. The forecast improvements of the feedforward network regression with the past buy–sell signals are substantial and dominate all other specifications including the feedforward network regression with past returns. Between the two moving average rules, rule *A* provides more accurate out-of-sample predictions relative to rule *B*. This may be due to the fact that rule *B* oversmooths the data.

The analysis above uses the continuous buy–sell signals as regressors. To compare the performance of the discrete buy–sell signals as predictors of the security returns, an indicator variable, $d_t^{n_1,n_2}$ is formed. Let

$$d_t^{n_1,n_2} = \begin{cases} 1 & (m_t^{n_1} - m_t^{n_2}) > 0 \\ 0 & \text{otherwise} \end{cases} \tag{7}$$

Table 4
The ratio of the MSPE's of the feedforward model with past buy–sell signals to the MSPE's of the benchmark model

Date	Obs.	Rule *A*	Rule *B*
2/1/1963–29/12/1967	1259	**0.856,7**	**0.871**,7
2/1/1968–31/12/1971	983	0.934,7	**0.853**,7
3/1/1972–31/12/1975	1009	**0.854**,8	0.904,9
2/1/1976–31/12/1979	1010	**0.831**,9	**0.856**,7
2/1/1980–30/12/1983	1012	**0.854**,8	**0.827**,8
3/1/1984–30/06/1988	1137	**0.894**,8	**0.876**,8
3/1/1939–31/12/1942	1204	0.994,7	1.012,8
2/1/1943–31/12/1946	1166	**0.862**,7	0.996,8
2/1/1947–30/12/1950	1129	0.967,8	0.966,9
2/1/1951–31/12/1954	1057	**0.885**,9	0.900,8
3/1/1955–31/12/1958	1007	0.927,7	0.939,8
2/1/1959–31/12/1962	1007	**0.835**,9	**0.866**,7
2/1/1915–31/12/1918	1198	**0.721**,7	**0.869**,7
2/1/1919–30/12/1922	1190	0.951,8	0.936,7
2/1/1923–31/12/1926	1129	0.976,7	0.955,8
3/1/1927–31/12/1930	1185	**0.863**,9	**0.833**,8
2/1/1931–31/12/1934	1188	**0.891**,9	**0.881**,8
2/1/1935–31/12/1938	1202	**0.815**,8	**0.832**,8
2/1/1897–31/12/1900	1191	**0.786**,8	**0.882**,8
2/1/1901–31/12/1904	1189	**0.879**,7	**0.885**,8
3/1/1905–31/12/1909	1503	0.927,8	**0.851**,8
3/1/1910–31/12/1914	1385	0.972,8	0.984,7
Average		0.885	0.899

The linear benchmark model is $r_t = \alpha + \sum_{i=1}^{p} \beta_i r_{t-i} + \epsilon_t$. The feedforward network with buy–sell signals is $r_t = \alpha_0 + \sum_{i=1}^{p} \beta_{ij} r_{t-i} + \sum_{j=1}^{d} \eta_j G(\alpha_j + \sum_{i=1}^{p} \gamma_{ij} s_{t-i}^{n_1,n_2}) + \epsilon_t$. Rules A and B refer to the $(n_1, n_2) \equiv [(1, 50), (1, 200)]$ where n_1 and n_2 are in days. The first entry under each rule is the ratio of MSPE's. The second entry is the number of hidden units of the feedforward network model indicated by the cross-validation technique.

where n_1 and n_2 are the short and long moving averages, respectively. The linear test regression is then written as

$$r_t = \alpha + \sum_{i=1}^{p} \beta_i r_{t-i} + \sum_{i=1}^{p} \eta_i d_{t-i}^{n_1,n_2} + \epsilon_t \qquad (8)$$

where $\epsilon_t \sim ID(0, \sigma_t^2)$. The predictive performance of the linear models with discrete buy–sell signals is within one percent of the predictive performance of the model in Eq. (2). This suggests that the predictive performance of discrete buy–sell signals are similar to that of the continuous buy–sell signals.

The results of this paper confirm the findings of Brock et al. (1992) that the linear conditional mean estimators may not successfully characterize the temporal dynamics of the security returns. The results show that the use of the past buy–sell signals in the linear conditional mean estimations provides only a slight improvement in the forecast performances relative to the benchmark model. The main findings of this paper indicate that the nonlinearities in the stock return series play an important role in modelling the conditional mean of the security returns. The forecast gains originate from the utilization of the past buy–sell signals as inputs in the nonparametric conditional mean specifications. In the feedforward network model with past buy–sell signals, the forecast gains over the benchmark model are much more pronounced.

5. Conclusions

This paper has used the daily Dow Jones Industrial Average Index from 1897 to 1988 to examine the linear and nonlinear predictability of stock market returns with simple technical trading rules. Evidence of nonlinear predictability in stock market returns is found by using the past buy and sell signals of the moving average rules. The forecast gains originate from the utilization of the past buy–sell signals as inputs in the feedforward network specifications. The evidence across subsamples indicates that the two moving average rules studied here provide at least a 10% forecast improvement in the volative years of the Great Depression and in the trendy years of 1980–1988. The performance of these rules is more moderate in the 1939–1950 period in which there is no clear trend in either a positive or negative direction.

The technical trading rules used in this paper are very popular and very simple. The results here suggest that it is worthwhile to investigate more elaborate rules and the profitability of these rules after accounting for transaction costs and brokerage fees.

Acknowledgements

I thank the editor, the associate editor and two anonymous referees for helpful comments which improved the results in the paper. I also thank Buz Brock, Blake LeBaron, Chung-Ming Kuan, Tung Liu for discussions and Blake LeBaron for providing the data; Xian Yang for providing excellent research assistance; and the Natural Sciences and Engineering Research Council of Canada and the Social Sciences and Humanities Research Council of Canada for financial support.

References

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

Chopra, N., Lakonishok, J., Ritter, J.R., 1992. Performance measurement methodology and the question of whether stocks overreact. Journal of Financial Economics 31, 235–268.

Conrad, J., Kaul, G., 1988. Time-variation in expected returns. Journal of Business 61, 409–425.

Cutler, D.M., Poterba, J.M., Summers, L.H., 1991. Speculative dynamics. Review of Economic Studies 58, 529–546.

Cybenko, G., 1989. Approximation by superposition of a sigmoidal function. Mathematics of Control, Signals and Systems 2, 303–314.

De Bondt, W.F.M., Thaler, R.H., 1985. Does the stock market overreact. Journal of Finance 40, 793–805.

Fama, E.F., French, K.R., 1986. Permanent and temporary components of stock prices. Journal of Political Economy 98, 246–274.

French, K.R., Roll, R., 1986. Stock return variances: The arrival of information and the reaction of traders. Journal of Financial Economics 17, 5–26.

Funahashi, K.-I., 1989. On the approximate realization of continuous mappings by neural networks. Neural Networks 2, 183–192.

Gallant, A.R., White, H., 1988. There exists a neural network that does not make avoidable mistakes. Proceedings of the Second Annual IEEE Conference on Neural Networks, San Diego, CA. IEEE Press, New York, pp. I.657–I.664.

Gallant, A.R., White, H., 1992. On learning the derivatives of an unknown mapping with multilayer feedforward networks. Neural Networks 5, 129–138.

Hecht-Nielsen, R., 1989. Theory of the backpropagation neural networks. Proceedings of the International Joint Conference on Neural Networks, Washington, DC. IEEE Press, New York, pp. I.593–I.605.

Hornik, K., Stinchcombe, M., White, H., 1989. Multilayer feedforward networks are universal approximators. Neural Networks 2, 359–366.

Hornik, K., Stinchcombe, M., White, H., 1990. Universal approximation of an unknown mapping and its derivatives using multilayer feedforward networks. Neural Networks 3, 551–560.

Jegadeesh, N., 1990. Evidence of predictable behavior of security returns. Journal of Finance 45, 881–898.

Kuan, C.-M., White, H., 1994. Artificial neural networks: An econometric perspective. Econometric Reviews 13, 1–91.

Lehmann, B.N., 1990. Fads, martingales and market efficiency. Quarterly Journal of Economics 105, 1–28.

Lo, A.W., MacKinlay, A.C., 1988. Stock market prices do not follow random walks: Evidence from a simple specification test. Review of Financial Studies 1, 41–66.

Lo, A.W., MacKinlay, A.C., 1990. When are contrarian profits due to stock market overreaction?. Review of Financial Studies 3, 175–205.

Poterba, J.M., Summers, L.H., 1988. Mean reversion in stock prices: Evidence and implications. Journal of Financial Economics 22, 27–59.