

**“DIRECTION-OF-CHANGE FORECASTING USING A VOLATILITY-  
BASED RECURRENT NEURAL NETWORK”**

Stelios Bekiros\*

*Center for Nonlinear Dynamics in Economics and Finance, Faculty of Economics  
University of Amsterdam, Roetersstraat 11, 1018 WB Amsterdam, The Netherlands*

Dimitris Georgoutsos†

*Department of Accounting and Finance, Athens University of Economics and  
Business, 76 Patission str, GR104 34, Athens, Greece*

November 2006

CeNDEF working paper 06-16

---

\* (Corresponding author); Tel.: + 31 20 525 5375; fax: +31 20 525 5283; E-mail address:  
S.Bekiros@uva.nl

†Tel.: +30 210 8203441; fax: +30 210 8214122; E-mail address: dgeorg@aueb.gr.

### *Abstract*

This paper investigates the profitability of a trading strategy, based on recurrent neural networks, that attempts to predict the direction-of-change of the market in the case of the NASDAQ composite index. The sample extends over the period 2/8/1971 – 4/7/1998, while the sub-period 4/8/1998 – 2/5/ 2002 has been reserved for out-of-sample testing purposes. We demonstrate that the incorporation in the trading rule of estimates of the conditional volatility changes strongly enhances its profitability during “bear” market periods. This improvement is being measured with respect to a nested model that does not include the volatility variable as well as to a buy & hold strategy. We suggest that our findings can be justified by invoking either the “volatility feedback” theory or the existence of portfolio insurance schemes in the equity markets. Our results are also consistent with the view that volatility dependence produces sign dependence.

*JEL classification:* G10; G14; C53

*Keywords:* Technical trading rules; Neural networks; Volatility trading

## 1. Introduction

In the present paper we explore the predictive return sign ability of trading rules that rely on a simple switching strategy: positive predicted returns are executed as long positions and negative returns as short positions. A similar strategy has been employed, with considerable success, by a number of other researchers. Gençay (1998b) examines the profitability of a simple trading rule, applied on the DJIA index, where signs are modeled as a function of the past returns and are estimated by a feedforward network, a class of artificial neural networks (ANN). The results of this simple model indicate that nonparametric models with technical rules provide excess returns when compared to a simple buy-and-hold strategy. Fernández-Rodríguez *et al.* (2000) conduct a similar exercise for the Madrid stock market general index and show that a simple trading rule based on ANNs is always superior to a buy-and-hold strategy during “bear” market conditions. Pesaran and Timmermann (1994) examine whether the predictability of the Standard and Poor’s 500-index returns could have been historically exploited by investors to earn profits in excess of a buy-and-hold strategy. In general terms they find that the returns from the switching strategy are higher than those from the passive one for annual returns, even when transactions costs are high. They also find that the predictive power of various economic factors is increased during volatile periods.<sup>1</sup>

The present paper advances the existing literature by exploring the predictive ability of trading rules that incorporate, among others, forecasts of the conditional volatility changes over the next trading period. The empirical investigation of the relation between stock return volatility and stock returns has a long tradition in finance (Bekaert and Wu, 2000) According to the “time-varying risk premium theory” the return shocks are caused by changes in conditional volatility. When news arrives

in the market the current volatility increases and this causes upward revisions of the conditional volatility since there is a well-documented fact that volatility is persistent. This increased conditional volatility has to be compensated by a higher expected return, leading to an immediate decline in the current value of the market. So in the case of bad news the volatility feedback effect reinforces the initial drop in stock market prices. However when good news arrives in the market and volatility increases, prices decline to induce higher expected returns offsetting thus the initial price movement.<sup>2</sup> An alternative rationalization for the presence of conditional volatility revisions in the trading rule may be offered by invoking trigger strategies in the equity markets (Krugman 1987). Participants in portfolio insurance schemes react whenever the maximum expected loss, as measured for example by the Value-at-Risk (VaR), reaches a predetermined level and therefore share price dynamics are being driven, partly, by revisions in the measured conditional volatility.<sup>3</sup> If we assume a continuity of portfolios that deviate to a varying degree from their pre-determined level of VaRs then each time the conditional volatility rises, a number of those portfolios will hit their risk limits and this will generate a re-allocation of assets towards safer ones. Each time portfolio insurers leave the market the stock prices must fall in order for the other investors to be given an incentive to hold a larger quantity of stock. If we further assume a rational expectations world then investors take into account the effects of portfolio insurance schemes and no step drop in stock prices is being observed. In an intriguing recent paper Christoffersen and Diebold (2003) show that volatility dependence produces sign dependence, and therefore forecastability, as long as expected returns are nonzero. The intuition behind this relationship is that volatility changes will alter the probability of observing negative or positive returns. More specifically, the higher the volatility, the higher the

probability of a negative return, as long as the expected returns are positive. Moreover, they show that this result is entirely consistent with the existence of no conditional mean dependence, or the absence of conditionally Gaussian distributions.

Another branch of the literature studies the contemporaneous relationship between the one-day stock index returns and the associated changes in the level of implied volatility indexes. The results indicate the existence of a negative and statistically significant relationship between the returns of the S&P100 (Nasdaq 100) and the implied volatility VIX (VIX) index (Whaley (2000), Simon (2003), Giot (2000)). For the S&P100 index this relationship has been also found to be asymmetric in the sense that negative stock index returns are associated with greater proportional changes in implied volatility measures than are positive returns. The explanation offered for this opposite response is that option traders react to negative returns by bidding up the implied volatility.

Although the contemporaneous negative association between the returns of the volatility indices and the corresponding equity indices is well documented empirically there is a growing debate whether the implied volatility can be used as a forward indicator of the underlying equity index. This issue has not been treated properly in the literature with the exception a paper by Giot (2005) who regressed the forward looking S&P100 index returns, over various time intervals, on 21 dummy variables representing equally spaced percentiles of a rolling two-year history of VIX. Giot (2005) concludes that positive forward-looking returns are to be expected for long positions at high levels of the implied volatility indexes. In our paper we examine the trading implications of conditional volatility changes within a broader framework as concerns the functional form of the forecast generating mechanism as well as the presence of past returns that might have forecasting power.

In the next section we discuss the construction of the trading rule and the way this is incorporated into an ANN as well as the estimation techniques that have been applied. We then proceed with the presentation of the statistical and financial criteria that have been adopted to evaluate the forecasting ability of the various models. Finally, we comment on the empirical results we obtained from the empirical analysis.

## 2. Predicting Stock Index Returns with Neural Networks

Recent research into the time series properties of stock market indices returns by Pesaran and Timmermann (1994) and Abhyankar, Copeland and Wong (1997), to name a few, have indicated the presence of non-linear dynamics. Neural networks use a non-parametric method of forecasting which means that the underlying non-linear function is not prescribed, ex-ante, explicitly. Thus, the model is not limited to a restrictive list of non-linear functions.<sup>4</sup> In financial applications the most popular class of ANN models has been the single-layer feedforward networks (FNN). In a FNN, information suitably weighted is passed from the point of entry (the input layer) to a further layer of hidden neurons. This hidden information is also assigned a weight and finally reaches the output layer that represents the forecast.

Let  $p_t, t = 1, 2, \dots, T$  be the daily stock index price. The daily returns are then calculated by  $r_t = \log(p_t) - \log(p_{t-1})$ . The output,  $y$ , of a single layer FNN is then given by:

$$y_t = S \left[ a_0 + \sum_{i=1}^q a_i \cdot g_{t,i} \right] \quad (1)$$

where:

$$g_{t,i} = G \left( b_{i0} + \sum_{j=1}^n b_{ij} r_{t-j} + \sum_{k=1}^p c_{ik} \Delta h_{t-k}^{1/2} \right), \quad i = 1, \dots, q. \quad (2)$$

The inputs in the suggested FNN correspond to the daily returns over the previous  $n$  days, following Gençay (1998a, b) and Fernandez-Rodriguez *et.al.* (2000), and the daily revisions of the estimated conditional volatility,  $\Delta h^{1/2}$ , over the past  $p$  days.<sup>5</sup> As concerns the transfer functions  $G$  and  $S$  we use the tansig function. The tan-sigmoid function normalizes the values of each neuron to be in the interval  $(-1, +1)$ . The problem we are faced with is that the FNN is given the correct weights such that  $y$  has the correct value corresponding to the inputs. This is being accomplished by the *error backpropagation* method under which the neural network runs through all the input data over an initial “training” period and produces a list of outputs. Then the weights are reevaluated, by using a recursive “gradient” descent method, so that the mean-squared error between the observed output and the predicted one is minimized.<sup>6</sup> Once the neural network has been trained, it is applied over a different data set covering the so-called “validation” period. The purpose here is to evaluate the generalization ability of a supposedly trained network in order to avoid overfitting.

In a dynamic context it is natural to include lagged dependent variables as explanatory variables in the FNN in order to capture dynamics. This problem is being addressed in the relevant literature by constructing recurrent networks, i.e. networks with feedbacks from the hidden neurons, to the input layer with delay. The recurrent neural networks (RNN) memorize thus information since its output depends on both current and prior neuron inputs. In this paper we apply the Elman (1990) RNN with a single hidden layer and feedback connection from the output of the hidden layer to its input. In a RNN model, equation (2) can be re-written as:

$$g_{t,i} = G \left( b_{i0} + \sum_{j=1}^n b_{ij} r_{t-j} + \sum_{k=1}^p c_{ik} \Delta h_{t-k}^{1/2} + \sum_{l=1}^m \delta_{il} g_{l,t-1} \right), \quad i = 1, \dots, q. \quad (3)$$

It is easy to show, with back-substitution, that the output  $y_t$  depends on the entire history of the inputs  $r$  and  $\Delta h^{1/2}$ .

The trading rule over the testing period works as follows. At the end of each trading day the RNN is being re-estimated over a rolling sample that is equal to the training period set. The output unit, eq. (1), receives the weighted sum of the signals, from eq. (3) and produces a signal through the output transfer function ( $S$ ). If the value of the signal is greater than zero it is interpreted as a “buy” signal for the next trading day while a value less than zero as a “sell” signal. Then, the total return of the strategy, when transaction costs are not considered, is estimated as:

$$R_0^N = \sum_{t=0}^N \hat{y}_t r_t, \quad (4)$$

where  $\hat{y}_t$  is the recommended position which takes the value of (-1) for a short position and (+1) for a long position (e.g., Gençay, 1998b and Fernandez-Rodriguez *et al.*, 2000).

### **3. The market timing and investment performance of alternative trading rules**

We have estimated the RNN model of equations (1) and (3) on daily returns of the Nasdaq composite index that span the period 2/8/1971 to 2/5/2002.<sup>7</sup> The Nasdaq 100 composite index measures all Nasdaq and international based common type stocks listed on the Nasdaq Stock Market. Today includes more than 3000 companies and for that reason is one of the most widely followed and quoted major market indices.

*Insert Figure 1 here*



The testing, out-of-sample, period has been split into two subperiods; a “bull” market period from 4/2/1998 to 3/12/2000 and a “bear” market period from 3/1/2000 to 2/5/2002. The training and validation period account for the rest of the sample with the validation period covering almost 30% of the entire data set.

In order to rationalize the use of neural network models we have tested for the presence of non-linear dependence in the series. To that end, we have made use of the well known BDS test statistic which under the null of i.i.d. is given by (Brock et.al., 1991):

$$W_{m,T(\varepsilon)} = T^{1/2}[C_{m,T(\varepsilon)} - C_{1,T(\varepsilon)}] / \sigma_{m,T(\varepsilon)}. \quad (5)$$

$C_{m,T(\varepsilon)}$  is the correlation integral from  $m$  dimensional vectors that are within a distance  $\varepsilon$  from each other, when the total sample is  $T$ , and  $\sigma_{m,T(\varepsilon)}$  is the standard deviation of  $C_{m,T(\varepsilon)}$ . Under the null hypothesis,  $W_{m,T(\varepsilon)}$ , has a limiting standard normal distribution. The BDS test has been applied on: (a) the original data, (b) the residuals from an autoregressive filter, in order to ensure that the null is not rejected due to linear dependence, and (c) the natural logarithm of the squared standardized residuals from a GARCH-M (1,1) model, in order to ensure that rejection of the null is not due to conditional heteroscedasticity (De Lima (1996)).

*Insert Table 1 here*

In all three cases we were unable to accept the null of i.i.d. at the 1% marginal significance level and the evidence seems to suggest that a genuine non-linear dependence is present in the data.

The results relating to the predictability as well as the profitability of the Elman network we estimated appear in Table 2. They correspond to a specification where two lags of the returns and one lag of the conditional volatility changes appear in equation (3), ( $n=2$  and  $p=1$ ).<sup>8</sup> In addition there is one hidden layer with ten neurons ( $g$ ) and one output layer with a single neuron ( $y$ ). Conditional volatility estimates,  $h_t^{1/2}$ , have been obtained from: a rolling 20-day standard deviation of returns; an exponentially weighted moving average (EWMA) with a decay factor equal to 0.94;<sup>9</sup> a GARCH (1,1) model; and a Glosten, Jagannathan and Runkle (1993) (GJR) GARCH (1,1) model that allows for an asymmetric response of volatility to positive or negative shocks. The choice of these estimating techniques is rationalized on the basis of findings in a number of studies that show that the forecasting power of conditional volatility models is higher than the one obtained from implied volatility indices. For example Simon (2003) reports, for the Nasdaq 100 index over the period 1995-2002, that the mean absolute forecast error of the VIX is about a third greater than those of both the out-of-sample GJR GARCH (1,1) volatility forecasts and the EWMA, with the decay factor equal to 0.94, volatility forecasts.

*Insert Table 2 here*

The adequacy of the chosen specification, without the presence of the volatility changes, is considered as satisfactory. As the first column of Table 3 shows the total return of the trading strategy is 29.2% for the entire testing period when a buy-and-hold (B&H) policy would have earned only 4.5%. Moreover, the proportion of the correctly predicted signs is above 50% and this is reflected in a significant value, at the 5% level, of the Henriksson- Merton (1981) (HM) test statistic. Finally, the Sharpe ratio (SR) and the Ideal profit (IP) index are both positive, although rather

small in value. As concerns the two testing sub-periods, the chosen strategy behaves better during the bull period, according to the Pesaran-Timmerman (1992) (PT) and HM tests as well as the SR and IP indices. However, the overall return compared to the B&H policy is superior during the “bear” market sub-period. This image accords with previous results derived by Fernández *et. al.* (1999, 2000) from a similar model applied on the Nikkei and the Madrid stock market general indices.

Next, we evaluate the trading strategy with the conditional volatility variable included. The first evidence that emerges from this change is that the returns improve significantly, over the “bear” market period, independently of the model we used to produce the conditional volatility estimates. Over, the “bull” market period the strategy does not seem capable of succeeding the profits of the “no volatility” strategy while it is always worse than the B&H policy. Similarly, the other performance indices, PT, HM, SR, PI, and the sign rate show an improvement over the “no volatility” case for the “bear” market period. The significant increase in the profitability of the suggested trading rule may be compromised with the marginal improvement of the sign rate index by the substantial improvement of the quantitative importance of the correctly forecasted signs. Those same indices for the “bull” market period produce values that are similar or worse than those under the “no volatility” specification. The comparison between the four different specifications for the volatility estimation show that simple models of historical volatility measurement, like the equally weighted and the exponentially weighted moving averages, produce sign forecasts that are no worse than those obtained from more complicated econometric models that are often used to model conditional volatility.<sup>10</sup> The results we obtain are in accordance with the conclusions reached by Christoffersen and Diebold (2003). The profitability of our trading rule as well as the sign rate indicator

improve substantially over the “bear” market period as implied by the aforementioned paper for a higher volatility period that characterize the post-bubble time interval in our test. Moreover, this improvement over one of the two testing periods is not depicted strongly into the results of the HM and PT tests since they have no power to detect sign dependence in the face of non zero expected returns (Christoffersen and Diebold, 2003).

#### **4. Conclusions**

In the present paper we expand the literature that evaluates the return sign forecasting ability of trading rules based on neural networks over simple alternative strategies like a Buy & Hold. A B&H policy cannot be consistently outperformed from any trading rule, no matter how elaborate this is, in a random walk market. We first replicate previous evidence coming from other stock market indices, according to which the forecasting ability of simple rules outperform the B&H profits over “bear” market conditions although the evidence from various profitability indices is positive for the “bull” period as well. Then we included in the trading rule revisions of the conditional volatility of the Nasdaq index that have been produced from alternative estimating techniques. This change generated a substantial improvement of the profits and the profitability per unit of risk over the “bear” market period. These results seem to indicate that the neural network has been “trained” to relate correctly changes in conditional volatility with the “sign” of the market one day ahead. This may be attributed to a number of reasons. The first associates increases in volatility to higher expected returns. In the case when increases in volatility are generated from “bad” news we will experience lower prices the next trading day. However, when increases in volatility are generated from “good” news it is not clear what the net effect on

prices will be. This explanation seems to accord with the enhanced predictability of our model, which incorporates volatility revisions, over the “bear” market period when “bad” news dominate. The second reason associates increases in volatility with trigger strategies followed by many portfolio managers. Every time volatility rises, the risk limit is being hit for some portfolios and then liquidation follows. This puts a pressure on the market that is more severe during “bear” market conditions. Finally, our results are in broad accordance to the conclusions reached from a “statistical” perspective according to which there is a close relationship between asset return signs and asset return volatilities.

## References

- Abhyankar, A., Copeland, L. S. and Wong, W., 1997, Uncovering Nonlinear Structure in Real-time Stock Market Indexes: The S&P 500, the DAX, the Nikkei 225, and the FTSE-100, *Journal of Business & Economic Statistics* 15, 1-14.
- Bekaert, G. and G. Wu, 2000, Asymmetric Volatility and Risk in Equity Markets, *The Review of Financial Studies* 13, No 1, 1-42.
- Brock, W., D. Hsieh and B. LeBaron, 1991, *Nonlinear Dynamics, Chaos and Instability*, The MIT Press.
- Cheng, B., and D.M. Titterington, 1994, Neural networks: a review from a statistical perspective, *Statistical Science* 9, 2-54.
- Christie, A.A., 1982, The Stochastic Behavior of Common Stock Variances-Value, Leverage and Interest Rate Effects, *Journal of Financial Economics* 10, 407-432.
- Christoffersen, P., and F. Diebold, 2003, Financial Asset Returns, Direction-of-change Forecasting, and Volatility Dynamics, working paper 10009, NBER.
- Cybenko, G., 1989, Approximations by superpositions of a sigmoidal function, *Mathematics of Control, Signals and Systems* 2, 203-230.
- De Lima, P.J.F., 1996, Nuisance parameter free properties of correlation integral based statistics, *Econometric Reviews* 15, 237-259.
- Elman, J. L. 1990, Finding structure in time, *Cognitive Science* 14, 179-211.
- Fernández-Rodríguez, F., S. Sosvilla-Rivero, and M. D. Garca-Artiles, 1999, Dancing with Bulls and Bears: Nearest-neighbour forecasts for the Nikkei index, Japan and the World Economy 11, 395-413.
- Fernández-Rodríguez, F., C. Gonzalez-Martel and S. Sosvilla-Rivero, 2000, On the Profitability of Technical Trading Rules based on Artificial Neural Networks: Evidence from the Madrid Stock Market, *Economic Letters* 69, 89-94.

Gençay, R., 1998a, The Predictability of Security Returns with Simple Technical Trading Rules, *Journal of Empirical Finance* 5, 347-359.

\_\_\_\_\_, 1998b, Optimization of Technical Strategies and the Profitability in Security Markets, *Economic Letters* 59, 249-254.

\_\_\_\_\_ and T. Stengos, 1998c, Moving Average Rules, Volume and the Predictability of Security Returns with Feedforward Networks, *Journal of Forecasting* 17, 401-414.

Giot, P., 2005, Relationships Between Implied Volatility Indexes and Stock Index Returns, *The Journal of Portfolio Management*, 92-100.

Glosten, L., R. Jaganathan and D. Runkle, 1993, On the Relation between the Expected Value and the Volatility of the Nominal Excess Return on Stocks, *The Journal of Finance* 48, 1779-1801.

Henriksson, R.D. and R.C. Merton, 1981. On the Market timing and Investment Performance II: Statistical procedures for evaluating forecasting skills, *Journal of Business* 54, 513-533.

Hornik, K., M. Stinchcombe and H. White, 1989. Multi-layer feedforward networks are universal approximations, *Neural Networks* 2, 359-366.

Krugman, P., 1987, Trigger Strategies and Price Dynamics in Equity and Foreign Exchange Markets, NBER Working Paper No 2459.

Pesaran, M. H., and A. Timmermann, 1992, A simple non-parametric test of predictive performance, *Journal of Business and Economic Statistics* 10, 461-465.

\_\_\_\_\_, 1994, Forecasting Stock Returns. An Examination of Stock Market Trading in the Presence of Transaction Costs, *Journal of Forecasting* 10, 335-367.

\_\_\_\_\_, 1995, Predictability of Stock Returns: Robustness and Economic Significance, *The Journal of Finance* 50, 1201-1228.

Schwert, W., 2002. Stock Volatility in the new millennium: how wacky is Nasdaq?, *Journal of Monetary Economics* 49, 3-26.

Simon, David, 2003, The Nasdaq Volatility Index During and After the Bubble, *The Journal of Derivatives*, 9-23.

Whaley, Robert, 2000, The Investor fear gauge, *The Journal of Portfolio Management* 26, 12-27.

White, H., 1992. *Artificial Neural Networks: Approximation and Learning Theory*. Cambridge, MA, Blackwell Publishers.



## Endnotes

<sup>1</sup> The above-mentioned literature is part of a more extensive one, on asset return predictability, that incorporates the buy and sell signals from simple technical trading strategies into an ANN specification. In Gençay (1998a), moving average rules in an ANN provide a forecast improvement, as measured by the mean square prediction errors (MSPEs), when they are compared to the predicted Dow Jones index daily returns from a linear regression or a GARCH-M(1,1) process. Gençay and Stengos (1998) extend the previous work by incorporating a volume average rule in their trading strategy.

<sup>2</sup> An asymmetric nature of the volatility response to return shocks emerges from the above discussion. Bad news generates an increase in conditional volatility while the net impact of good news is not clear. An alternative explanation to the asymmetric reaction of the conditional volatility may be offered through the “leverage effects” (e.g., Christie 1982). A negative (positive) return increases (reduces) financial leverage, which makes the stock riskier (less risky) and increases (reduces) volatility. The causality however here is different: the return shocks lead to changes in conditional volatility, whereas the time-varying premium theory contends the opposite (e.g., Bekaert and Wu, 2000).

<sup>3</sup> The VaR depends entirely on a multiple of the estimated conditional volatility under the assumption of normally distributed returns.

<sup>4</sup> Cybenko (1989) and Hornik *et.al.* (1989) have demonstrated that ANN models can approximate, under certain regularity conditions, any continuous function. This unveils the main weakness of the ANNs since they may end up fitting the noise in the data rather than the underlying statistical process. Cheng and Titterton (1994) have shown that ANNs are equivalent to non-linear non-parametric models while they claim that most forecasting models (ARMA, autoregressive with thresholds, non-parametric with kernel regression, etc.) can be written in the form of a network of neurons.

<sup>5</sup> Whaley (2000) has used a similar approach where revisions of the implied volatility index, VIX, are significantly related, in an asymmetric way, to the S&P 500 index returns.

<sup>6</sup> As White (1992) has shown the existence of such weights is guaranteed since any non-linear function can be approximated as above, with a single layer, to an arbitrary degree of accuracy with a suitable number of neurons.


<sup>7</sup> Although the test statistics are based on nominal stock index returns, similar results would obtain with excess returns since the volatility of daily nominal returns is so much larger than that of Treasury-bill rates.

<sup>8</sup> The procedure for the selection of the lags involved the estimation of autoregressive (AR) models and the calculation of the Ljung-Box statistics for the first 16 lags of the series. Significant autocorrelations of up to the second lag of the return series were identified. Additionally, the Akaike Information Criterion (AIC) that was estimated for the first six lags provided the minimum value at the second lag. As concerns the conditional volatility variable, sensitivity analyses for different number of lags were conducted on the RNN but the results were not found to be qualitatively different from those presented in Table 2. Similar exercises were conducted for a different number of lagged returns but again the results we obtained are not better than those shown in Table 2. The results of the sensitivity analyses are available upon request.

<sup>9</sup> The exponentially moving average corresponds to the approach adopted by *RiskMetrics* and for that reason it is denoted here as RM (0.94).

<sup>10</sup> This has not been surprising since it is documented that forecasts of volatility, for the NASDAQ index, from MA rules closely approximate those from GARCH (1,1) models (e.g., Schwert, 2002). Simon (2003) also reports that the GJR (1993) GARCH volatility forecasts of the Nasdaq 100 average 3.0 percentage points higher than the actual when the EWMA volatility forecasts are only 1.5 percentage points below actual volatility.

**Figure 1:** Daily closing prices and historical volatility, annualized, of the Nasdaq composite Index based on a rolling 21-day standard deviation (02/08/1971 – 02/05/2002)



**Table 1: BDS test**

Series	$m=2$		$m=3$		$m=4$	
	$\varepsilon = 1$	$\varepsilon = 1.5$	$\varepsilon = 1$	$\varepsilon = 1.5$	$\varepsilon = 1$	$\varepsilon = 1.5$
O.D.	5.18*	4.93*	8.98*	8.24*	10.74*	9.54*
RAF	4.74*	4.53*	8.30*	7.80*	9.87*	9.05*
NLSNR	3.15*	3.64*	3.97*	4.22*	4.20*	4.19*

Notes: O.D. = original data (daily returns of the Nasdaq index),

RAF = residuals from an autoregressive filter,

NLSNR = natural logarithm on standardized normalized residuals.

$m$  = the value of the dimension,  $\varepsilon$  = the number of standard deviations of the data.

Brock *et. al.*, (1991) suggest that the standardized normal distribution is a good approximation of the finite sample distribution for a sample of 500 or more observations, values of the dimension  $m$  below 5 and values of the distance  $\varepsilon$  between 0.5 and 2 standard deviations of the data.

(\*) indicates significance at the 1% significance level (the critical value is 2.58).

**Table 2:** Out-of-sample tests. Testing period: 4/2/1998 – 2/5/2002. “Bull” market period: 4/2/1998 to 3/12/2000. “Bear” market period: 3/1/3/2000 to 2/5/2002.

Sub-period	RNN (no volatility)		RNN – MA (20)		RNN – RM (0.94)		RNN - GARCH (1,1)		RNN - GJR GARCH		
	Testing period	Bull	Bear	Bull	Bear	Bull	Bear	Bull	Bear	Bull	Bear
Total Return	0.292	0.569	-0.277	0.518	1.087	0.400	1.207	0.177	0.623	0.303	0.606
B&H Return	0.045	1.027	-0.982	1.027	-0.982	1.027	-0.982	1.027	-0.982	1.027	-0.982
Sign Rate	0.525	0.543	0.507	0.535	0.533	0.517	0.547	0.499	0.521	0.507	0.513
PT test	1.480 <sup>***</sup>	1.665 <sup>**</sup>	0.382	1.477 <sup>***</sup>	1.513 <sup>***</sup>	0.645	2.195 <sup>**</sup>	-0.008	0.971	0.265	0.594
Merton test	1.992 <sup>**</sup>	2.110 <sup>**</sup>	0.549	1.875 <sup>**</sup>	2.242 <sup>**</sup>	0.831	3.155 <sup>*</sup>	-0.010	1.751 <sup>**</sup>	0.344	1.104
MSPE	0.029	0.021	0.035	0.023	0.037	0.022	0.035	0.022	0.046	0.021	0.034
Sharpe Ratio	0.189	0.995	-0.316	0.900	1.230	0.695	1.370	0.300	0.711	0.521	0.679
Ideal Profit	0.016	0.081	-0.026	0.074	0.102	0.057	0.113	0.025	0.058	0.043	0.056

Notes: RNN = Recurrent Neural Network. Methods for forecasting volatility: MA(20) = Moving Average with a 20 days window, RM(0.94) = RiskMetrics’ exponentially weighted MA rule (decay factor = 0.94), GJR = Glosten-Jagannathan-Runkle (1993) GARCH model.

PT test = the Pesaran and Timmerman (1992) test. HT test = the Henriksson and Merton (1981) test. Both tests are asymptotically distributed as  $N(0,1)$ .

The sign rate measures the proportion of correctly predicted signs. The Sharpe ratio is defined as the ratio of the mean return of the strategy over its standard deviation (it has been annualized by multiplying it with the squared root of 250). The Ideal Profit is the ratio of the returns of the trading strategy over the returns of a perfect predictor.

(\*), (\*\*), (\*\*\*) indicate significance at the one sided 1%, 5% and 10% levels.