Neural Networks and Genetic Algorithms: Another Tool For The Technical Analysis of Financial Markets

Donn S. Fishbein, MD, PhD

The spectacular demise of Enron (ENRNQ, formerly ENE) is one of the more notable corporate failures in financial history. A bizarre sideshow to the collapse was the testimony of four noted Wall Street analysts before a Congressional committee in February of 2002 explaining how they could possibly have maintained “strong buy” recommendations on the stock up to and even after its demise was a virtual certainty. To the person, each said they relied on data supplied by Enron in their forecasts, and were thus mislead and as surprised as the general public. What they should have said was that despite their lofty status and salaries, they had little or no knowledge of even the basics of technical analysis of financial markets (in addition to the fact that their recommendations were based more on enhancing their firms’ investment banking businesses than enriching the sorry souls who entrusted their money to their firms’ brokerage businesses.)

Basic technical indicators foretold the downtrend of Enron’s stock price long before its accounting practices made the evening news. **Chart 1** shows the price of Enron stock plotted along with three simple moving averages (SMA) of different lengths. The n-period SMA is one of the most basic technical indicators, and for each period is simply the average price of the stock over the past n periods. N is typically between 5 and 200, and for the daily closing price chart shown below the period is one trading day. Thus to calculate a 10 period SMA of the daily closing price of a stock, add the past ten closing prices and divides by ten. The calculation is repeated for each day shown on the chart, and the result can be plotted as a line superimposed on the daily price chart. SMAs smooth the underlying price data, and function as low pass filters.

The direction of a SMA can be used to signal a trade direction, for example buying when the slope of the SMA turns positive. The SMA can also be used in relation to the underlying price, with a buy signaled when the price crosses above the SMA and a sell signaled under the opposite conditions. If SMAs of different lengths are plotted together, a buy is signaled when the shorter length SMA crosses above the longer length SMA. Using these rules and referring to **Chart 1**, a sell is clearly indicated in February, 2001, with no subsequent buy indicated. Note that the 10 day and 50 day SMAs converge in February, 2002 only because the stock simply can't go any lower.

The greatest weakness of SMAs is that they tend to arrive late to the dance (suffer from lag.) A part or the entirety of a market move may have already occurred before a SMA signals a buy or sell. This is most pronounced in a market with no major trend, where reliance on SMAs alone can lead to a long succession of small losing trades. In the worst case, when a SMA has a period equal to half of the period of a market instrument with a regular period, the signals are 180 degrees out of phase, and every trade signaled is a loser.

In an attempt to address weaknesses in the fixed period SMA, numerous variations have been developed. Two main strategies are to vary the weighting of individual prices, usually giving more weight to more recent prices, or to vary the period over which the average is calculated, either in a static or dynamic fashion. The exponential moving average (EMA) is expressed as the sum of a percentage of the previous period's EMA and a percentage of the current period's value. The effect is to deemphasize older data and more heavily weight current data. The relative weights determine the effective length of the EMA. Another approach is to change the length of the moving average based on the underlying data it is calculated upon. The calculated period may be fixed or vary from period to period.

The TRIX (TRIple eXponential moving average) is derived from the EMA, and is defined as the percentage rate of change of an EMA of an EMA of an EMA of the underlying price. The TRIX provides a smoothed version of the price, which can then be used in a number of ways. **Chart 2** shows a daily price chart for the S&P 500 market index, along with a TRIX indicator.

TRIX is an unbounded oscillator that varies about the zero line. Its simplest interpretation is that a zero crossing to the upside is a buy indicator, and a zero crossing to the downside a sell indicator. This simple interpretation yields only modest results, with an annualized yield of just 3.6%. **Chart 3** shows the trades (up arrow buy, down arrow sell) generated by this system over a five year period.

One of the difficulties in using this indicator is that there are a number of parameters to choose, namely the length of each of the three EMAs and the period over which the percentage change is calculated. One can rely on standard parameters (usually 9 or 12 periods) for each EMA, or determine them on a case-by-case basis. Then there is the problem of how to employ this indicator to generate buy and sell signals.

The application of neural networks and genetic algorithms (NNGA) can help with this task. Neural networks (NN) can assist in finding a relationship, or mapping, between inputs (price, TRIX) and an output (price sometime in the future.) Genetic algorithms (GA) can assist in the optimization of this relationship, particularly in selecting the optimal parameters for each EMA and the percent change indicator. A basic introduction to artificial neural networks, genetic algorithms, and technical analysis can be found in an earlier paper (*Fishbein, DS. Trading Financial Markets With Neural Networks and Genetic Algorithms, TCF Proceedings 2001*), as well as in many other locations (see reference list). The following discussion will explore various interpretations of the TRIX indicator as applied to trading instruments based on the Standard and Poor's 500 Index (S&P 500, SPX), along with the integration of several other market indices as potential NN inputs.

In the first implementation, a GA is used to select the optimal parameters for the TRIX indicator, and a training process is used to develop weights for the interconnections within a NN that map the input, the value of TRIX, to the output, the predicted opening value two days after the closing value. This predicted value is then used to signal buys and sells. Selling in this case refers to selling the index short, not simply closing the long position. The trades can be executed by either trading the SPDR (S&P 500 Depository receipts), or S&P 500 futures. The network is trained over a four year period, from the second quarter of 1996 to the first quarter of 2001. During the training period, the network showed an annualized return of 32.1%, more than double the return on a buy-and-hold strategy during the same period.

One danger with neural networks is that the network doesn't "learn" anything or develop the ability to generalize, but simply memorizes its input/output pairs. While a complete discussion of this problem is beyond the scope of this paper, one simple way to test if this has occurred is to train a network on a portion of the available input data, and then without any further network training (network weights frozen), apply inputs the network hasn't seen yet and measure network performance. If the network remains profitable in its forward looking evaluation period, its validity is more certain. In financial NNs, this is done by training the NN on a portion of the price data, and then testing on a later period. Using the TRIX NN described in the preceding paragraph, and without any further training, the NN is used to trade over a one year period immediately after the four year training period. During this one year evaluation period, the network continues to perform well, with an annual return of 42.0%, suggesting the network is able to generalize to data it hasn't been trained on.

A strategy for improving the performance of a network is to add additional inputs. Several additional inputs could be considered for the TRIX network. First, the closing price of the index could be added. Traditional descriptions of the TRIX (*Penn, D. TRIX is for Traders, Technical Analysis of Stocks and Commodities, Mar 2002, p 32-37*) include interpretations based on convergences, crossings, and divergences with the closing price, so this seems reasonable. However, the closing price represents only a single point in the day. A more comprehensive picture of the day's activity might be given by using the open, high and low values for the day, in addition to the close. **Chart 5** shows a network trained with the inputs O, H, L, C, and TRIX. While this network shows a higher annual return (54.1%) during the training period, the one year evaluation period shows a lower return of 22.6%. Analysis of the contribution of input factors shows that the training process gave the highest weighing to the O, H, L and C values, and used only a minor contribution from the TRIX indicator.

The volatility index (VIX) is a general measure of large cap stock volatility, and is based on the volatility of certain S&P 100 options. In general, volatility spikes are seen near market bottoms, while low volatility readings are of less certain significance. The number of advancing issues on the New York Stock Exchange (ADV_X) forms the only input to a daytrading system for the S&P 500 named "Oddball". This system has been widely discussed on the Omega Tradestation mailing list (omega-list@eskimo.com). While it is not immediately intuitive why ADV_X would be a good predictor of the S&P 500 index, "Oddball" does appear to be a profitable daytrading system.

When TRIX, O-H-L-C, VIX, and ADV_X are used as inputs to a network, the results are encouraging. However, analysis of the input contributions from each input shows that TRIX is not being given any weight at all. This is confirmed by repeating the analysis without TRIX, which yields the same results. Although this is unexpected, given the early success with TRIX, the synergy of O-H-L-C, VIX and ADV_X outperform the prior networks, showing a 60% annual yield during training and a 43.1% yield during the one year evaluation period (**Chart 6**).

The use of separate training and evaluation periods, with profitable results during the out-of-sample evaluation period, does not guarantee network validity. It is possible that profitable results could be seen in both periods due to random chance. The validity of the network can be further demonstrated by successive walk-forward testing. In this process, a number of input groups are held back from network training. The first held-back group is then used to test the network, and then is used along with the original training set to retrain the network. The new network is then tested on the next held-back group. Repeated profitability of networks tested on unseen data increases confidence in the overall validity of the network design. This process is shown using the OHLC-VIX-ADV network with six walk-forward periods of two months each (**Figure 7**), which shows an annual yield for out-of-sample evaluation periods of 32.6%. It is expected that the results should be slightly less favorable than a network trained over the entire period, but that the walk-forward tests would continue to show network validity (profitability).

Although the inputs to this NN have been chosen with the S&P 500 index in mind, particularly with regards to VIX and ADV, it is instructive to see how the NN performs when trained with and used to trade other market instruments. Boise Cascade Corp (BCC) is an integrated paper and wood products corporation whose stock shows little volatility and trades in a fairly narrow range. The Nasdaq-100 index (NDX) has shown much greater volatility over the last five years compared to either the S&P 500 or Boise Cascade.

Table 1 shows the results of training the NNs discussed above with the S&P 500 index (as already discussed), Boise Cascade, and the Nasdaq-100 index. Most of the networks remain profitable. Boise Cascade, whose market profile is more similar to the S&P 500 index, performs more consistently across the five systems than does the Nasdaq 100 index, particularly within the evaluation period. These findings give further confidence in the validity of these network designs. Selection of inputs more related to the underlying fundamentals of the Nasdaq 100 market could probably improve on results in that market.

Conclusion

The recent Enron debacle has served to highlight the shortcomings of traditional fundamental analysis of stocks as performed by most Wall Street analysts. Technical analysis can provide a methodical, emotionless approach to trading financial markets. The use of neural networks and genetic algorithms improves on traditional technical analytic techniques. Networks of fairly simple design can achieve quite satisfactory results.

References

Acheli, Steven B. *Technical Analysis from A to Z.* **http://www.equis.com/free/taaz/**

Fosback, Norman G. *Stock Market Logic*. Dearborn Financial Publishing, Inc., 1991.

Zirilli, Joseph S. *Financial Prediction Using Neural Networks*. International Thomson Computer Press, 1997.

Zweig, Martin. *Winning On Wall Street*. Warner Books, 1986.

TABLE 1

Market/ Network	Training, Annualized Return %	Evaluation, Annualized Return (%)
S&P 500 Index		
TRIX NO OPT	3.6	n/a
TRIX OPT	32.1	42.0
TRIX OHLC	56.1	22.6
OHLC VIX ADV	60.0	43.1
OHLC VIX ADV 6 walk	60.6	32.6
Boise Cascade Corp		
TRIX NO OPT	-8.4	n/a
TRIX OPT	36.3	2.5
TRIX OHLC	54.1	49.4
OHLC VIX ADV	76.6	73.5
OHLC VIX ADV 6 walk	75.1	37.1
Nasdaq 100 Index		
TRIX NO OPT	12.7	n/a
TRIX OPT	118.0	85.8
TRIX OHLC	155.2	2.5
OHLC VIX ADV	107.8	37.4
OHLC VIX ADV 6 walk	180.5	-9.0

CHARTS

Chart 1 – Enron and 3 Simple Moving Averages

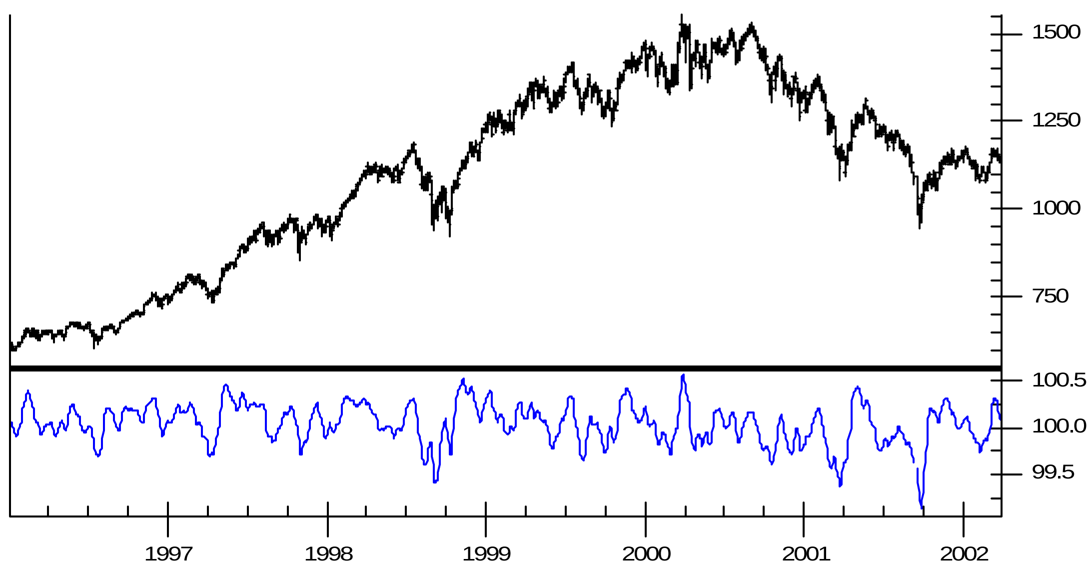

Chart 2 – SPX and TRIX

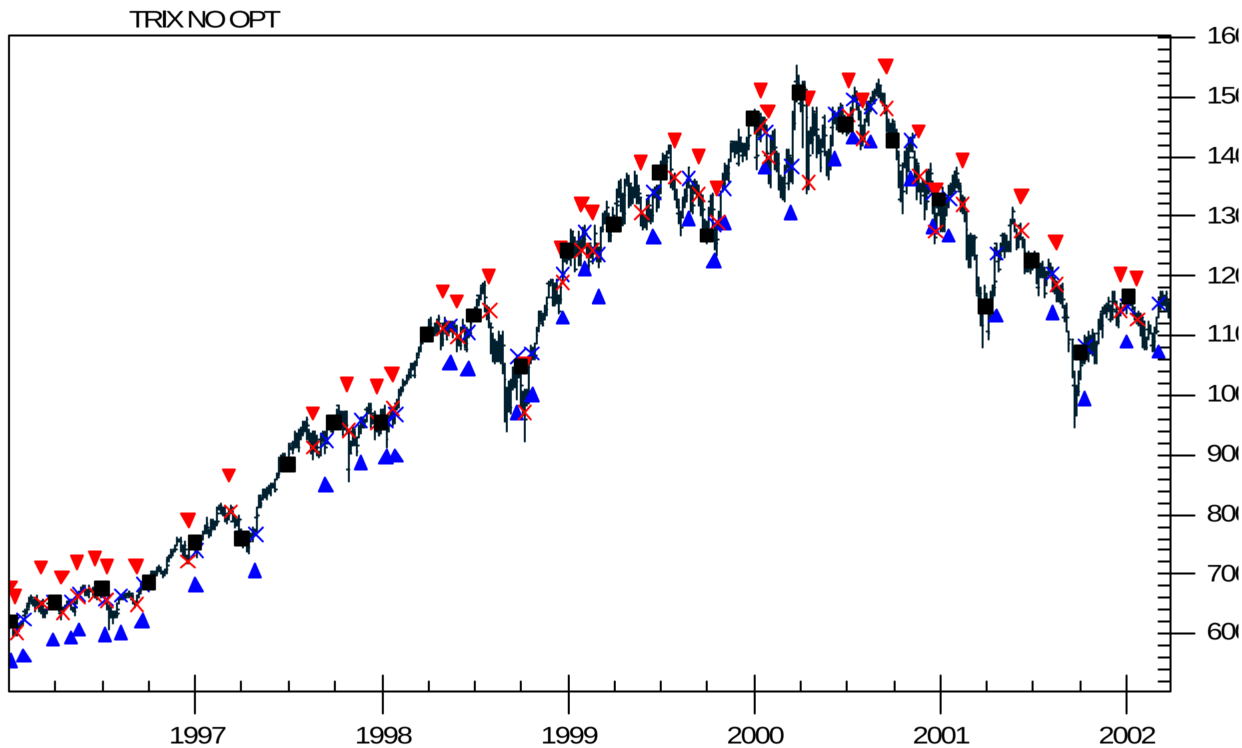

Chart 3 – Trading the SPX with the unoptimized TRIX indicator

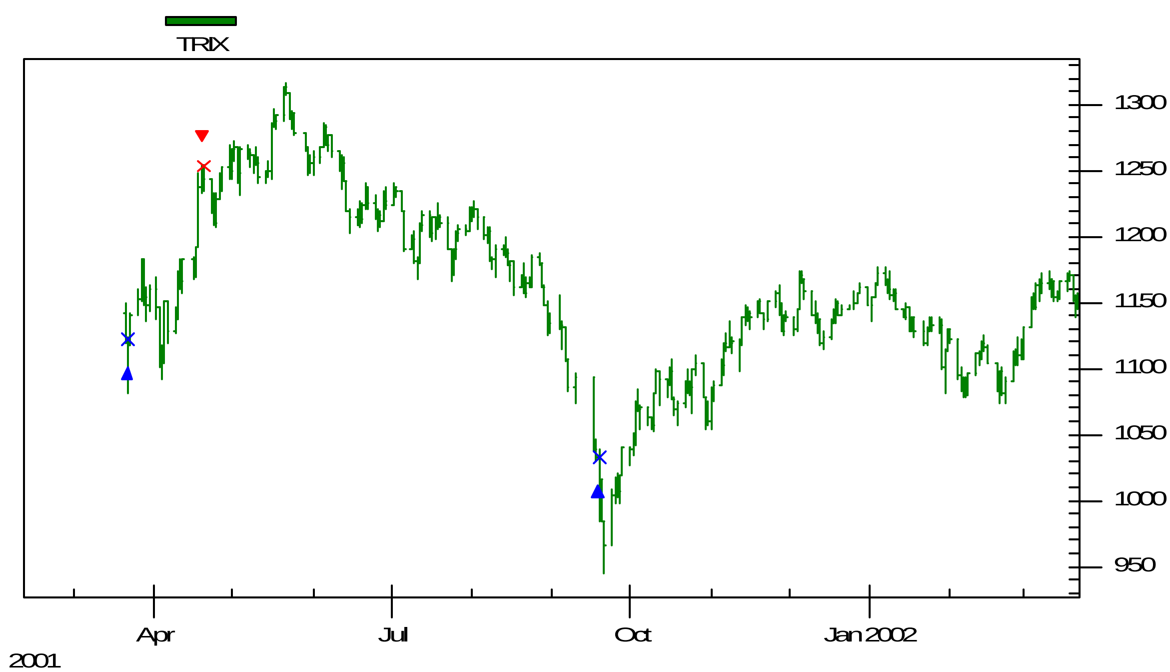

Chart 4 – Trading the SPX with an optimized TRIX indicator

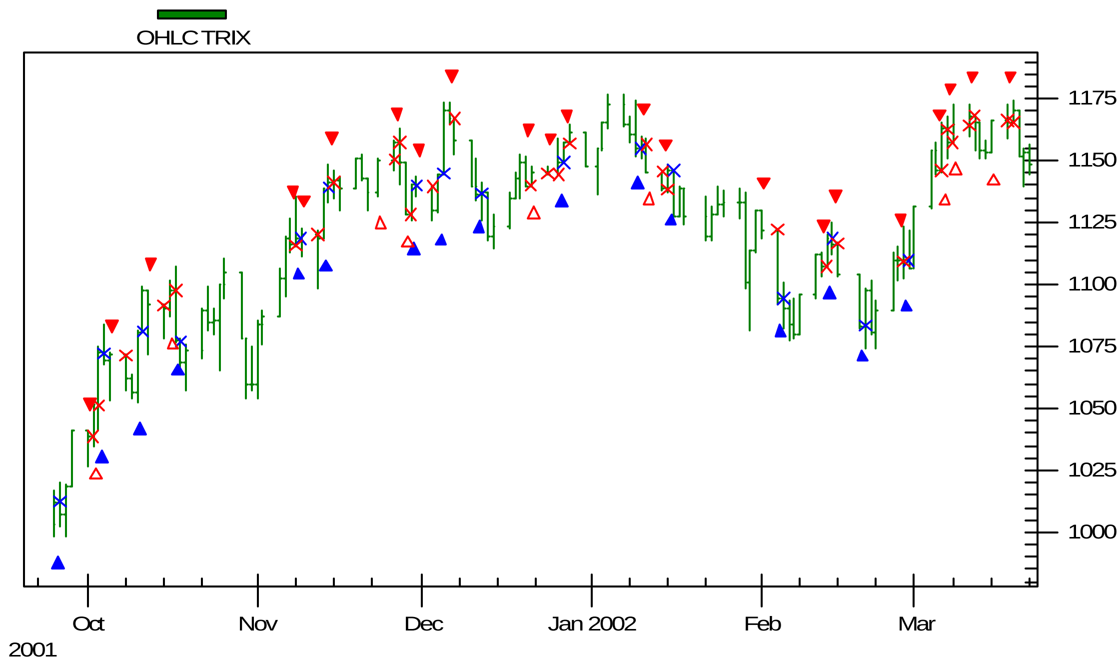

Chart 5 – Trading the SPX with an OHLC and TRIX network

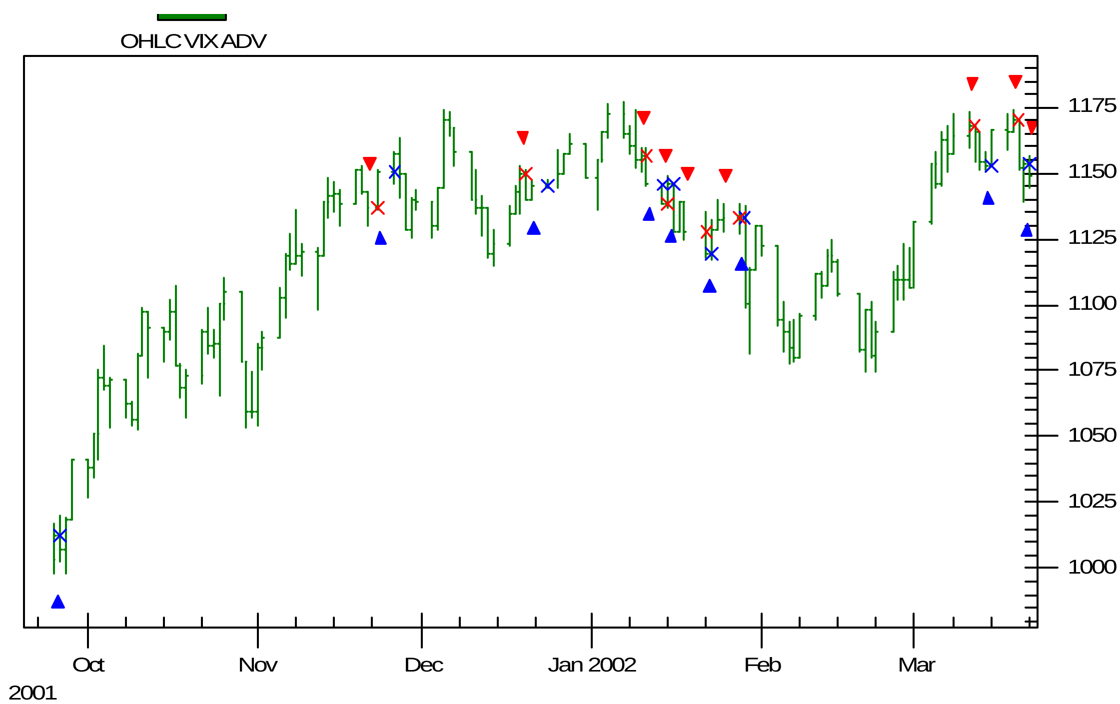

Chart 6 – Trading the SPX with an OHLC, VIX and ADV network

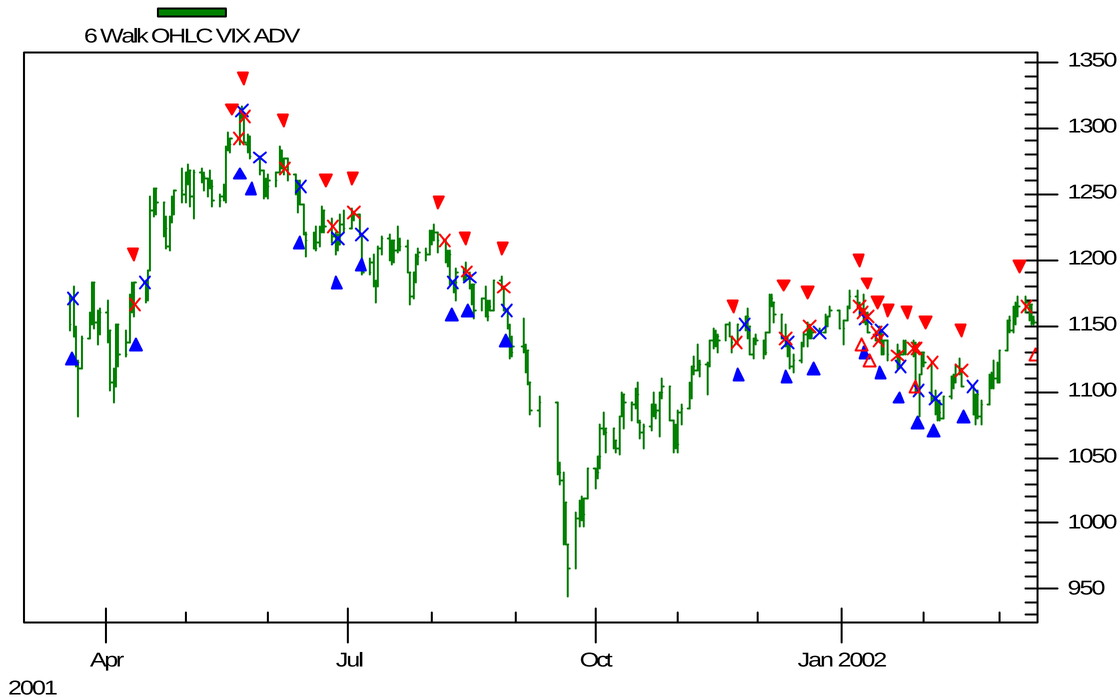

Chart 7 – OHLC, VIX, ADV network, 6 walk-forward periods