Neural Network Stock Trading Systems
Donn S. Fishbein, MD, PhD
Neuroquant.com

There are at least as many ways to trade stocks and other financial instruments as there are traders. Remarkably, most people trade without any system at all, relying on intuition, hunches, hearsay, and random choice to guide transactions. These rudderless traders provide the losing trades to make possible the winning trades of more directed traders.

Among those who follow a system, three basic approaches are often used. Buy-and-hold traders make the assumption that the market goes up in the long term, following an ever-strengthening economy. They may also subscribe to the efficient market theory, a corollary of which states that efforts to time the market will be futile. Therefore, it is adequate to purchase a stock and sit back and wait. Unfortunately, the wait might be a long one. An investor who purchased the NASDAQ index near its peak of over 5000 in 2000 has not yet recovered half of their investment as of March, 2006.

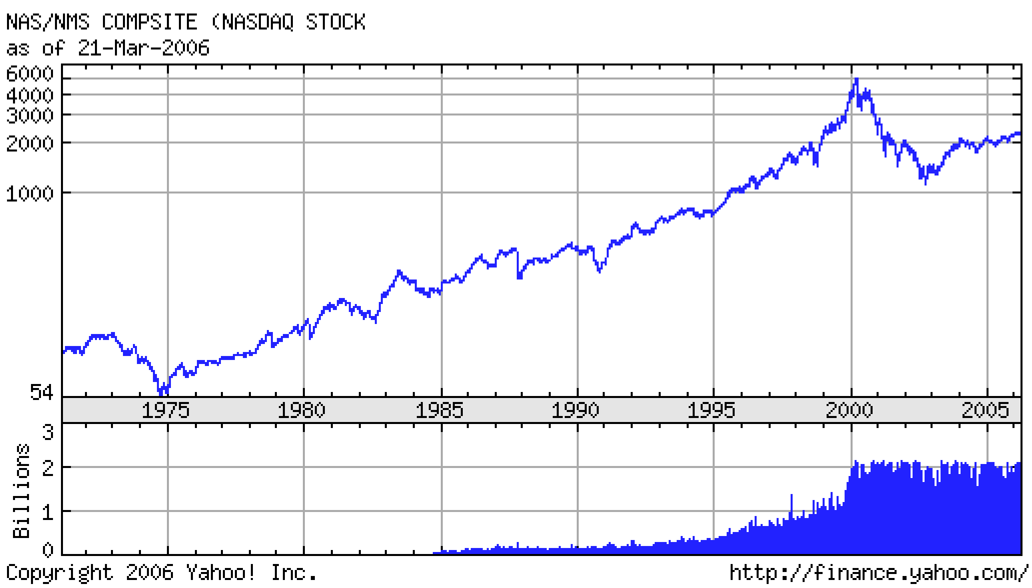

Another approach attempts to select stocks based on their fundamental characteristics, such as major markets and products, earnings, dividends, cash flow, sales, and profits. Two of the most successful stock market investors in recent years, Warren Buffett and Peter Lynch, were primarily fundamental analysts. Fundamental analysis relies on extensive research, and may utilize information not available to the ordinary investor. It is time consuming, difficult to automate, and may require resources beyond those available to the individual investor.

Technical analysis makes the assumption that everything known about a stock is reflected in its market data, primarily its price but also volume, open interest (options), etc. The intrinsic, or fundamental value of the stock is not considered. Mathematical functions of independent variables such as price are frequently employed. Technical analysis offers the advantages that data is available to all at little or no cost, that the analyses can be automated, and that assumptions can be tested on older market data before being used for live trading.

At first glance technical analysis can seem overwhelming. A typical technical analysis software package offers hundreds or thousands of technical indicators, each with their own parameters. Selecting the independent variables to use, which indicators based on these variables, and what trading rules to employ are just part of the task. As is often the case, simpler is frequently better.

One often heard objection to technical analysis is that is just doesn't work. Technical analysis, according to its detractors, is an excellent tool to analyze the past, but simply can't predict the future. Bolstered by the academic efficient market theory and the random walk hypothesis, it is argued that any attempt to predict future market direction is futile. These naysayers are wrong, and a simple example can prove this.

Consider a basic trading system, first described decades ago. If a stock increases by 4% in value, buy it. If it declines by 4% in value, sell it short. That's it. A few problems are apparent. By waiting for a 4% move in either direction, we're giving up a portion of the potential profit. If the market moves in a sawtooth fashion, 4% up and then 4% down repeatedly, we'll slowly lose our money, 4% per trade. But in fact this simple system makes money year after year, handily beating market averages. The chart below shows buys and sells for the period from 2000 to 2004 for the Nasdaq-100 exchange traded fund (QQQ (now QQQQ)).

If such a simple system can outperform buy-and-hold investing, a reasonable question to answer is if a more sophisticated system can perform even better. Technical analysis is an extensive field, with numerous indicators, systems, and theories. The choices seem overwhelming. How does one select and combine various tools to form a winning system? Having selected indicators and rules, how does one optimize the system for the specific market being traded?

Artificial neural networks offer one solution to this problem. ANNs, a form of non-linear modeling, represent a very simple approximation of biological neural networks, the neural architecture the human is based upon. Without extensively reviewing the subject of ANNs, it is noted that ANNs consist of a multi layered network of artificial neurons. Each neuron has inputs and outputs, and its connections (synapses) to neurons in the adjacent layers are weighted. The knowledge stored in an ANN is contained both within the overall architecture of the network and within the synaptic weights.

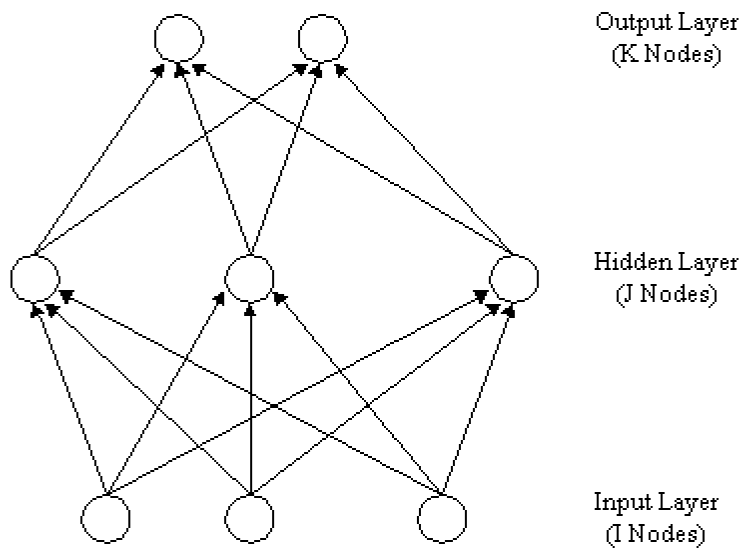

The strength of ANNs lies in their ability to represent non-linear relationships and in their ability to learn these relationships directly from the data being modeled. In the case of financial modeling, learning is accomplished by training the network with historical data. The network is shown past data, along with the desired prediction. The ANN output is compared to the actual value, and the error is then back propagated through the network and used to adjust synaptic weights. Precautions are taken to assure the network is generalizing and not simply memorizing the data. The process is continued until the error reaches the desired threshold. Once trained, the network can then be tested with data it has not been trained on (out-of-sample data) to test its predictive ability. This architecture and training method, referred to as "feed forward back propagation," is a common ANN architecture.

A Neural Network Trading System Based on Common Technical Analysis Indicators

Moving averages, stochastics, the average directional indicator (ADX), and the moving average convergence-divergence (MACD) indicator are commonly used in technical analysis. Each has its strengths and weaknesses, and each has deficiencies if used alone. Could they be combined to yield a more useful trading system?

Each indicator will be discussed in turn. Next a simple neural network will be constructed using these four indicators. The network will be allowed to decide how to weight each indicator, and a genetic algorithm will select the most profitable parameters. Finally, the results of trading several common stock index exchange traded funds with the neural network system will be reviewed.

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.

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 stochastic oscillator quantifies the current market close with relation to the upper and lower bounds of a defined period. Closes near the period's high are thought to represent overbought positions, and closes near the period's low oversold conditions. Another interpretation is that rising stocks tend to close near the period's high, and declining stocks tend to close near the period's low. The stochastic oscillator varies from 0 to 100, with 20 typically thought of as oversold and 80 as overbought. Stochastics are generally used in combination with other indicators.

The ADX (Directional Movement Indicator) is a trend measurement system. It is an oscillator bounded by 0 and 100. Values under 20 indicate a weak trend, and values over 40 a strong trend. By itself it does not indicate the direction of the trend.

The moving average convergence/divergence (MACD) indicator assists in identifying changes in trends. It is based on the difference between 26-day and 12-day exponential moving averages. The difference is referred to as the fast line. The slow line is a 9-day EMA of the fast line. Crossovers of the fast and slow lines may signal trends.

Each of these four indicators have strengths and weaknesses. For example, moving averages tend to perform well in a trending market, stochastics generate many false signals during a strong trend, and ADX identifies the presence or absence of a trend but not its direction. Perhaps using these indicators in combination could improve a trading system.

A simple neural network trading system can be constructed using these four indicators as inputs. The neural network will decide through a training process how to best combine the indicators, while a genetic algorithm will optimize the parameters of each indicator. The chart below shows each indicator plotted separately, along with the composite neural network. The chart is plotted for the S & P 500 exchange traded fund SPY, with buys indicated with a blue up arrow and sells with a red down arrow.

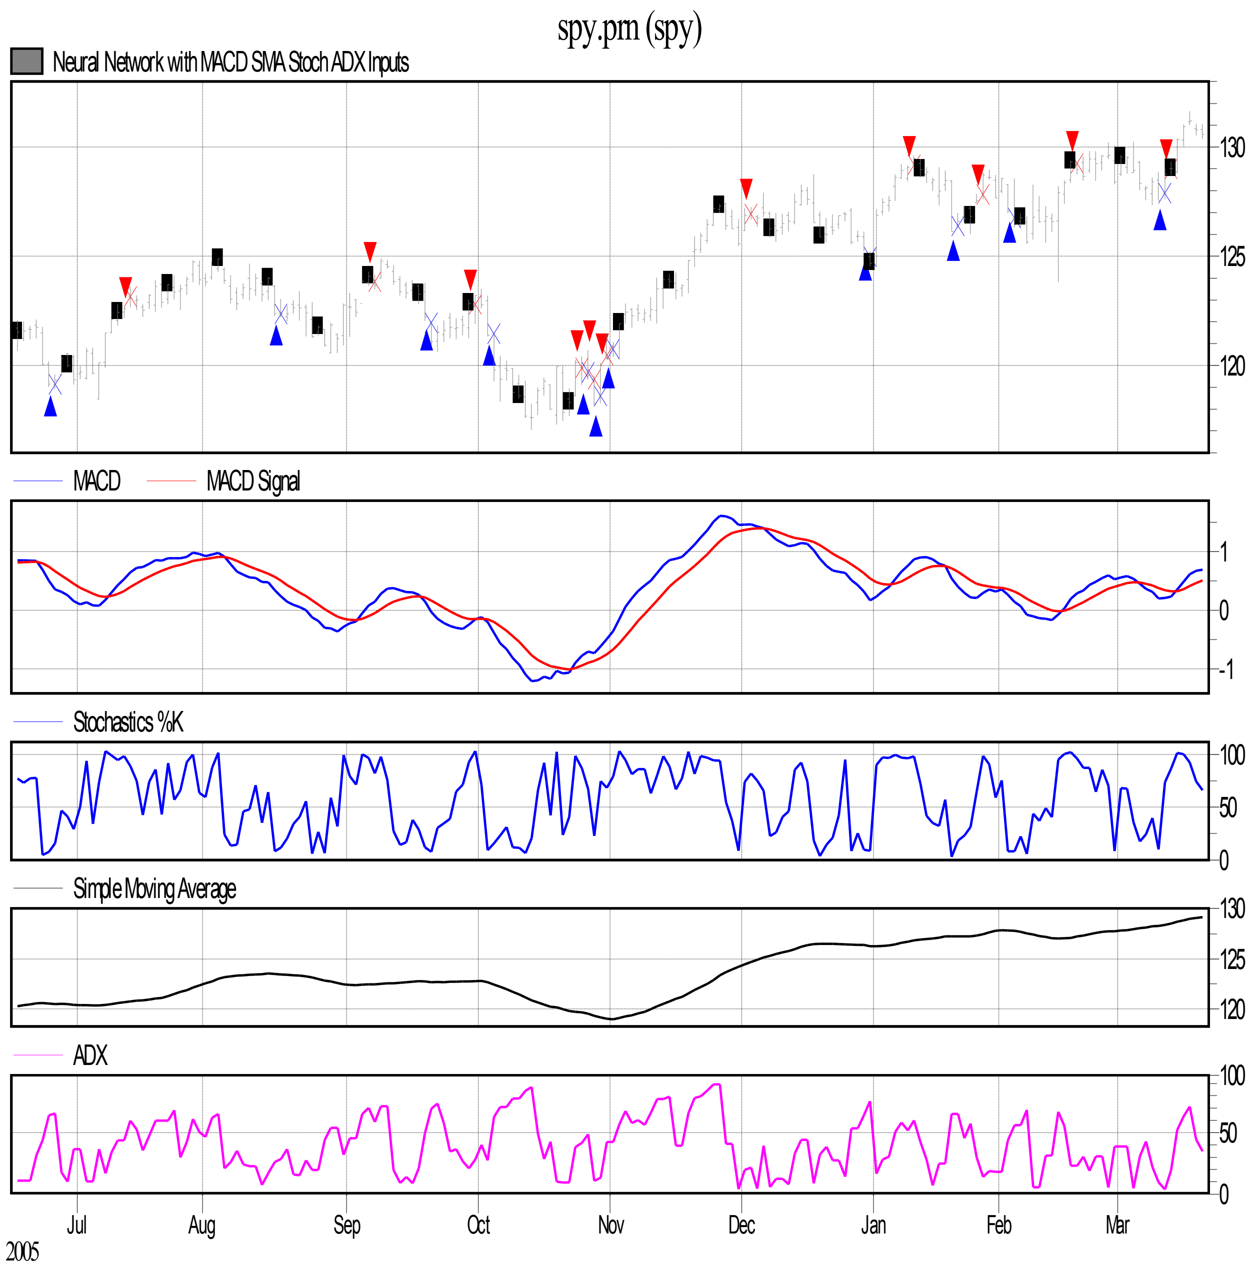

The neural network was initially trained on data from March 2004 to March 2005. It was evaluated on data from March 2005 to March 2006 with a 1 day walk-forward period. The account began with $10,000, and all profits were reinvested with each transaction. The system traded both long and short, and allowed $25 in transaction fees for each trade. Trades were signaled after market close and executed at market open the next day. For the one year evaluation period, the following results were recorded:

Symbol	***1 year return***	***# Trades***	***% WinningTrades***
DIA	31.1%	21	81%
IJR	78.2%	30	90%
QQQQ	63.0%	21	90%
SPY	38.9%	38	76%

This neural network system with genetic optimization proved profitable with each instrument with a large percentage of winning trades. The more volatile exchange traded funds provided a higher profit margin. Results such as these were not obtainable with any single conventional indicator. The neural network system had the ability to adapt to and remain profitable in changing market conditions.

References

Arnold, Curtis M. Timing The Market. Weiss Research, Inc., 1993

Chen CH. Fuzzy Logic and Neural Network Handbook. IEEE Press, 1996.

Fishbein DS. Trading Financial Markets with Neural Networks and Genetic Algorithms. Trenton Computer Festival Proceedings, 2001.

Fishbein DS. Neural Networks and Genetic Algorithms: Another Tools for the Technical Analysis of Financial Markets. Trenton Computer Festival Proceedings, 2002.

Fishbein DS. Mechanical Stock Trading Systems That Really Work. Trenton Computer Festival Proceedings, 2003.

Fishbein DS. How to time the stock market Using Artificial Neural Networks and Genetic Algorithms. Trenton Computer Festival Proceedings, 2004.

Fishbein DS. Neural Networks and Stock Market Timing. Trenton Computer Festival Proceedings, 2005.

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

Murphy John. Intermarket Analysis. Wiley Trading, 2004.

Neuroquant.com: www.neuroquant.com

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

Zweig, M. Winning of Wall Street. 1986, 1997(revised), Warner Books.