Mechanical Stock Trading Systems That Really Work
Neural Networks and Genetic Algorithms

Donn S Fishbein, MD, PhD
Neuroquant.com

People who invest with an eye towards making money face a gauntlet of those who think they are more deserving of your money than you are. These villains range from the blatantly illegal (and obvious) Nigerian spam scam practitioners and pyramid schemes, to equally illegal but not so obvious stock market manipulation schemers, to completely legal purveyors of trading systems, courses, and manuals. With regards to the latter, it is staggering to see how much misinformation is sold, sometimes at exorbitant prices. Much of this material serves only to enrich the vendor.

The only defense against overpaying for financial advice, and worse, following this advice only to lose money, is to fully understand the basis of the advice you are following. To this end, the purpose of this monograph and those which have preceded it (Fishbein 2001, 2002) is to fully disclose trading methods which have been developed and tested over a number of years. All of these systems utilize traditional technical analysis and artificial neural networks (ANN) and/or genetic algorithms (GA). They represent one way to trade that has proved to be consistently profitable.

The most basic question to ask when evaluating a trading system is whether or not it is profitable. This question in itself is a complex topic, with innumerable statistics described to evaluate trading system. Just a few points will be made here. First, it is important that for whatever period of time and number of trades one considers, the profit should be distributed fairly evenly. A system that makes the majority of its profit over a five year period in a single, 4 day trade is less likely to repeat this performance over another period. Second, the maximum draw-down, or largest negative excursion that the account balance makes, is critical in that it determines the likelihood that one will actually stay with a system. A wildly profitable system that happens to have several 85% draw-downs is likely to shake out the most stalwart investor before it has a chance to show its dazzling return.

Regardless of how profit is measured or evaluated, it is meaningless as a measure unless the system has been tested on out-of-sample data. Out-of-sample data refers to data not used during development and optimization of the trading system. A trading system that can perform well using data that has already been used to optimize it may simply represent curve-fitting, or in the case of a neural network, rote memorization. If a trading system can perform well on data it has not seen before, the odds that if can perform well trading in the future are increased. It is not necessary, as some purists claim, that the out-of-sample data occur in the future and the system be tested in real-time, as this restriction greatly increases the development period of a system, and logically makes no sense.

To this end, all systems which have been presented at the three immediately previous Trenton Computer Festivals on this same topic will be examined to determine their profitability on previously unseen data. No alterations will be made to the systems as originally presented, although it is noted that artificial neural networks have a natural ability to adapt to changing conditions.

A variety of stocks were used to demonstrate the systems when initially presented. For the sake of consistency, each system is retested using the S&P 500 and Nasdaq 100 tracking stocks (SPY and QQQ). Each system was trained on 48 months of data. Network weighting and indicator parameters were fixed, and then the systems were tested out-of-sample (evaluation period) on the following 6 months of data, ending in March of 2003. All systems traded both long and short. Returns are stated on an annualized basis. More detail on each of the systems can be found in the respective TCF Proceedings, or at www.neuroquant.com. These references also contain more background information on technical analysis and artificial neural networks.

Trenton Computer Festival 25 May, 2000

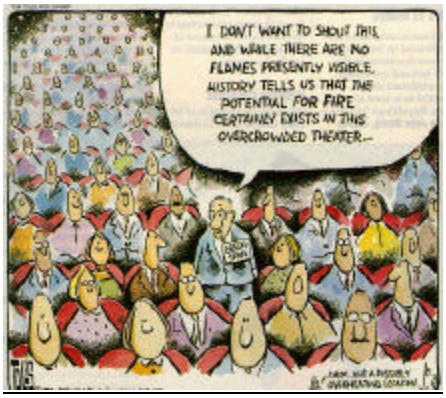

May, 2000 saw the great stock market run-up peaking, with many investors staring blindly and unknowingly into the abyss ahead. In a cartoon shown during the 2000 TCF presentation, Chairman Greenspan, whose comments in those days were felt to significantly influence markets, stands in a crowded theater and says, “I don’t want to shout this, and while there are no flames presently visible, history tells us that the potential for fire certainly exists in this overcrowded theater.” Truer words were never spoken.

System 1: An ANN/GA used two lags of the closing price and two lags of the money flow index (a volume weighted form of the relative strength index.) The system was demonstrated initially using SPY, the S&P 500 tracking stock. The network was allowed to choose the parameters and the weight for each input.

Test Date	Symbol	Annual Return Training	Annualized Return Evaluation
4/2000	SPY	Not stated	29.8%
4/2003	SPY	22.3%	4.6%
4/2003	QQQ	51.4%	44.7%

System 2: An ANN/GA used a variety of price momentum (direction movement indicator, accumulation distribution indicator, and stochastic %k) and volume (money flow index, volume price trend) indicators, with the network allowed to choose the parameters and the weight for each input.

Test Date	Symbol	Annual Return Training	Annualized Return Evaluation
4/2000	XRX	Not stated	56.7%
4/2003	SPY	33.1%	20.0%
4/2003	QQQ	74.3%	44.7%

System 3: A trading system used three ANN/GA predictions as inputs, and signaled a trade when 2 or 3 of the inputs signaled a trade. The first network used multiple lags of linear regression of the closing price line. The second network used five price momentum indicators (accumulation distribution, commodity channel index, moving average convergence divergence, relative strength index, and stochastic %k). The third network used five volume indicators (negative volume index, positive volume index, Williams's variable accumulation distribution, volume life force, and volume price trend).

Test Date	Symbol	Annual Return Training	Annualized Return Evaluation
4/2000	BGEN	74.7%	34.5%
4/2003	SPY	34.7%	13.1%
4/2003	QQQ	15.6%	45.8%

In summary, the systems initially presented in May, 2000 were retested using two index tracking stocks. Experience has shown that index tracking stocks tend to be less volatile than many individual stocks, and at least as far as ANN/GA systems are concerned, are some of the most difficult instruments to trade profitably. Each of the systems was profitable to varying degrees when retested, and considerably outperformed buy-and-hold strategies, which were uniformly negative the test period.

Trenton Computer Festival 26: May, 2001

May, 2001 saw many investors badly bloodied by the market decline of 2000-2001, and hoping wistfully for the return of the E-trade "money coming out the wazoo" guys,

whose passing from the airwaves was lamented during the TCF 2001 talk. Unfortunately, the market decline continued. ANN/GA systems continued to make money.

System 1: A trading strategy (N-Cross) triggered when at least two of eight crossover indicators within the price momentum category triggered. These indicators included stochastic %K/%D crossover, commodity channel index crossover, exponential moving average envelope breakout, %R breakout, Bollinger band breakout, high channel breakout, closing price moving average vs. opening price average, and exponential moving average crossover. This system is a pure genetic algorithm, which selects both indicator parameters and rule weighting, and involved no prediction.

Test Date	Symbol	Annual Return Training	Annualized Return Evaluation
4/2001	HGSI	Not stated	143%
4/2001	BCC	Not stated	71.5%
4/2003	SPY	45.9%	35.3%
4/2003	QQQ	78.2%	1.9%

System 2: In a departure from standard ANN/GA systems used above, the AMA-SMA system employs a general regression neural network (GRNN) to calculate a smoothed, adaptive moving average of the closing price. The relationship $SMA_1 < AMA < SMA_2$ is evaluated, where AMA refers to the adaptive moving average, and SMA_1 and SMA_2 are simple moving averages whose parameters are determined by a GA. The same expression is evaluated for triggering long and short trades, with only the parameters for each SMA differing. Separate training and evaluation periods are typically not employed with systems based on GRNNs, because these systems are never predicting the training set. Rather, they perform a comparison of today's bar with all of the bars in the training set. The output they produce is in effect a weighted average of all of the outputs of those bars which are similar to today's bar.

Despite its apparent simplicity, this system has proved robust and profitable over a variety of financial instruments and timeframes.

Test Date	Symbol	Annualized Return
4/2001	HGSI	45.8%
4/2001	BCC	37.3%
4/2003	SPY	35.5%
4/2003	QQQ	80.0%

In summary, the two systems presented in 2001 have remained profitable when tested with data unseen at the time of system design. Buy-and-hold during the test periods resulted in losses. System 1 (N Cross) showed only marginal profitability during the

evaluation period trading QQQ, but still considerably outperformed a buy-and-hold strategy.

Trenton Computer Festival 27: May, 2002

In May, 2002, this Farside cartoon was shown at the TCF talk, relabeled "Market Prospects 2002" by the speaker. Market predictions are generally unwise and unnecessary for those who use mechanical systems which trade both sides of the market, but sometimes humor is difficult to resist. The markets did in fact continue to decline.

System 1: The TRIX indicator was used as the single input to an ANN/GA system. The TRIX indicator is defined as the rate of change of a triple-smoothed exponential moving average of the closing price. Rather than use the parameters and interpretation defined by its author, the system was allowed to choose the indicator parameters and thresholds for long and short triggers.

Test Date	Symbol	Annual Return Training	Annualized Return Evaluation
4/2002	BCC	30.3%	2.5%
4/2002	SPY	37.1%	7.1%
4/2002	QQQ	120.5%	30.3%
4/2003	SPY	30.1%	7.1%
4/2003	QQQ	170.9%	30.7%

System 2: The second system presented was the AMA – 2 SMA system originally presented at TCF 2001. No alteration was made to the system construction or operation from the 2001 version.

Test Date	Symbol	Annualized Return
4/2002	SPY	33.4%
4/2003	SPY	35.5%
4/2003	QQQ	80.0%

Both systems showed consistent results from year to year. As is usually the case, the more volatile instrument (QQQ) showed greater returns. Although it is not

highlighted in these tables, QQQ also suffered a greater maximum drawdown than SPY, consistent with greater risk.

Trend Detection Index evaluated with an adaptive neural indicator

This system is presented for the first time in this series. M. H. Pee described the Trend Detection Index (TDI) in an article in 2001 (TASC, 2001). The 20 day TDI was defined as the absolute value of the sum of 20-day momenta (price today minus the price 20 days ago) of the last 20 days (AV20), minus the sum of 20-day absolute momenta of the last 40 days (SumAM40), minus the sum of 20-day absolute momenta of the last 20 days (SumAM20). A positive value indicates a trend, and a negative value a consolidation. The direction index (DI) was defined as the sum of the 20-day momentum of the last 20 days. The trading rules were defined as enter long if both TDI and DI were positive, and enter short if TDI was positive and DI was negative. When tested on a basket of commodities, the system proved profitable, relying on a high profit/loss ratio to compensate for between 45% and 36% winning trades.

Rather than rely on the fixed parameters and thresholds in Pee's system, the components of his system (AV20, SmAM40, SumAM20, and DI) were used as inputs to a GRNN, an adaptive neural indicator. The system was tested on three index tracking stocks (DIA, SPY and QQQ), and three technology stocks (MSFT, INTC, CSCO). The results were uniformly positive, with all instruments showing profitable trading over a four year period, a time when the overall market was in decline (March 1999 to March 2003). Results are calculated based on reinvesting all profits, a $10/trade commission, with trades signaled after market close and executed at the following day's opening price.

Symbol	Annual return	Max Drawdown	% Profitable	Profit Factor*
DIA	19.6%	18%	58.1%	3.3
SPY	8.6%	17%	67.4%	2.81
QQQ	89%	42%	59.4%	1.13
CSCO	101%	31%	51%	1.89
INTC	79%	32%	55.7%	2.3
MSFT	83.7%	29%	49.6%	1.7

*profit factor = average win / average loss

A few observations can be offered on the TDI/GRNN system above. As with many trading systems that go long and short, the more volatile stocks were more profitable, albeit with higher risk. The systems trade frequently, with average trades lasting four to five days. There is a high percentage of losing trades, which is counterbalanced by a profit factor significantly greater than one. This system was evaluated during a period of high market volatility. Whether the system would perform as well in a more placid market has not been evaluated yet.

Wise advice which bears repeating

The following suggestions for trading with neural nets are from Ward Systems, and have been offered before in this series:

- Pick an instrument that is rising and falling because of trading, not fundamental factors, if you want to try to predict with technical indicators.
- Pick a volatile instrument to predict; one that repeatedly rises and falls is the best way to phrase it. Don't pick something that has only gone up in recent memory. This is the most important thing you can do; I can't emphasize it enough.
- Predict the percent change rather than the change or the actual price. It is more normalized.
- Use only about three to six indicators as input. This is the best way to prevent over-fitting problems.
- Make sure that your inputs are all as different as possible when graphed. Don't pick an indicator and 4 lags of it; they will be too similar.

Summary

Trading financial markets can be approached as an art, a science, or a crapshoot. This paper has tried to focus on a very small branch of the science of trading, that of the use of artificial neural networks and genetic algorithms. These techniques provide one path to approaching financial markets with a methodical, analytical method, and have proved profitable. It is the responsibility of the successful trader to fully understand his or her tools, to not take any claim for granted without proof, to innovate, and to teach and learn from others.

Acknowledgement

I would like to thank Don Libes for proofreading this paper and offering many helpful suggestions, as well as raising questions for future examination.

References

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

Fishbein, D. *Trenton Computer Festival Proceedings*, 2001, 2002.
www.neuroquant.com

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

Pee, MH. Trend Detection Index. *Technical Analysis of Stocks and Commodities,* 19:10, 54-61, 2001.

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