The Value of Technical Analysis

Matthew C. Roberts[1]

December 5, 2002
Keywords: Technical Analysis, Genetic Algorithms, Equity Forecasting

[1]Assistant Professor, Department of Agricultural, Environmental, and Development Economics, The Ohio State University. 320 Ag. Admin. Bldg., 2120 Fyffe Rd., Columbus, Ohio 43210-1067. Phone: (614) 688-8686. Fax: (614) 292-0078. email: roberts.628@osu.edu

Abstract

Even though there is little academic research that supports the usefulness of technical analysis, its use remains widespread in financial markets. One explanation previously offered is the ability of technical methods to identify periods of high volatility. Using Genetic Programming techniques, *ex ante* optimal technical trading strategies are generated. Because they are mechanically generated from simple arithmetic operators, they are free of the data-snooping bias inherent in technical analysis research. The fee an investor would be willing to pay for these trading rules is estimated. The results indicate that highly risk-averse agents could significantly benefit from technical strategies.

Technical analysis is the forecasting of prices based upon previous prices and/or volumes. There are many academic studies that test the profitability of technical trading strategies. Early studies such as Alexander and Fama and Blume apply simple filter rules, which are long when prices rise $x\%$ above their n-day low, to equities and equity indices and find that although such models may have predictive power, they are not capable of generating excess returns after transactions costs. Later studies have supported these findings.

Dooley and Shafer (1976) and (1983) find "remarkable profits" from incorporating filter rules into exchange-rate models, although these results do not account for transactions costs. Sweeney (1986) confirms that these results continued to hold in later periods. Sweeney (1988) finds that filter rules can produce excess returns net of transactions costs for NYSE equities. Neftci finds that certain moving average rules, which are long when prices rise above an n-day moving average, applied to the DJIA do generate excess returns over certain time periods, though the issue of transactions costs is not addressed. Brock, Lakonishok, and LeBaron find that filter and moving average rules can generate positive returns, but these may be overshadowed by transactions costs.

Lo, Mamaysky, and Wang analyze the effectiveness of the 'head and shoulders' pattern for individual US equities. They conclude that

> "We find that certain technical patterns, when applied to many stocks over many time periods, do provide incremental information, especially for Nasdaq stocks. Although this does not necessarily imply that technical analysis can be used to generate 'excess' trading profits, it does raise the possibility that technical analysis can add value to the investment process."

Osler (1998) also examines the head and shoulders chart formation in foreign exchange markets, and is particularly noteworthy as it is one of only a few studies that identifies profitable trading strategies. Osler finds that "these aggregate profits would have been both statistically and economically meaningful regardless of transactions costs, interest differentials or risk." In application to equities, however, Osler and Chang find that that "head and shoulders trading is not profitable."

Regardless of its standing in academic circles, technical analysis is very common among practitioners. Oberlechner (2001) surveys traders on their use of technical analysis, and finds that "Only a very small minority of foreign exchange traders demonstrate an exclusively fundamental or exclusively chartist overall forecasting approach." This is consistent with the previous survey research performed by Taylor and Allen (1992), Menkhoff (1997), and Lui and Mole (1998). These findings are a stark contrast to the general disdain of technical analysis held by most academic practitioners. Malkiel (1981) sums up the consensus view, "Obviously, I am biased against the chartist. This is not only a personal predilection, but a professional one as well. Technical analysis is anathema to the academic world." This quote, along with most academic research, epitomizes the academic view of technical analysis, that it is primarily a manifestation of investor irrationality. Neftci offers one possible explanation, that technical analysis might be a practical method for reducing highly nonlinear prediction problems to more tractable forms. Another possible explanation is offered by Brock, Lakonishok and LeBaron, who note the consistent ability of simple moving average rules to forecast periods of high returns and low volatility, although the profitability of such forecasts in the presence of transactions costs is questionable. The subsequent findings of Sullivan, Timmermann and White and Allen and Karjalainen support this hypothesis. It is also worth noting that while technical methods date at least to the 1920's, and anecdotal evidence points to use as far back as the 17th century, it was not until Engle that relatively low-cost methods to estimate time-varying volatility became available.

Therefore, this research attempts to evaluate technical analysis on its ability to increase the utility of agents of varying degrees of risk aversion. Using the method of Fleming, Kirby and Ostdiek, the fee that a risk-neutral investor would be willing to pay for an *ex ante* optimal technical strategy is estimated for 33 two-year periods between 1935 and 2000.

To eliminate potential data-snooping biases, the technical trading strategies are produced through genetic programming (GP). Technical analysis research prior to Neely, Weller and Dittmar and Allen and Karjalainen use trading rules that were already in common use, which introduces the risk of data snooping. According to Sullivan, Timmermann and White, data snooping occurs

> "when a given set of data is used more than once for purposes of inference or model selection. When such data reuse occurs, there is always the possibility that any satisfactory results obtained may simply be due to chance rather than to any merit inherent in the method yielding the results."

In essence, previous research has focused on the performance of technical strategies in common use by applying them to historical data. However, the rules' common usage is *prima facie* evidence that the rules have been successful when applied to historical prices, and therefore might be little more than *ex post* model fitting. Brock, Lakonishok and LeBaron acknowledge this risk, conceding that "a complete remedy for data-snooping biases does not exist." In GP, the rules are combinations of simple operators and historical prices that are 'optimal' by some measure. Because they are constructed using a mechanical process, the population from which they are drawn is much larger than the population of previously successful trading strategies and therefore are free from data-snooping biases.[1]

The use of genetic programming also readily permits the use of a greater variety of data in the formation of trading rules. With the notable exception of Fiess and MacDonald, previous research into technical analysis has also been limited by the exclusive use of closing prices in trading-rule construction. Casual perusal of any technical analysis textbook will confirm the use of open, high, and low prices in addition to closing prices in common techniques.[2]

This paper has five sections, the first section discusses the theory and practice of genetic programming algorithms. The second section introduces the framework used to measure the value of technical trading rules. The third section discusses the data and rolling estimation procedures. The fourth section presents the results of the genetic programs. The fifth section summarizes and offers concluding remarks.

[1] Although genetic programming eliminates data snooping in rule selection, the construction of the genetic programming algorithm itself could still induce data-snooping bias.

[2] Trading volume is also a frequently-used statistic in the formation of technical trading strategies. However, reliable volume data is not available for US equities over the entire time period used in this study, and is therefore not incorporated in the analysis.

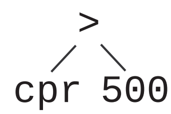

Figure 1: Sample Trading Rule

1 Trading Rules as Genetic Programs

According to Ladd, "…a genetic algorithm creates a set of solutions, testing them against a given problem, and then 'breeding' a new set of solutions based on some measure of success." Genetic programming is the application of genetic algorithms to the creation of computer programs. A genetic program (GP) is best viewed as a hierarchy of simple components, or *nodes*. Figure 1 is a trading rule which specifies that if the closing price on a given date is above 500, a long position should be taken. The *root* node of this rule, $>$, returns a boolean `TRUE` or `FALSE`, which represents holding a long or zero position in an asset, respectively. In this rule, the root node is also a *functional* node in that it operates as a mathematical function, mapping the results of subnodes into its own results. `CPR` and `500` are *terminal* nodes, because they require no arguments.

In principle, nodes may require and return any kind of value. In the above example, both of the terminal nodes return real values, but the functional node returns a logical value. In order to mix real and logical values, a model requires a typing schema. (Zumbach, Pictet, and Masutti) In a weakly-typed model, the syntax is constructed so that logical and real values map smoothly into one another. If an operator that requires a real argument is passed a logical argument, then the logical `TRUE`(`FALSE`) becomes the real value `1`(`0`). If the reverse occurred, a positive (negative) real becomes the logical `TRUE`(`FALSE`).[3] Weakly-typed models always satisfy the closure property of Koza (1992), i.e. that all permissible combinations of operators can be evaluated. In strongly-typed models, the form of logical and real arguments is preserved, and syntactic rules ensure that only the proper type of values are passed. Although it increases complexity, strong typing produces clearer rules and is also more naturally aligned with the

[3] This is only the most common convention. Any syntax that provides such a logical to real mapping will suffice.

method in which compiled languages, such as C or Fortran, treat data.

The choice of nodes in building the genetic programs is similar to those used in Neely, Weller, and Dittmar and Allen and Karjalainen. Terminal nodes may be real [-2,2], boolean (`TRUE`,`FALSE`), or return price data, `OPR`, `HPR`, `LPR`, and `CPR` represent the opening, high, low, and closing price. Function nodes can be the arithmetic operations, $+$, $-$, $\times$, and $\div$, boolean operators, `IF-THEN-ELSE`, `AND`, `OR`, `NOT`, inequalities, $<$, $>$, square, square root, and the 1-norm (distance). Additionally, four functions are included that operate on lagged data, each of which requires two arguments, a data series (`OPR`, `HPR`, `LPR`, or `CPR`) and a real value, k, which indicates the number of prior observations over which to operate. `LAG` returns the k-period lagged price, `MIN` and `MAX` return the minimum and maximum values over the k previous periods, and `AVG` returns the k-period average. NWD and AK also include `POS`, which indicates the evaluation of the rule in the previous period, which is not used here, as it increases the computational time required for evaluation by an order of magnitude, and it is not obvious that it should play a role in the formation of technical trading rules.

The mapping of genetic programs into technical trading rules is relatively simple. The root node of the rules is required to return a logical value when evaluated. This logical value maps into either long/short, long/neutral, or short/neutral, but the binary nature of the root node precludes more complicated results, such as long/neutral/short. For this reason, when programs are generated below, both long/neutral and long/short rules are generated.

2 Evolving Optimal Trading Rules

The primary contribution of genetic algorithms lies in the methods' ability to solve highly non-linear problems using techniques modelled on biological processes. The 'evolution' of a genetic program begins with the random generation of a population of rules such as that in figure 1. Typically, these initial rules are limited in both levels and number of nodes. Each rule is evaluated based upon some fitness criteria. Then, a second generation of rules is created from the first, based upon the members' fitness, using operations that are analogous to evolutionary processes, i.e. where the fittest members of the population are more likely to reproduce. This

Table 1: Nesting of Common Technical Indicators within Functional/Terminal Node Sets.

Technical Indicator	Nested in Node Sets AK	Nested in Node Sets Current
Trend Lines	-	-
Support/Resistance	X*	X
Channel Line	-	-
Percentage Retracements	X	X
Speedlines	X	X
Gaps	-	X
Head and Shoulders	-	-
Double Tops/Bottoms	-	-
Triangles	-	-
Moving Average	X	X
Envelopes	X	X
Bollinger Bands	X*	X
Momentum	X	X
RSI	X	X
Stochastics	-	X
% R	-	X
MACD	X	X
Candlesticks	-	X

* Although these indicators can be based only on closing prices, high and low prices are most commonly used.

Figure 2: Recombination of Parent Rules into Child Rules.

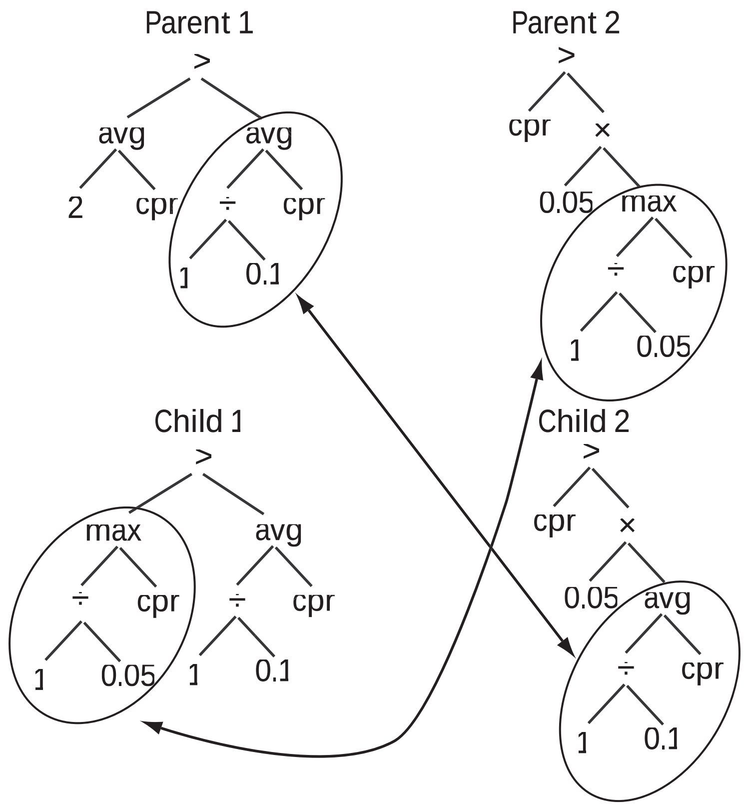

second generation is evaluated, from which a third is generated, and so on, until some termination criteria is met

Only the two most common evolutionary operators in genetic programming, reproduction and recombination (Koza), are used in this study. Reproduction copies a rule, unaltered, from the parent generation into the child generation. Recombination is the 'swapping' of sub-trees of two parent rules to create two new children. In figure 2, a sub-tree of each parent rule is chosen and interchanged to create two new child rules. In strongly-typed models, the sub-trees must be chosen so that both return the same kind of argument (logical or real).

As in Koza, succeeding populations are generated through 10% reproduction, and 90% recombination. The parent rules used in the genetic operations, instead of being chosen randomly, are chosen with probability proportional to their fitness rank in the population, as per the suggestion of Whitley. The probability of rule i being chosen as a parent for reproduction or recombination is

$$p_i = \frac{r_i^3}{\sum_j r_j^3} \tag{1}$$

where r_i is the fitness rank of rule i in the population of N rules, i.e. for the fittest rule, $r_i = N$, and for the least fit rule, $r_i = 1$.

To avoid generating rules that are too specifically tailored to the dataset used in their creation, the data for the generation of each rule is divided into training, selection, and testing periods of two years, as in Allen and Karjalainen. Each generation of rules is evolved using the fitness measures obtained by evaluation with the training data. The fittest rule of the generation is then evaluated using the selection data and stored. A new generation is created and evaluated, and the fittest rule from this generation is evaluated using the selection data. If it is fitter than the fittest rule of the previous generation, it is stored, otherwise, the prior rule is retained. The GP terminates after 50 generations without an improvement in the stored rule, or after 200 generations. When the GP process terminates, one trading rule has been generated for each training/selection/testing period.

Because GP cannot guarantee convergence, either locally or globally, the quality of a solution is a monotonic function of its computational cost; as larger populations of larger rules are allowed to evolve longer, the probability of convergence increases. Balancing this need is the time required for estimation. The population size is 1000 rules, each of which is constrained to 100 nodes. In the initial rule generation, the rules are constrained to be no more than seven levels deep, but in recombination, the rules can grow to be 16 levels deep. To further improve the results, ten optimizations are performed over each set of training/selection data, differing only in the seed value to the random number generator, and the best rule of ten is used in subsequent testing.

3 Valuing Technical Analysis

Previous research in technical analysis has used a variety of measures of the merit of technical analysis. Most studies up to and including BLL focus on the ability of technical trading rules to deliver returns in excess of the market or risk-free return. More recently, Neely, Weller and Dittmar, *inter alia*, adjusted returns for risk via regression of trading returns on the market return to identify premia garnered for bearing systematic risk. Neely used four measures to adjust trading returns for risk and/or correlation to the market return: changes in the Sharpe Ratio of a portfolio when a traded asset is added, Sweeney and Lee's X^* statistic, which adjusts returns for risk compared to a random portfolio, and the X_{eff} statistic of Dacoragna, *et al.* which measures the utility that the trading strategy provides over a weighted average of return horizons.

This study also utilizes multiple measures to assess the usefulness of technical trading. Mean returns are used for their simplicity. Sharpe ratios are also used to measure the risk-adjusted return. This article also adapts the methodology of Fleming, Kirby and Ostdiek to estimate the willingness of investors to pay for technical analysis.

The simplest measure of a strategy's usefulness is its ability to generate returns in excess of a competing asset. Until the 1990's, the ability to generate large mean returns net of transactions costs was the benchmark for technical analysis research. For a technical trading program in which $I(t)$ is an indicator with value of 1 if long, 0 if neutral, or -1 if short, $I^*(t) = 1 - \text{abs}(I(t))$, the total return to a trading strategy, $\pi = \sum \tilde{r}_t$ where

$$\tilde{r}_t = (r_t - r_t^{\text{rf}})I(t) + \left[I(t) \oplus I(t-1) \times \frac{1-c}{1+c}\right] \tag{2}$$

and $r_t = \log(p_{t+1}/p_t)$, the continuously compounded return, n is the number of transactions, c is the transactions cost, $\oplus$ is the exclusive-or operator, and r_t^{rf} is the risk-free rate. In the results below, π is reported as the simple annualized return. Based upon the results of Sweeney (1988) and AK, one-way transactions costs of 0.25% for out-of-sample testing are used. Neely, Weller, Dittmar (1997) suggest using higher transactions costs in training and selection to reduce the danger of over-fitting, and so one-way transactions costs of 1% are used in these periods.

Ignoring risk in evaluating trading rules is clearly suspect. AK note that "Even though the rules do not lead to higher absolute returns than a buy and hold strategy, the reduced volatility might still make them attractive to some investors on a risk-adjusted basis." Neely performs an analysis similar to AK, but includes fitness measures which account for risk, such as the Sharpe ratio. Although it fails to incorporate correlation to the market return, the Sharpe ratio is the benchmark measure of risk-adjusted returns. The Sharpe ratio of a trading strategy, $\psi = \text{Mean}(\tilde{r}_t)/\text{StdDev}(\tilde{r}_t)$. By convention, the data are annualized before the mean and standard deviations are computed.

The final measure is drawn from Fleming, Kirby and Ostdiek, who use it to measure the value of volatility-timing strategies. In a departure from previous research, we estimate the fee a risk-averse investor would pay for a technically-managed asset. Not only does this result in an intuitive, economic measure of the value of technical analysis in the investment decision, it is a direct test of the oft-postulated view that technical analysis itself may represent some form of volatility-timing strategy.

If quadratic utility is viewed as a second-order approximation to the investor's utility function, and the investor is assumed to have risk aversion represented by γ, which is invariant to wealth, then the investor's utility function in investment returns can be represented as

$$U(\tilde{r}|\gamma) = \sum_{t=1}^{T} \tilde{r}_t - \frac{\gamma}{2(1+\gamma)}\tilde{r}_t^2 \tag{3}$$

for $\tilde{r}_t$ the nominal return net of transactions cost in time t and with initial wealth normalized to unity.

The goal of this study is to estimate the level of fee, expressed in basis points subtracted from the dynamic portfolio and denoted ϕ, that would leave the investor indifferent between the static and dynamic portfolios, i.e. the value of ϕ that solves

$$\sum_{t=1}^{T}(\tilde{r}_{d,t} - \phi) - \frac{\gamma}{2(1+\gamma)}(\tilde{r}_{d,t} - \phi)^2 = \sum_{t=1}^{T} \tilde{r}_{s,t} - \frac{\gamma}{2(1+\gamma)}\tilde{r}_{s,t}^2 \tag{4}$$

where $R_{s,t}$ and $R_{d,t}$ are the returns on the static and dynamic portfolios at time t,

Using equation 4, ϕ is estimated for the case of the static portfolio 100% invested in the market return, and the dynamic portfolio is a one asset portfolio 100% invested in the technically-traded market return.

$$\tilde{r}_{s,t} = r_t \quad (5)$$

$$\tilde{r}_{d,t} = r_t I(t) + I^*(t) r_t^{\text{rf}} + \left[I(t) \oplus I(t-1) \times \frac{1-c}{1+c} \right] \quad (6)$$

Note that the product $I^*(t)r_t^{\text{rf}}$ is added to compute ϕ, as the utility measure is a function of nominal returns.

4 Data and Estimation

This study uses the open, high, low, and closing prices of the Dow Jones Industrial Average from 1 January 1931 through 31 December 2001 from the Bridge Database. Interest rates are drawn from the Fama Risk-Free Rate series. Because equity prices are non-stationary over the sample, they are normalized on the 100-day moving average of closing prices. This results in all prices being of the same order of magnitude, as well as preserving the relationship of open, high, low, and closing prices.

Two years of data were used for each of the training, selection and testing periods. The training periods start every two years, so that none of the training, selection, or testing periods overlap. This results in 33 rules being generated for each fitness criteria.

5 Results

Table 2 reports summary statistics for the fitness measures of the optimal rules using selection and out-of-sample data. Even though all of fitness measures for the selection data-based results must be non-negative[4], none of the mean values are more than one standard deviation from zero. The standard deviations and extrema indicate that the rule performance is highly variable. The number of rules which result in trading vary from 18 for ψ/L/S to 9 for $\phi, \gamma = 10$/L/S. The out-of-sample results are very different. Only one of the fitness criteria generates more rules

[4]Rules which are unable to generate a positive fitness value on the selection data are replaced by a rule which is always long.

with positive fitness than negative, and only two fitness measures have means greater than zero.

The rules themselves also vary significantly, though the summary statistics reported in table 2 obscure this fact. Most of the rules have no intuitive meaning. Equation 7 is one example; this rule is long if

$$0 < |(\texttt{OPR} \times 0.4113) - |\texttt{CPR} - 0.5420| \quad | \tag{7}$$

This rule resulted from the Sharpe ratio fitness criteria, long/neutral positions, and training and selection periods of 1985-86 and 1987-88. Figure 3 charts prices and the positions taken by the rule during the month of October, 1987–during the selection period. During the entire two-year selection period the rule conducted 20 transactions and was long 486 out of 506 days. This trading rule returned 19.5%; the market returned 3.4%. The standard deviation of the dynamic strategy was also lower, 21% rather than 27.5%. This rule was applied to 1989-90 for out-of-sample evaluation, which included August, 1990–the invasion of Kuwait. As expected, the rule is less successful when applied to out-of-sample data. Figure 4 displays the positions taken during August and September of 1990, in which the rule has much greater difficulty predicting the returns of the next period. During the entire out-of-sample period, the rule under-performed the market, returning -0.2% rather than 3.2%, while long 473 days out of 505, using 40 trades.

Tables 3 and 4 report the cumulative performance of the optimal trading rules over the whole data span. These are the same 33 trading rules described in table 2, but the cumulative results assume that an investor updates the trading rule every two years. If no rule is able to produce returns greater than the market in a selection period, then a buy-and-hold strategy is adopted for the succeeding testing period. These results measure the performance of this strategy over the entire 66 years of DJIA data. Total trading returns, with and without transactions costs are reported, as is the market return, along with standard deviations. Also reported is the average return and standard deviation for periods in which the combined trading rule is 'long' and 'neutral/short.' The utility based fee, ϕ is also reported for all of the rules, at $\gamma = 1$ and $\gamma = 10$. In addition, the CAPM α and β are reported, along with the regression's R^2.

Table 3 reports the cumulative performance using selection data–each rule is evaluated using the trading strategy used in its generation. Certain consistencies across the results are immediately noticeable. As should be expected from results generated with in-sample data, the results in table 3 are very positive. From comparison of the trading return to the market return, and the long return to the neutral/short return, the rules are clearly able to distinguish between periods of high and low returns. When transactions costs are omitted, seven of the eight strategies are able to generate above-market returns, with transactions costs, the number is only reduced to six. Of the rules, all are able to reduce the variability of returns beneath the market, although some only by the barest of margins. When evaluated by the level of ϕ, all of the cumulative rules are able to improve on the buy-and-hold strategy for both $\gamma = 1$ and $\gamma = 10$.

The results reported in table 4, computed with out-of-sample data, are less impressive. None of the dynamic strategies produce returns greater than the buy-and-hold strategy after transactions costs, although two are quite close. Omitting transactions costs, three strategies produce above-market returns. Five of the strategies remain able to distinguish between periods of high and low returns, and in all cases, the higher returns also exhibit higher variability.

The return-maximizing (π) long/neutral strategy and the Sharpe Ratio-maximizing long/short strategies exceed the market return before transactions costs, and nearly match it when transactions costs are included. However, neither significantly reduces the variability of returns, and so both fail to improve on the Sharpe Ratio of the market. Each of the rules generates a positive ϕ at $\gamma = 1$, of 96 basis points for π, L/N, and 70 basis points for ψ, L/S. Both of these rules are highly correlated to the market asset, with β of .93 and .82. This is not surprising; these rules are long 90% and 88% of the time, respectively.

The out-of-sample results continue to provide evidence that technical methods can discriminate between periods of high and low returns, as the returns on long days remain higher than those of neutral/short days. Unlike the selection results, in which five specifications had lower volatility on long days than short days, the variance on the short/neutral days is lower for all specifications.

If technical methods are viewed as volatility-timing strategies instead of excess-return generating strategies, then it is natural to examine the impact of these strategies on the mean-variance frontier. As in the calculation of ϕ, investors can choose between two risky assets, the market portfolio and the technically-traded market portfolio, in addition to a risk-free asset. Portfolio weights are chosen to

$$\min_{\boldsymbol{\omega}} \boldsymbol{\omega}' \Sigma \boldsymbol{\omega} \tag{8}$$

subject to

$$\boldsymbol{\omega}' + (1 - \boldsymbol{\omega}'\boldsymbol{\iota}) r^{\mathrm{rf}} = \hat{r} \tag{9}$$

where Σ is the covariance matrix of asset returns, $\boldsymbol{\omega}$ are the asset weights, r^{rf} is the risk-free rate, $\hat{r}$ is the target rate, and $\boldsymbol{\iota}$ is a conformable vector of ones. The solution is

$$\boldsymbol{\omega} = \frac{\left(\hat{r} - r^{\mathrm{rf}}\right) \Sigma^{-1} \left(\boldsymbol{r} - r^{\mathrm{rf}}\boldsymbol{\iota}\right)}{\left(\boldsymbol{r} - r^{\mathrm{rf}}\boldsymbol{\iota}\right)' \Sigma^{-1} \left(\boldsymbol{r} - r^{\mathrm{rf}}\boldsymbol{\iota}\right)} \tag{10}$$

where $\boldsymbol{r}$ is a vector of asset returns. For this analysis, r^{rf}, $\hat{r}$, and $\boldsymbol{r}$ are the mean risk-free rate, the mean market return before transactions costs, and the mean asset returns over the period 1931-2000, respectively.[5]

Table 5 reports the weights for each rule specification, as well as return and volatility results for the out-of-sample data. These results demonstrate that technical strategies have the potential to offer significant improvements in portfolio performance. The value of optimally-weighted portfolios is inversely related to the correlation of the traded asset to the market portfolio. Creating optimally-weighted portfolios marginally improves the performance based upon the first three fitness measures, π, ψ, and $\phi, \gamma = 1$. The variance of π, L/S is greatly reduced, to 8.7% from 14.6% for the simple portfolios. More interesting are the results of forming optimal portfolios with the $\phi, \gamma = 10$ strategy. After transactions costs, the returns of the L/N and L/S rules are 2.396% and 2.518%, but the standard deviations are only 6.949% and 7.539%. The Sharpe Ratio of the market over this time period is 0.2231, but the Sharpe Ratios for these two portfolios are 0.3448 and 0.3340, respectively–a substantial improvement.

[5]Because the weights are being created using *ex post* rules and returns, these weighted results do constitute data-snooping. Caveat lector.

6 Conclusion

Technical analysis played a prominent role in financial markets decades before the birth of mean-variance analysis or Sharpe Ratios, and it remains in widespread use today. In spite of this, the use of technical methods to predict future prices remains in conflict with weak-form market efficiency, and there remain no theoretical underpinnings to technical analysis. If, instead of increased first-moment returns, technical analysis is viewed as a method of increasing risk-adjusted returns, there is no longer a conflict with weak-form efficiency.

The results in this paper indicate that this may be a plausible explanation of technical analysis in financial markets. In the presence of trading costs, the rules generated in this study cannot consistently generate returns in excess of the market. This study demonstrates that these rules can consistently discriminate between periods of high and low returns and volatilities in out-of-sample data. Rules can that are free of data-snooping bias can be constructed that significantly increase risk-adjusted returns, as measured by the fee that a risk-averse investor would pay to own the portfolio. Some evidence is also introduced that suggests that these rules may be even more valuable when they are part of an optimally-weighted portfolio.

The results of this study must be considered in the context of its limitations. Genetic programming, used in order to eliminate potential data-snooping biases common in technical analysis research, cannot guarantee convergence to global minima or maxima. Therefore these results can only be properly interpreted as lower bounds of GP performance in this application. The optimally-weighted portfolio returns reported in table 5 are constructed using *ex post* optimal weights. Finally, these results are contingent on the ability to execute at the daily closing price. The last two points are prime opportunities to extend this research, by using *ex ante* optimal weights and assessing the impact of executing at the following day's open price. Inclusion of other conditioning information, such as interest rates or other domestic or European index levels are further possible avenues of research.

Table 2: Optimal Technical Trading Rules: Summary Statistics

	π		ψ		$\phi, \gamma = 1$		$\phi, \gamma = 10$	
	L/N	L/S	L/N	L/S	L/N	L/S	L/N	L/S
Selection Data								
Mean	5.395	4.103	0.442	0.375	2.264	4.758	9.884	4.261
Std. Dev.	7.570	5.875	0.605	0.523	4.304	10.198	15.035	11.952
Max	29.149	20.892	2.120	1.974	16.292	41.952	64.912	63.448
Min	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
$n > 0$	16	15	17	18	13	13	18	9
$n < 0$	0	0	0	0	0	0	0	0
Average Nodes	46.823	53.133	42.105	46.777	35.933	40.615	46.428	45.444
Average Levels	7.823	6.866	7.421	7.944	7.200	7.538	6.761	7.333
Out of Sample Data								
Mean	0.484	-2.686	-0.001	-0.063	-0.367	-4.011	0.633	-4.648
Std. Dev.	4.243	7.292	0.611	0.515	4.192	11.842	7.110	11.058
Max	16.141	14.643	1.999	1.653	15.209	8.632	23.316	8.501
Min	-11.082	-26.013	-1.625	-1.315	-10.869	-60.256	-14.502	-48.120
$n > 0$	4	2	7	4	4	1	5	1
$n < 0$	4	8	5	5	6	7	9	8
Average Nodes	46.823	53.133	42.105	46.777	35.933	40.615	46.428	45.444
Average Levels	7.823	6.866	7.421	7.944	7.200	7.538	6.761	7.333

The statistics reported refer to the fitness achieved by the 33 technical trading rules.

Table 3: Cumulative Performance of Optimal Technical Trading Rules on Selection Data.

	π		ψ		$\phi, \gamma = 1$		$\phi, \gamma = 10$	
	L/N	L/S	L/N	L/S	L/N	L/S	L/N	L/S
Total Return	3.662	7.559	3.595	5.371	5.404	9.716	5.263	9.268
Total Return*	3.994	8.577	4.148	6.207	5.821	10.789	6.050	9.634
Std. Dev., Total	13.775	15.483	12.856	15.487	13.490	15.479	8.814	15.481
Market Return				4.086				
Market Std. Dev.				15.490				
Return, Long	4.450	7.417	4.821	5.943	8.014	8.676	12.809	8.679
Std. Dev., Long	14.524	15.248	13.837	15.589	15.745	14.828	12.614	15.487
Return, Neutral/Short	1.042	-13.600	-0.319	-7.400	-5.969	-19.600	-3.552	-11.814
Std. Dev. Neutral/Short	22.354	16.812	23.362	14.791	14.747	18.909	17.786	15.459
$\phi, \gamma = 1$	0.520	4.128	0.333	2.158	2.170	6.040	1.506	5.031
$\phi, \gamma = 10$	2.873	4.166	3.676	2.193	4.916	5.557	9.108	5.078
α	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
β	0.790	0.659	0.688	0.762	0.758	0.572	0.324	0.588
R^2	0.790	0.435	0.688	0.581	0.758	0.327	0.324	0.346
% Days in Mkt:	89.948	85.604	86.316	86.980	73.393	85.731	48.742	79.408
Trades	42	62	70	52	52	64	98	22

Trading positions for each of the 33 rules are generated for the selection data and concatenated in order to produce the cumulative results. For this table, 1.0=1%.

Table 4: Cumulative Performance of Optimal Technical Trading Rules on Out-of-Sample Data.

	π		ψ		$\phi, \gamma = 1$		$\phi, \gamma = 10$	
	L/N	L/S	L/N	L/S	L/N	L/S	L/N	L/S
Total Return	3.250	-1.108	2.425	3.130	2.677	-1.155	0.206	-1.471
Total Return*	3.690	0.527	3.521	3.948	3.099	2.857	1.231	-0.326
Std. Dev., Total	14.137	14.684	13.865	14.682	13.000	14.683	11.558	14.684
Market Return	3.275							
Market Std. Dev.	14.682							
Return, Long	4.090	2.197	4.085	4.126	4.294	3.597	2.409	1.843
Std. Dev., Long	14.869	14.935	14.914	14.961	15.257	14.957	16.122	15.320
Return, Neutral/Short	-3.938	10.648	-1.622	-2.591	0.679	1.479	4.231	9.160
Std. Dev. Neutral/Short	12.790	12.922	13.109	12.459	13.038	12.919	12.988	11.855
$\phi, \gamma = 1$	0.958	-3.053	0.114	0.697	0.242	-2.272	-2.344	-3.990
$\phi, \gamma = 10$	1.740	-3.008	1.226	0.763	2.455	-2.244	1.413	-3.940
α	0.000	-0.000	0.000	0.000	0.000	0.000	-0.000	-0.000
β	0.927	0.792	0.891	0.826	0.783	0.778	0.619	0.734
R^2	0.927	0.627	0.891	0.682	0.783	0.605	0.619	0.539
% Days in Mkt:	90.396	86.598	86.417	87.925	72.600	85.676	51.398	79.659
Trades	56	108	140	52	54	262	134	76

Trading positions for each of the 33 rules are generated for the selection data and concatenated in order to produce the cumulative results. For this table, 1.0=1%

Table 5: Cumulative Performance of *Ex Post* Optimally-Weighted Trading Rules on Out-of-Sample Data.

	π		ψ		$\phi, \gamma = 1$		$\phi, \gamma = 10$	
	L/N	L/S	L/N	L/S	L/N	L/S	L/N	L/S
Total Return	2.717	1.949	2.400	2.858	2.810	1.335	2.396	2.518
Total Return*	3.289	3.293	3.292	3.291	3.277	3.285	3.308	3.311
Std. Dev., Total	13.684	8.732	14.177	13.988	12.632	13.902	6.949	7.539
Market Return				3.275				
Market Std. Dev.				14.682				
$\phi, \gamma = 1$	-0.002	-0.179	0.001	0.096	-0.017	-0.119	0.008	-0.173
$\phi, \gamma = 10$	1.267	6.162	0.665	0.995	2.414	0.815	7.699	7.093
Mkt. Weight	-0.330	0.973	0.207	0.483	1.567	1.266	0.755	0.740
Dyn Weight	1.294	-0.779	0.812	0.513	-1.046	-0.469	-0.860	-0.647

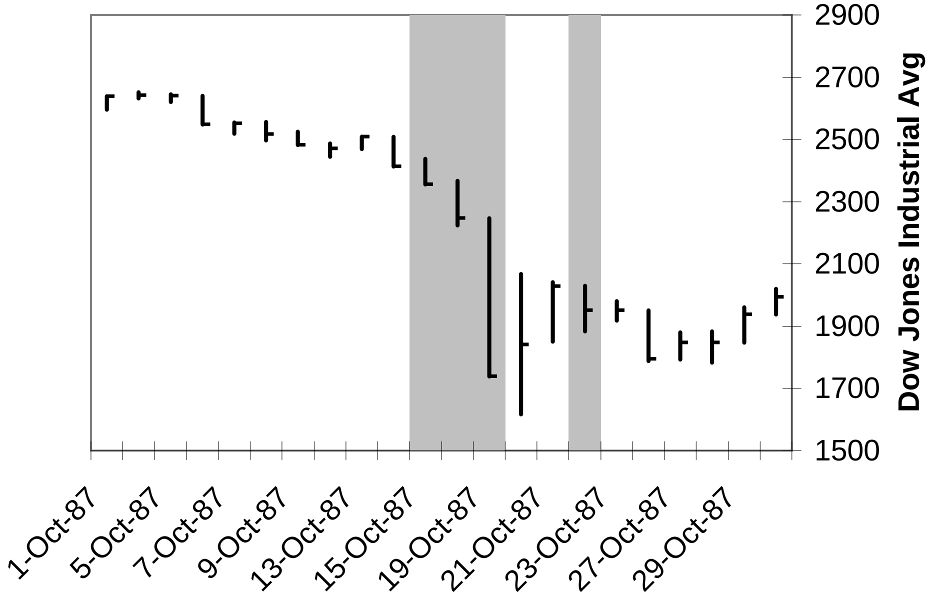

Figure 3: Positions of Selected Trading Rule for the Month of October, 1987.
Note: The portfolio is long except for the shaded regions, which represent neutral positions. These positions result from the rule generated using 1985-86 training data and 1987-88 selection data, therefore these positions are generated 'in-sample,' using the Sharpe Ratio fitness criteria.

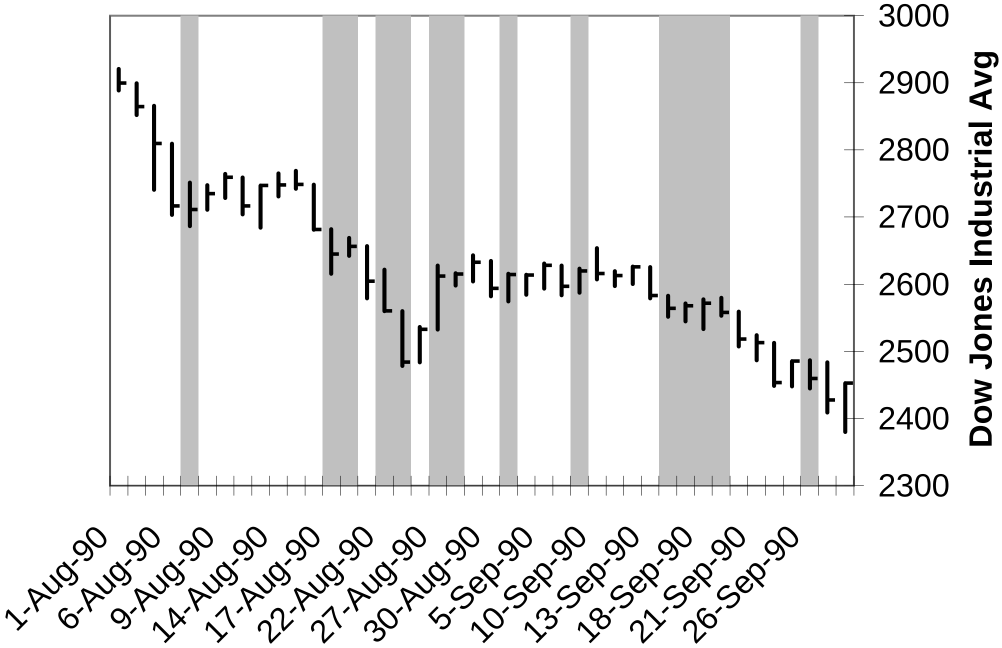

Figure 4: Positions of Selected Trading Rule for the Months of August and September, 1990. Note: The portfolio is long except for the shaded regions, which represent neutral positions. These positions result from the rule generated using 1985-86 training data and 1987-88 selection data, therefore these positions are generated 'out-of-sample,' using the Sharpe Ratio fitness criteria.

References

Alexander, S. S., "Price Movements in Speculative Markets: Trends or Random Walks, No. 2," in Paul Cootner, ed., *The Random Character of Stock Market Changes*, MIT Press, 1964, pp. 338–372.

Dacoragna, M.M., R. Gençay, U.A. Müller, and O.V. Pictet, "Effective Return, Risk Aversion, and Drawdown," *Physica A*, 2001, *289*, 229–248.

Dooley, Michael P. and Jeffrey R. Shafer, "Analysis of Short-Run Exchange Rate Behavior: March 1973 to September 1975," 1976. International Finance Discussion Paper 123, Federal Reserve Board, Washington, D.C.

___ **and** ___ , "Analysis of Short Run exchange rate behavior: March 1973 to November 1981," in D. Bigman and T. Taya, eds., *Exchange Rate and Trade Instability*, Cambridge, MA: Ballinger Publishing, 1983.

Engle, Robert F., "Autoregressive Conditional Heteroscedasticity with Estimates of the Variance of United Kingdom Inflation," *Econometrica*, 1982, *50*, 987–1007.

Fama, Eugene and Marshall Blume, "Filter Rules and Stock Market Trading," *Journal of Business*, 1966, *3*, 226–241.

Fiess, Norbert and Ronald MacDonald, "Technical Analysis in The Foreign Exchange Market: A Cointegration-Based Approach," *Multinational Finance Journal*, 1999, *3* (3), 147–172.

Fleming, Jeff, Chris Kirby, and Barbara Ostdiek, "The Economic Value of Volatility Timing," *Journal of Finance*, 2000, *forthcoming.*

Koza, John R., *Genetic Programming: On the Programming of Computers by Means of Natural Selection*, Cambridge: MIT Press, 1992.

Ladd, Scott Robert, "An Introduction to Genetic Algorithms," http://www.coyotegulch.com 2001.

Lo, Andrew W., Harry Mamaysky, and Jiang Wang, "Foundations of Technical Analysis: Computational Algorithms, Statistical Inference, and Empirical Implementation," *Journal of Finance*, 2000, *55* (4), 1705–1765.

Lui, Yu Hon and David Mole, "The Use of Fundamental and Technical Analyses by foreign exchange dealers: Hong Kong evidence," *Journal of International Money and Finance*, 1998, *17*, 535–545.

Malkiel, Burton G., *A Random Walk Down Wall Street*, 4th ed., New York: Norton, 1985.

Menkhoff, L., "Examining the use of technical currency analysis," *International Journal of Finance and Economics*, 1997, *2*, 307–318.

Neely, Christopher J., "Risk-Adjusted, Ex Ante, Optimal Technical Trading Rules in Equity Markets," Federal Reserve Bank of St. Louis Working Paper 1999-015D 2001.

——— **, Paul Weller, and Rob Dittmar**, "Is technical analysis in the foreign exchange market profitable? A genetic programming approach," *Journal of Financial and Quantitative Analysis*, 1997, *32* (4), 405–426.

Oberlechner, Thomas, "Importance of Technical and Fundamental Analysis in the European Foreign Exchange Market," *International Journal of Finance and Economics*, 2001, *6*, 81–93.

Osler, Carol, "Identifying Noise Traders: The Head-and-Shoulders Pattern in U.S. Equities," 1998. Staff Report No. 42, Federal Reserve Bank of New York.

——— **and Kevin Chang**, "Head and shoulders: Not just a flaky pattern," 1995. Staff Report No. 4, Federal Reserve Bank of New York.

Sweeney, Richard J., "Beating the foreign exchange market," *Journal of Finance*, 1986, *41* (1), 163–182.

———, "Some New Filter Rule Tests: Methods and Results," *Journal of Financial and Quantitative Analysis*, 1988, *23* (3), 285–300.

——— **and E.J.Q. Lee**, in "International Dimensions of Securities and Currency Markets," Vol. 4 of *Advances in Financial Planning and Forecasting*, Greenwich, CT: JAI Press, 1990, pp. 55–79.

Taylor, Mark P. and Helen Allen, "The Use of Technical Analysis in the Foreign Exchange Market," *Journal of International Money and Finance*, 1992, *11* (3), 304–314.

Zumbach, Gilles, Olivier V. Pictet, and Oliver Masutti, "Genetic Programming with Syntactic Restrictions applied to Financial Volatility Forecasting," April 2001. Working Paper, Olsen & Associates Research Institute for Applied Economics.