



Stock market trading rule discovery using technical charting heuristics

William Leigh^{a,*}, Naval Modani^b, Russell Purvis^c, Tom Roberts^d

^aDepartment of Management Information Systems, College of Business Administration, University of Central Florida,
P.O. Box 161400, Orlando, FL 32816-1400, USA

^bDepartment of Finance, College of Business Administration, University of Central Florida, Orlando, FL 32816, USA

^cDepartment of Management, College of Business and Public Affairs, Clemson University, Clemson, SC 29634-1305, USA

^dDepartment of Accounting and Information Systems, School of Business, The University of Kansas, 345N Summerfield Hall,
1300 Sunnyside Avenue, Lawrence, KS 66045-7585, USA

Abstract

In this case study in knowledge engineering and data mining, we implement a recognizer for two variations of the ‘bull flag’ technical charting heuristic and use this recognizer to discover trading rules on the NYSE Composite Index. Out-of-sample results indicate that these rules are effective. © 2002 Elsevier Science Ltd. All rights reserved.

Keywords: Stock market; Technical analysis; Financial expert system

1. Introduction

Many Wall Street financial advisors practice technical analysis. Charles Dow developed the original Dow Theory for technical analysis in 1884, and Edwards and Magee (1997) wrote a modern version. Charting, a technique of technical analysis, compares market price and volume history to archetypal chart patterns and predicts future price behavior based on the degree of match. The charting technique may be operationalized through trading rules of the form:

If charting pattern X is identified in the previous N trading days, then buy; and sell on the Y th trading day after that.
If charting pattern X is identified in the previous N trading days, then sell.

Obviously the difficult part is identifying ‘charting pattern X ’.

Charting is rarely tested in the academic literature. The practitioner literature frequently includes articles implying that trading success may be achieved with charting, for example, see Martinelli and Hyman (1998).

We work with the New York Stock Exchange Composite Index, a broad-based composite index, so that the overall average in the period is equivalent to the return from a buy-and-hold or a random-selection trading strategy. The aca-

demically accepted efficient markets hypothesis (Haugen, 1997) implies that buy-and-hold or, equivalently, random selection are optimal trading strategies (and that the efforts of Wall Street financial analysts are futile). Nobel prize-winner Paul Samuelson (1965, p. 44) claims to prove that, “...there is no way of making an expected profit by extrapolating past changes in the futures price, by chart or any esoteric devices of magic or mathematics. The market quotation already contains in itself all that can be known about the future, and in that sense it has discounted future contingencies as much as is humanly possible.” Thus, the finding of results using our charting method which are significantly better than the overall average for the period constitutes a failure to confirm our implicit null hypothesis, which is the efficient markets hypothesis.

Other expert system and knowledge engineering efforts to accomplish stock market prediction, all of which ignore the efficient markets hypothesis, are numerous. Many employ technical analysis momentum heuristics, and some use neural networks to learn patterns, but none that we have found use charting heuristics directly. Examples include: neural networks learning from past price history (Gencay, 1998); neural network using knowledge of real-world events (Kohara, Ishikawa, Fukuhara, & Nakamura, 1997); fuzzy expert system combining several sources of information (Lee & Kim, 1995); decision support system with influence diagram (Poh, 2000); rough sets to extract trading rules from price history (Kim & Han, 2001); data mining employing signal processing techniques (Last, Klein, & Kandel, 2001); grey theory and fuzzification (Wang, 2001); and rules

* Corresponding author. Tel.: +1-407-823-3173; fax: +1-407-823-2389.
E-mail address: leigh@pegasus.cc.ucf.edu (W. Leigh).

0.5		-1	-1	-1	-1	-1	-1	-1	
1	0.5		-0.5	-1	-1	-1	-1	-0.5	
1	1	0.5		-0.5	-0.5	-0.5	-0.5		0.5
0.5	1	1	0.5		-0.5	-0.5	-0.5		1
	0.5	1	1	0.5				0.5	1
		0.5	1	1	0.5			1	1
-0.5			0.5	1	1	0.5	0.5	1	1
-0.5	-1			0.5	1	1	1	1	
-1	-1	-1	-0.5		0.5	1	1		-2
-1	-1	-1	-1	-0.5		0.5	0.5	-2	-2.5

Fig. 1. Template grid for chart pattern ‘Bull Flag 1’. The charting pattern is coded using weights in a 10 × 10 grid. The consolidation part of the pattern, in the first seven columns, is downward-sloping. The breakout part of the pattern is in the last three columns. A blank cell signifies a weight of 0.0.

(Armstrong & Collopy, 1993) or agents (Skouras, 2001) using technical heuristics combined with other sorts of knowledge in an expert system framework.

2. Identifying the bull flag by template match

Charting, the sort of technical analysis that we use, is based on the recognition of certain graphical patterns in price and/or volume time series data. This work concentrates on one charting pattern, the *bull flag*. The definition of flag from Downes and Goodman (1998):

Flag: Technical chart pattern resembling a flag shaped like a parallelogram with masts on either side, showing a consolidation within a trend. It results from price fluctuations within a narrow range, both preceded and followed by sharp rises or declines.

A bull flag pattern is a horizontal or sloping flag of *consolidation* followed by a sharp rise in the positive direction, the *breakout*.

The template grids we use to identify the occurrence of two variations of the bull flag charting pattern are shown in Figs. 1 and 2. Each pattern is represented using a 10 × 10 grid with weights in the cells. (A blank cell signifies a weight of zero.) The weighting is used to define areas in the template grid to represent the consolidation and the breakout portions of this bull flag pattern.

We fit these templates to price history using a 120 trading day window. Within each fitting window, we ‘windsorize’ (Roberts, 1995, p. 150) to remove the worst noise by replacing every observation which is beyond two standard deviations from the mean (of the 120 index values in the window) with the respective two standard deviation boundary value.

In each application of the matching/fitting process, the 10 × 10 template grid is matched to an image of the time series of price data to get a template fit value for one trading day at a time, with the leftmost time series data point being the value for the trading day which precedes the current day by 1 day less than the fitting window width in trading days,

-1	-1	-1	-1	-1	-1	-1	-1	0.25	0.8	2
-1	-1	-1	-1	-1	-1	-1	-1	0.25	0.8	2
								0.25	0.8	
								0.25		
										-1
									-0.75	-1
1	1	1	1	1	1	1	1	-0.5	-0.75	-1
0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	-0.5	-0.75	-1

Fig. 2. Template grid for chart pattern ‘Bull Flag 2’. The consolidation part of the pattern, in the first seven columns, is horizontal. The breakout, the last three columns, is abrupt. A blank cell signifies a weight of 0.0.

so for our fitting window with of 120 trading days, the leftmost time series data point in a window precedes the current day by 119 trading days. Values for the earliest 10% of the trading days in the image (12 days of the 120) are mapped to the first column of an image grid, values for the next-to-earliest 10% of the trading days are mapped to the second column of the image grid, and so on, until the most recent 10% of the trading days are mapped to the rightmost column of the image grid.

The vertical translation process to put the time series data into the image grid is adaptive: the highest price value in the fitting window is made to correspond with the top of the image grid, and the lowest price value in the window is made to correspond with the bottom of the image grid. The price values in between the lowest and highest are allocated to the intervening cells accordingly.

To compute a total fit for a single trading day, the percentage of values that fall in each cell of a column in the image grid is multiplied by the weight in the corresponding cell of the bull flag template grid (a cross-correlation computation). For example, there will be price data for 12 trading days represented in each column of a single 120 trading day window. If all 12 of these trading days have price values which are in the lowest 10% of the difference between the lowest and highest prices in the fitting window for the day being fitted, then 100% (12 values out of a total of 12 in the column) will be the value in the lowest cell of the 10 cells in the column. If this column is the leftmost of the columns in the window, then this 100% will be multiplied by the value in the corresponding cell in the bull flag template grid (which is the one in the lowest left-hand corner), which has the value of -1.0 in Fig. 1, to result in a cell fit value of -1.0 × 100% = -1.0. This is done for the 10 cells in the column and summed, resulting in a fit value for the column of -1.0, since, in this case, there will be 0.0% in the other nine cells of the column.

In this way, 10 column fit values for price are computed for a trading day. Summing all 10 column total values for a trading day results in a total fit for the trading day. This process is an example of *template matching* as described in Duda and Hart (1973). The maximum of this cross-correlation




Fig. 3. Closing values of the New York Stock Exchange Composite Index for 4748 trading days in the period of the study, August 6, 1980–June 8, 1999. This is clearly a bull market period.

computation occurs when the image grid most closely matches the template grid.

We work with the data sample comprised of the closing price values for the New York Stock Exchange Composite Index for the 4748 trading days in the period, August 6, 1980–June 8, 1999 (and the 119 days preceding the sample period and the 99 days following that period, in order to accommodate the fitting windows and forecasting horizons for all of the trading days in the sample period, but these preceding and following days are not included in any of the results statistics). Fig. 3 shows the time series of closing prices for this period. This is clearly a bull market period.

The best total fit found for the charting pattern Bull Flag 1 for the 120 trading days width window in the period of study occurred for 5/10/1982 (ending date). Fig. 4 shows the values of the NYSE Composite Index for this 120 period and the following 100 day period. (Figs. 3 and 4 together can support the conjecture that this best fit of the 120 trading day width marks the beginning of the great bull market of the 1980s and 1990s. 5/10/1982 is the 444th trading day in the period of study.)

Implementation is accomplished using a spreadsheet data analysis tool.

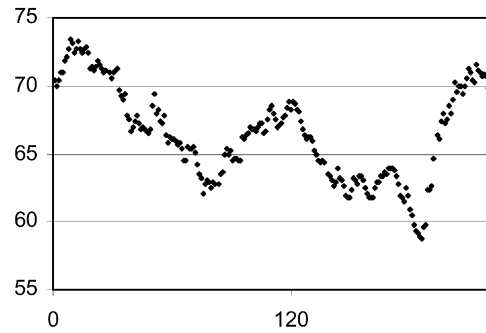


Fig. 4. Best total fit found for Bull Flag 1 in 120 trading day width window in period of study occurred for 5/10/1982. NYSE Composite Index is the vertical axis. The first 120 trading days plotted here end on 5/10/1982 and exhibit high congruence with the Bull Flag 1 pattern. Index values for the following 100 days are supplied to show the ensuing price activity.

3. Initial data mining

We use the first 500 trading days of the period of study for a mining and training sample. We compute total fit values for all 500 trading days in the period of the study, for a fitting window width of 120 trading days. We sort these results in descending order by total fit value and report average profit percentages for the top 5, 10, 20, and 30%, that is, for the 95th percentile, and for the 90th, 80th, and 70th deciles. We determine average profit change percentages for a forecast horizon of 100 trading days. We compare these average profit percentages against the overall average profit percentage for all of the trading days in the period of the study (which value may be found on each of the following tables as the profit percentage for all trading days greater than or equal to the 0th percentile, found in the bottom row of Table 1.) The statistical test used is the one-tailed, two-sample unequal variance (heteroscedastic) *t*-test.

Table 1 contains results found with the first 500 days in our study sample. The use of Bull Flag 1 for 95th percentile

Table 1

Results of applying a 120-day-wide fitting window to only the first 500 days in the period of study. (Profit figures are for a 100 trading day period. ‘Fit ptile’ refers to the percentile of the total fit values. ‘Prof% all \geq ptile’ refers to the average profit percentage difference between the average computed for the trading days with fit values greater than or equal to the fit value for the percentile and the overall average for the period of the study. ‘*P* value’ is the probability value, or significance level, from the *t*-test statistical computation comparing the average profit percentage value in the subset with the overall average profit value, which is shown as the 0th percentile ‘Prof% all \geq ptile’)

Fit ptile	Prof% all \geq ptile	Diff. (%)	<i>P</i> value	Decile prof%	Decile <i>P</i> value
<i>Bull Flag 1</i>					
95	6.62	6.70	0.0000		
90	2.61	2.69	0.0296	2.61	0.0296
80	-0.84	-0.76	0.2534	-4.41	0.0033
70	-2.03	-1.95	0.0216	-4.55	0.0015
0	-0.08				
<i>Bull Flag 2</i>					
95	5.04	5.12	0.0000		
90	3.24	3.32	0.0000	3.24	0.0000
80	1.80	1.88	0.0009	0.31	0.2691
70	-0.50	-0.42	0.2592	-5.08	0.0000




Fig. 5. Vertical axis is 100 trading day horizon profitability; horizontal axis is normalized volume window height; each indicated buy trading day in the learning sample, the first 500 days, is plotted. These charts reveal that the addition of consideration of the normalized volume window height may be valuable in improving the results of rules 1 and 2, and the effect appears to be more pronounced in the case of Rule 1.

fit values and better as a buy signal results in a 6.70% improvement over buying at random (which is equivalent to the overall average profit increase for every day in the period) with a 100 trading day forecast horizon. The use of Bull Flag 2 for 95th percentile fit values and better as a buy signal results in a 5.12% improvement over buying at random for holding for 100 trading days.

4. Trading rule development

The results in Table 1 reveal that excess profits might be realized through a trading rule using this approach. Trading rules which might be derived from these data mining results include:

Rule 1: If charting pattern Bull Flag 1 is identified in the

previous 120 days with a fit value \geq the 95th percentile fit value for the previous trading days, then buy; and sell on the 100th trading day after that.

Rule 2: If charting pattern Bull Flag 2 is identified in the previous 120 days with a fit value \geq the 95th percentile fit value for the previous trading days, then buy; and sell on the 100th trading day after that.

Table 2 shows the results from applying these two rules to the trading days in our test sample, which is comprised of the days remaining in the period of the study after the first 500 days are removed. The 95th percentile fit value is computed from all days preceding the day of the fitting. So, for the first day in this test, day 501 of the period of the study, the first 500 days of the period are used in the computation of the 95th percentile fit value. For the last day, day 4748, the previous 4747 trading days in the period of the study are used.

The overall average 100 trading day increase in the index value for all 4248 trading days in the test period used for Table 2 was 5.5%. The average 100 trading day increase in the index for the days for which a buy was indicated by: Rule 1 was 11.5% and the *t*-test probability value comparing this average with the overall average is 0.000000000000007; Rule 2 was 7.8% and the *t*-test probability value comparing this average with the overall average is 0.000000006. Rule 1 indicated purchases on 124 trading days in the study period. Rule 2 was less selective, indicating purchase 220 times.

Table 2

Results of applying trading rules 1 and 2 to days in study period after the first 500 (*N* refers to the number of days for which buying was recommended. ‘All’ refers to all trading days in the interval)

	All		Rule 1		Rule 2	
	Profit (%)	<i>N</i>	Profit (%)	<i>N</i>	Profit (%)	<i>N</i>
1982	19.2	108	19.8	5	13.9	36
1983	2.1	251		0	13.59	2
1984	5.1	252		0	9.0	2
1985	9.9	250	20.3	10	6.4	42
1986	9.0	252		0		0
1987	-2.6	253		0	8.9	43
1988	5.4	253		0		0
1989	4.9	252		0	2.4	12
1990	3.1	253	15.9	25		0
1991	5.7	252	21.1	9	1.7	13
1992	3.0	254	2.5	2	-0.1	12
1993	1.9	253		0		0
1994	1.8	251	0.2	31		0
1995	10.6	251	13.5	22	10.2	33
1996	7.6	253		0	2.3	17
1997	12.1	252		0	16.3	7
1998	4.8	252	11.6	20	-15.0	1
1999	1.8	106		0		0
	5.5	4248	11.5	124	7.8	220

5. Addition of trading volume to rules

Further exploration of the first 500 trading days found the possibility of a relationship between the range in trading volume during the 120 day period of a fitting window and the subsequent price value. Fig. 5 shows this relationship for the indicated buy trading days in the 500 day learning sample. Each point in Fig. 5 corresponds to one of the days identified as a buying opportunity by rules 1 and 2 in the first 500 days of the period of this study, which was used for data mining previously. The vertical axis in Fig. 5 is the 100 trading day horizon profitability for those buy days. The

horizontal axis is a normalized volume window height. That normalized volume window height is computed by first windsorizing (in the 120 day window) to two standard deviations the 120 trading share volume values for the trading days in the fitting window, subtracting the lowest windsorized volume from the highest, and then dividing that difference by the trading volume value for the last day in the window, the day of the fitting. It appears that trading days which are identified as buying opportunities by rules 1 and 2 are more likely to be true buying opportunities if they have higher normalized volume window height values.

Examination of Fig. 5 led us to choose a cut-off value for normalized volume window height of 0.75 and to modify Rule 1 and Rule 2:

Rule 1v: If charting pattern Bull Flag 1 is identified in the previous 120 days with a fit value \geq the 95th percentile fit value for the previous trading days AND the normalized height of the volume window for the previous 120 days is greater than 0.75, then buy; and sell on the 100th trading day after that.

Rule 2v: If charting pattern Bull Flag 2 is identified in the previous 120 days with a fit value \geq the 95th percentile fit value for the previous trading days AND the normalized height of the volume window for the previous 120 days is greater than 0.75, then buy; and sell on the 100th trading day after that.

Application of these new rules, in the manner used previously, to the 4248 trading days in the test sample resulted in more selective buying and improved 100 day profitability for the days identified as buying opportunities: 55 buys and 14.0% profitability for Rule 1v and 132 recommended buys and 8.6% for Rule 2v. Respective *t*-test probability values were 0.0000000000 and 0.0000000012.

6. Conclusion

This case study employs classic knowledge engineering methods. We identify a heuristic from a domain of expert knowledge, devise computer algorithms for implementation, identify a performance criterion, obtain sample data, tune the mechanism and learn parameter values on a learn-

ing sample, and validate using a test sample. The results with the test sample imply that the charting heuristics and this implementation have validity. Perhaps this purely empirical paper, which comprises ‘measurement without theory’ (Koopmans, 1947), will hasten the academic acceptance of technical analysis and, especially, the charting heuristics of technical analysis.

References

- Armstrong, J., & Collopy, F. (1993). Causal forces: structuring knowledge for time-series extrapolation. *Journal of Forecasting*, 12, 103–115.
- Downes, J., & Goodman, J. (1998). *Dictionary of finance and investment terms*, (5th ed.). New York: Barron’s Educational Series, Inc.
- Duda, R., & Hart, P. (1973). *Pattern classification and scene analysis*, New York: Wiley.
- Edwards, R., & Magee, J. (1997). *Technical analysis of stock trends*, (7th ed). New York: Amacom.
- Gencay, R. (1998). The predictability of security returns with simple technical trading rules. *Journal of Empirical Finance*, 5, 347–359.
- Haugen, R. (1997). *Modern investment theory*, Upper Saddle River, NJ: Prentice-Hall.
- Kim, K., & Han, I. (2001). The extraction of trading rules from stock market data using rough sets. *Expert Systems*, 18 (4), 194–202.
- Kohara, K., Ishikawa, T., Fukuhara, Y., & Nakamura, Y. (1997). Stock price prediction using prior knowledge and neural networks. *Intelligent Systems in Accounting, Finance and Management*, 6, 11–22.
- Koopmans, T. C. (1947). Measurement without theory. *The Review of Economic Statistics*, 29 (3), 161–172.
- Last, M., Klein, Y., & Kandel, A. (2001). Knowledge discovery in time series databases. *IEEE Transactions on Systems, Man, and Cybernetics—Part B: Cybernetics*, 31 (1), 160–169.
- Lee, K., & Kim, W. (1995). Integration of human knowledge and machine knowledge by using fuzzy post adjustment: its performance in stock market timing prediction. *Expert Systems*, 12 (4), 331–338.
- Martinelli, R., & Hyman, B. (1998). Cup-with-handle and the computerized approach. *Technical Analysis of Stocks and Commodities*, 16 (10), 63–66.
- Poh, K. (2000). An intelligent decision support system for investment analysis. *Knowledge and Information Systems*, 2, 340–358.
- Roberts, T. (1995). *Bounded rationality in macroeconomics*, Oxford: Oxford University Press.
- Samuelson, P. (1965). Proof that properly anticipated prices fluctuate randomly. *Industrial Management Review*, 6, 41–49.
- Skouras, S. (2001). Financial returns and efficiency as seen by an artificial technical analyst. *Journal of Economic Dynamics and Control*, 25, 213–244.
- Wang, Y. (2001). Predicting stock price using fuzzy grey prediction system. *Expert Systems with Applications*, 22 (1), 33–38.