



ELSEVIER

Economics Letters 77 (2002) 55–63

**economics
letters**

www.elsevier.com/locate/econbase

Market timing: a test of a charting heuristic

William Leigh^{a,*}, Noemi Paz^a, Russell Purvis^b

^a*Department of Management Information Systems, College of Business, University of Central Florida, Box 161400, Orlando, FL 32816-1400, USA*

^b*Clemson University, Clemson, SC, USA*

Received 16 August 2001; received in revised form 4 March 2002; accepted 12 March 2002

Abstract

We implement a graphical (or ‘charting’) heuristic, the ‘bull flag’, which accepts a particular pattern of historical prices as a signal for a future market price increase, test it with several years of New York Stock Exchange Composite Index history, and find positive results. The results support the validity of technical analysis for stock market price prediction and fail to confirm the efficient markets hypothesis.

© 2002 Elsevier Science B.V. All rights reserved.

Keywords: Technical trading; Security markets

JEL classification: G14

1. Introduction

One interpretation of the efficient markets hypothesis is that market prices follow a random walk and cannot be predicted based on their past behavior. Discoveries of ‘anomalies’, relationships that can be used to earn abnormal returns and appear to violate the efficient markets hypothesis, are numerous in the finance literature. Well-known anomalies involve: unexpected earnings announcements, firm size, month of January, day of the week, analysts’ recommendations, impact of the federal budget deficit announcement, and others. Frankfurter and McGoun (2001) survey the anomalies literature and discuss the paradigmatic crisis in academic finance which they represent.

Stock market forecasters who practice technical analysis concern themselves with the dynamics of the market price and volume behavior itself, rather than with the fundamental economic nature of specific securities that are traded. Charles Dow published the original Dow Theory for technical analysis in 1884, and a modern explication is found in Edwards and Magee (1997). The efficient

*Corresponding author. Tel.: +1-407-823-3173; fax: +1-407-823-2389.

E-mail address: leigh@pegasus.cc.ucf.edu (W. Leigh).

markets hypothesis implies that the technical approach to market price prediction is invalid. However, positive reports as to the effectiveness of metric-based trading rules, such as momentum measures and moving averages, which use only historical price and volume information and are considered by many to be ‘technical’ rules, have appeared more than once and recently in the most respected finance journals. Hong et al. (2000) and Hong and Stein (1999) examine trading rules that use measures of momentum. Gencay (1998) reports positive results with metric-based technical trading rules implemented with nonparametric models. Neftci (1991), Brock et al. (1992), and an increasing number of others, look at technical trading rules and report positive results.

In addition to the use of metric-based trading rules, technical analysis includes charting heuristics. Charting heuristics are used to identify certain graphical patterns in historical price and volume time series data that are considered to be signals to buy (or sell). Lo et al. (2000) test charting heuristics using kernel regression for pattern identification and find marginally positive results. Neftci (1991) discusses the difficulties that the conventional methods of finance, based on linear models, have in describing typical stock market activity of interest to stock traders: the recognition of sporadic buy and sell signals, and the recognition of patterns in time series. In this paper we illustrate the use of template matching, a basic technique from pattern recognition, to implement a charting trading heuristic from technical analysis. This approach may provide the nonlinear and rule-based method that Neftci (1991) is seeking.

We present results from testing one variation of one technical analysis 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 then a horizontal or downward sloping flag of ‘consolidation’ followed by a sharp rise in the positive direction, the ‘breakout.’ In this paper we concentrate on this particular pattern, the bull flag, because the results are crisp and persuasive. (In addition to the particular variation of the bull flag pattern presented in this paper, we looked at other patterns applied to both price and volume and achieved mixed results.) We have found no rigorous testing of this particular charting pattern anywhere in the academic literature.

We implement the bull flag charting heuristic through use of a template matching technique from pattern recognition and test the resulting market price forecasts against the overall average price increase experienced in the period we are using for several prediction horizons (10, 20, 40, and 80 trading days). We work with the New York Stock Exchange Composite Index, and for this work with a broad-based composite index, the overall average price (index value) increase/decrease in the period is equivalent to the return from a buy-and-hold or random-selection trading strategy, which are implied as optimal by the random walk model of the efficient markets hypothesis.

We use the values of the New York Stock Exchange Composite Index for the period from 8/6/80 to 9/15/99 for testing and compute a value for how well a template representation of the bull flag pattern fits or matches the 40 trading day window ending with each of the 4817 trading days in the test period. The computed fit values are used in trading rules of the sort: ‘If the fit value for a trading day exceeds a set value then buy on that trading day and hold for some number of days.’ If the average of the returns in a test period from simulated trades using this rule exceeds the average of returns which would have accrued to buying on every day in the period of comparison by a statistically significant amount, then we have found a successful forecasting method based only on price history; this finding fails to confirm the efficient markets hypothesis (specifically, the weak form) and contributes in an important way to the growing ‘anomalies’ literature.

T =

-1.5	-1.5	-1.5	-1.5	-1.5	-1.5	-1.5	.25	1.25	4
-1	-1	-1	-1	-1	-1	-1	.25	1	2
-5	-5	-5	-5	-5	-5	-5	.25	.75	1
-5	-5	-5	-5	-5	-5	-5	.25	0	0
-5	-5	-5	-5	-5	-5	-5	0	0	0
-25	-25	-25	-25	-25	-25	-25	0	0	-1
-25	-25	-25	-25	-25	-25	-25	0	-.75	-1
.5	.5	.5	.5	.5	.5	.5	0	-.75	-1
2	2	2	2	2	2	2	-.5	-.75	-2
2	2	2	2	2	2	2	-.5	-.75	-2

Fig. 1. 10×10 grid of weights used in this study to represent the bull flag charting pattern. This template is fitted or matched to 4817 40 trading day wide windows, each fitting window ending on one of the 4817 trading days in the period of the study. The hypothesis is that good fits are indications of buying opportunities.

2. Method

Fig. 1 shows the template, T, that we use for the bull flag charting pattern. This is a 10 × 10 grid with weights, w_{ij} , ranging from −2 to +4 in the cells. The weighting values define areas in the template for the horizontal consolidation (first seven columns) and for the upward-tilting breakout (last three columns) portions of this bull flag pattern, which are also indicated by the graying in the figure.

The bull flag pattern template, T, is fitted/matched to the NYSE Composite Index’s time series closing price data 4817 times by fitting a window of 40 price values to the template starting with the oldest price and moving the window up one trading day for each of the next 4816 fittings. The procedure used to accomplish the fitting is *template matching* (Duda and Hart, 1973), a pattern recognition technique used to match a template to a pictographic image to identify objects. We let p_t be the composite’s price value on trading day t for the fitting window ending on trading day k , where $t = -39, \dots, 0, k = 1, \dots, 4817$, and $k = 1$ is the oldest price. For each trading day k we synthesize a 10×10 image grid, I_k , from each set of 40 closing price values. Next, we compute a cross-correlation of the bull flag template T with the image grid I_k and calculate two output values for each fitting: FIT_k and $HEIGHT_k$.

The following is a specification for the template matching process for a single 40 trading day window. Within each 40-day window of data, we ‘Winsorize variances’ (Roberts, 1995, p. 150) to remove the worst noise by replacing every observation which is beyond two standard deviations from the mean of the price values in the window with the respective two standard deviation boundary value.

The next step is to take the 40 days of closing prices and map the information into a 10×10 image grid for the fitting window ending with trading day k . Let the image grid’s gray scale values, g_{ij} , be the individual values computed into each cell of the 10×10 image grid, I_k . First we define how the price values will relate to the rows in the grid by calculating the range of the 40 prices and dividing the range by 10 to arrive at an increment value:

$$inc = (p_{max} - p_{min})/10$$

p_{max} and p_{min} are the maximum and minimum price values found within the 40 values in each window. Using this increment, we associate a row i with an interval:

$$[p_{\max} - i \cdot inc, p_{\max} - (i - 1) inc] \quad \text{for } i = 1$$

and

$$[p_{\max} - i \cdot inc, p_{\max} - (i - 1) inc] \quad \text{for } i = 2, \dots, 10$$

Each image grid's column j corresponds to four price values, p_t , at a time from the 40 in the fitting window. Specifically, prices $p_{-4(10-j)-3}$, $p_{-4(10-j)-2}$, $p_{-4(10-j)-1}$, and $p_{-4(10-j)}$ are associated with column j , where $j = 1, \dots, 10$. The image grid's gray scale values, g_{ij} , are found for a column j by determining what portion of each column's four price values fall into each of the 10 intervals identified by rows $i = 1, \dots, 10$:

$$g_{ij} = \begin{cases} 0 & \text{if none of the 4 } p_t \text{'s for column } j \text{ fall in interval } i \\ 0.25 & \text{if 1 of the 4 } p_t \text{'s for column } j \text{ fall in interval } i \\ 0.5 & \text{if 2 of the 4 } p_t \text{'s for column } j \text{ fall in interval } i \\ 0.75 & \text{if 3 of the 4 } p_t \text{'s for column } j \text{ fall in interval } i \\ 1.0 & \text{if 4 of the 4 } p_t \text{'s for column } j \text{ fall in interval } i \end{cases}$$

Finally, we calculate the FIT_k and $HEIGHT_k$ for the fitting window that ends with trading day k . FIT_k is a cross-correlation of the template grid's weights with the image grid's gray scale values, and $HEIGHT_k$ is the fitting window's price range value normalized by the fitting window's ending trade day price, p_k , for $k = 1, \dots, 4817$. Note that $p_k = p_t$ when $t = 0$. The calculated values are:

$$FIT_k = \sum_{i=1}^{10} \sum_{j=1}^{10} (w_{ij} g_{ij})$$

$$HEIGHT_k = (p_{\max_k} - p_{\min_k}) / p_k$$

3. Results

We test the results of applying a trading rule by comparing to the results of buying on every day in the period of comparison and holding for the number of trading days in the horizon specified in the trading rule.

Let p_k = a NYSE Composite Index price value on trading day k ; $HEIGHT_k$ = a HEIGHT value computed as described above for trading day k ; FIT_k = a FIT value computed as described above for trading day k ; h = number of trading days in the forecast horizon, where $h = 10, 20, 40, 80$ where $k = 1, \dots, 4817$ for the trading days in the period of the study; m = the first trading day k in a subperiod of comparison; n = the last trading day k in a subperiod of comparison.

Calculate results for a subperiod, s , when buying every day:

$$\text{Market Average Return}_s = \sum_{k=m}^n [(p_{k+h} - p_k) / p_k] / (n - m + 1)$$

and results when buying as specified by the trading rule:

$$\text{Number of Buys}_s = \sum_{k=m}^n R_k$$

where

$$R_k = \begin{cases} 1 & \text{if trading rule is true for FIT}_k \text{ and/or HEIGHT}_k \\ 0 & \text{otherwise} \end{cases}$$

$$\text{Trading Rule Average Return}_s = \sum_{k=m}^n [((p_{k+h} - p_k)R_k)/p_k] / \text{Number of Buys}_s$$

Finally, we have ‘excess’ profits for the subperiod:

Table 1

‘Excess’ profits as obtained from applying the results of the fitting process in a trading rule with height, fit, and horizon parameters to the 4817 trading days in our test period

Height ≥	Fit ≥	Horizon (%)			
		10	20	40	80
0.10	0	0.9	1.7	2.2	2.3
	2	0.8	1.6	2.0	2.4
	4	0.6	1.4	1.8	2.3
	6	0.6	1.3	2.1	2.4
	8		1.2	2.1	2.2
0.12	0	1.1	2.5	3.6	3.9
	2	1.0	2.4	3.5	4.1
	4	0.8	2.0	3.1	4.0
	6	0.7	1.9	3.3	4.0
	8		2.4	4.1	5.1
0.14	0	1.2	2.5	4.1	5.3
	2	1.1	2.4	4.0	5.2
	4		1.9	3.8	5.1
	6		1.8	4.0	5.4
	8		2.6	4.9	6.6
0.16	0		4.0	5.8	7.3
	2		4.0	5.8	7.3
	4		3.6	5.7	7.3
	6		3.9	6.4	8.3

The excess profit value in the cell is the difference between the Market Average Return, which is the average profit realized by buying on every day, and the Trading Rule Average Return, which is the average profit realized by buying only on the rule-indicated days. Both market strategies buy and hold for the number of trading days in the horizon period. ‘Height’ refers to the difference between the maximum and minimum prices in the 40 trading day wide template fitting window normalized by the price on the trading day of the fitting. ‘Horizon’ is the number of trading days between buying and selling. The trading rule is, ‘For a trading day, if HEIGHT_k ≥ Height and FIT_k ≥ Fit then buy and hold for Horizon trading days.’ A cell value appears in this table only if: (1) there are more than 30 trading days that meet the HEIGHT and FIT requirements of the trading rule for that cell; and (2) the *P*-value for the *t*-test (between the Market Average Return and the Trading Rule Average Return for the horizon period) is less than 1%.

$$\text{Excess Profits}_s = \text{Trading Rule Average Return}_s - \text{Market Average Return}_s$$

We compare Market Average Return to Trading Rule Average Return using a two-sample, one-tailed, unequal variance (heteroscedastic) Student's t -test. Table 1 displays Excess Profits obtained from applying the results of the fitting process as a trading rule of the form, 'For a trading day, if $\text{HEIGHT}_k \geq 0.10 \geq H$ and $\text{FIT}_k \geq F$ then buy and hold for h trading days', to the 4817 trading days in our test period. The excess profit value in the table's cells is the difference between the average profit realized by buying on the rule-indicated days and holding for the number of trading days in the horizon period and by buying every day and also holding for the number of trading days in the horizon period. A cell value appears in Table 1 only if: (1) there are more than 30 trading days that meet the HEIGHT and FIT requirements of the trading rule for that cell; and (2) the P -value for the t -test (between the average profit realized when using the trading rule and when buying every day) was less than 1%.

Table 2 considers the trading rule corresponding to the first row of results in Table 1: 'For a trading day, if $\text{HEIGHT}_k \geq 0.10$ and $\text{FIT}_k \geq 0.0$ then buy and hold for h trading days'. The table cells contain excess profits and t -test P -value results for subperiods of lengths of 400 trading days with the last subperiod length equal to 417 trading days. The 'Over All' row at the bottom of Table 2 compares the return from application of this trading rule for a subperiod equal to the complete 4817 trading days in

Table 2

Excess profits and t -test P -values for 400 day intervals (last interval is 417 days) for application of the trading rule 'For a trading day, if $\text{HEIGHT}_k \geq 0.10$ and $\text{FIT}_k \geq 0.0$ then buy and hold for h trading days'

Period begin	Buy days	Horizon							
		10		20		40		80	
		Excess profit (%)	P -value	Excess profit (%)	P -value	Excess profit (%)	P -value	Excess profit (%)	P -value
08/06/80	16	-1.6	0.1230	0.0	0.4865	-3.2	0.0003	2.4	0.0136
03/09/82	65	1.4	0.0089	2.3	0.0003	4.3	0.0000	4.7	0.0000
10/06/83	45	-0.3	0.824	-1.2	0.0000	-1.2	0.0001	-1.0	0.0146
05/13/85	33	0.3	0.0228	-0.1	0.4428	0.2	0.3614	-0.7	0.2913
12/12/86	57	1.0	0.0034	2.8	0.0000	4.2	0.0000	3.7	0.0000
07/14/88	0								
02/12/90	32	0.7	0.0277	1.5	0.0002	1.6	0.0009	0.1	0.4641
09/12/91	16	-0.8	0.0046	-1.2	0.0000	-3.3	0.0000	-3.2	0.0000
04/14/93	0								
11/11/94	0								
06/17/96	39	1.3	0.0002	2.4	0.0000	2.3	0.0012	-0.1	0.4448
01/16/98	61	1.7	0.0000	3.0	0.0000	3.2	0.0000	4.3	0.0000
Overall:	364	0.85	0.0000	1.67	0.0000	2.17	0.0000	2.33	0.0000
Trading Rule Avg. Return:			1.34		2.64		4.12		6.24
Market Average Return:			0.49		0.97		1.94		3.91

The column 'Buy Days' contains the number of days for which the trading rule is true in the interval. The 'Over All' row compares the Trading Rule Average Return (realized by buying on 364 indicated days out of the 4817) with the Market Average Return realized by buying on all (4817) days in the period of the study (8/6/80 to 9/15/99).

Table 3

Results of a regression analysis using the HEIGHT to forecast the price at a horizon of 80 trading days^a

Multiple <i>R</i>	0.38095					
<i>R</i> square	0.14512					
Adjusted <i>R</i> square	0.14276					
Standard error	0.05150					
Observations	364					
	df	SS	MS	<i>F</i>	Significance <i>F</i>	
Regression	1	0.16299	0.16299	61.45238	0.00000	
Residual	362	0.96015	0.00265			
Total	363	1.12315				
	Coefficients	Standard error	<i>t</i> -Stat	<i>P</i> -value	Lower 95%	Upper 95%
Intercept	−0.05205	0.01485	−3.50562	0.00051	−0.08124	−0.02285
Height	0.88834	0.11332	7.83916	0.00000	0.66549	1.11120

^a This uses the HEIGHT_{*k*} and FIT_{*k*} values computed for the 364 trading days with a HEIGHT_{*k*} ≥ 0.10 and a FIT_{*k*} ≥ 0.0. This is the regression analysis display from Microsoft Excel and includes an analysis of variance: ‘df’ is ‘degrees of freedom’, ‘SS’ is ‘sums of squares’, ‘MS’ is ‘mean squares’, ‘*F*’ is ‘*F*-ratio’, ‘Lower’ and ‘Upper 95%’ refer to the confidence interval values, and so forth.

the test period. The difference is significant for the 10-, 20-, 40-, and 80-day horizons tested in this study.

Table 2 results are for the trading rule corresponding to the top row of Table 1. Further inspection of Table 1 reveals that analyses of the form of Table 2 for the trading rules represented by the other rows of Table 1 would show results superior to those in Table 2. Examination of Tables 1 and 2 reveals that these results would admit various presentations which include significant results with a holdout sample.

Tables 3 and 4 contain results from linear regression analyses. In Table 3 the HEIGHT is used to forecast the price at a horizon of 80 trading days. In Table 3, the regression is computed for the 364 trading days which had HEIGHT_{*k*} ≥ 0.1 and FIT_{*k*} ≥ 0.0. Table 4 shows the results of a like analysis but for the 69 trading days which had HEIGHT_{*k*} ≥ 0.10 and FIT_{*k*} ≥ 10.0. The statistics from these regression analyses are significant, and are better with the higher minimum FIT value (Table 4). The FIT and HEIGHT values appear to have predictive value for price level, as well as for future price direction which is the only concern of the trading rules.

4. Conclusion

We test the bull flag charting heuristic for trading the NYSE Composite Index in a rigorous way, which has not been done before. Statistical results are significant and fail to confirm the null hypothesis that the markets are (efficient markets hypothesis weak form) efficient. We supply results for a long time period and in such a way that ‘data snooping’ charges may be deflected; the results are valid *ex ante*.

Stock market buy/sell transaction costs are not considered. However, gross profitability results with

Table 4

Results of a regression analysis using the HEIGHT to forecast the price at a horizon of 80 trading days^a

Multiple <i>R</i>	0.53830					
<i>R</i> square	0.28977					
Adjusted <i>R</i> square	0.27917					
Standard error	0.03103					
Observations	69					
	df	SS	MS	<i>F</i>	Significance <i>F</i>	
Regression	1	0.02632	0.02632	27.33573	0.00000	
Residual	67	0.06451	0.00096			
Total	68	0.09083				
	Coefficients	Standard error	<i>t</i> -Stat	<i>P</i> -value	Lower 95%	Upper 95%
Intercept	−0.08350	0.02329	−3.58602	0.00063	−0.12998	−0.03703
Height	0.97130	0.18578	5.22836	0.00000	0.60049	1.34211

^a This uses the HEIGHT_{*k*} and FIT_{*k*} values computed for the 69 trading days with a HEIGHT_{*k*} ≥ 0.1 and a FIT_{*k*} ≥ 10. This is the same analysis as carried out for Table 3.

the reported method for buy-and-hold periods from 10 to 80 days are found to be twice what is experienced with random buying. In this era of fixed-price fees, or free 401(k) transfers, excess profits after transaction costs with the reported method may be obtained with a suitable trade size. We do not provide an overall study period (over-19-year) cash flow comparison with buy-and-hold as only one charting heuristic has been investigated, and an overall study period comparison between buy-and-hold and market timing with charting heuristics would require that we develop, tune, and deploy many charting heuristics.

References

- Brock, W., Lakonishok, J., LeBaron, B., 1992. Simple technical trading rules and the stochastic properties of stock returns. *The Journal of Finance* 47, 1731–1764.
- Downes, J., Goodman, J., 1998. *Dictionary of Finance and Investment Terms*, 5th Edition. Barron's Educational Series, New York.
- Duda, R., Hart, P., 1973. *Pattern Classification and Scene Analysis*. Wiley, New York.
- Edwards, R., Magee, J., 1997. *Technical Analysis of Stock Trends*, 7th Edition. Amacom, New York.
- Frankfurter, G., McGoun, E., 2001. Anomalies in finance: what are they and what are they good for? *International Review of Financial Analysis* 10 (4), 407–429.
- Gencay, R., 1998. Optimization of technical trading strategies and the profitability in security markets. *Economics Letters* 59, 249–254.
- Hong, H., Lim, T., Stein, J., 2000. Bad news travels slowly: size, analyst coverage, and the profitability of momentum strategies. *The Journal of Finance* 55 (1), 265–295.
- Hong, H., Stein, J., 1999. A unified theory of underreaction, momentum trading, and overreaction in asset markets. *The Journal of Finance* 54 (6), 2143–2184.

- Lo, A., Mamaysky, H., Wang, J., 2000. Foundations of technical analysis: computational algorithms, statistical inference, and empirical implementation. *The Journal of Finance* 6 (4), 1705–1796.
- Neftci, S., 1991. Naïve trading rules in financial markets and Wiener–Kolmogorov prediction theory: a study of ‘technical analysis’. *The Journal of Business* 64 (4), 549–571.
- Roberts, T., 1995. *Bounded Rationality in Macroeconomics*. Oxford University Press, Oxford.