The Predictive Power of "Head-and-Shoulders" Price Patterns in the U.S. Stock Market[*]

Gene Savin [†]
The University of Iowa

Paul Weller [‡]
The University of Iowa

Jānis Zvingelis [§]
Mesirow Financial

Revised version
October 19, 2006

[*] We are grateful to Bruce Lehmann, to three referees and to seminar participants at the University of Iowa for very helpful comments, and to Ken French for making available the data on factor risk portfolios.

[†] Department of Economics, Henry B. Tippie College of Business, The University of Iowa, Iowa City, IA 52242, (319) 335-0855, (319) 335-1956 (fax), gene-savin@uiowa.edu

[‡] (Corresponding author), Department of Finance, Henry B. Tippie College of Business, The University of Iowa, Iowa City, IA 52242, (319) 335-1017, (319) 335-3690 (fax), paul-weller@uiowa.edu

[§] Investment Strategies, Mesirow Financial Investment Management, 350 N. Clark, Chicago, IL 60610, (312) 595-7299, (312) 595-6139 (fax), jzvingelis@mesirowfinancial.com

Abstract:

We use the pattern recognition algorithm of Lo et al. (2000) with some modifications to determine whether "head-and-shoulders" price patterns have predictive power for future stock returns. The modifications include the use of filters based on typical price patterns identified by a technical analyst. With data from the S&P 500 and the Russell 2000 over the period 1990-1999 we find strong evidence that the pattern had power to predict excess returns. Risk-adjusted excess returns to a trading strategy conditioned on "head-and-shoulders" price patterns are 5-7 percent per year. The success of the strategy rests in part on its identification of negative momentum stocks.

Keywords: technical analysis, kernel regression, stock prices
JEL subject numbers: G1, G11, G14

Introduction

Technical analysts use information about historical movements in price and trading volume, summarized in the form of charts, to forecast future price trends in a wide variety of financial markets. They argue that their approach to trading allows them to profit from changes in the psychology of the market. This view is summarized in the following quotation:

The technical approach to investment is essentially a reflection of the idea that prices move in trends which are determined by the changing attitudes of investors toward a variety of economic, monetary, political and psychological forces… Since the technical approach is based on the theory that the price is a reflection of mass psychology ("the crowd") in action, it attempts to forecast future price movements on the assumption that crowd psychology moves between panic, fear, and pessimism on one hand and confidence, excessive optimism, and greed on the other. (Pring, 1991, pp. 2–3).

The claims that technical trading rules can generate substantial profits are rarely, if ever, subjected to scientific scrutiny by the technicians themselves. The many books and trading manuals that have been written on the subject of technical analysis typically use a method of description and anecdote. By contrast, early academic work testing the efficient market hypothesis (Fama, 1965; Fama, 1970) concluded that there was no evidence that stock market prices were predictable, and that therefore there was no substance to technical analysis.

However, more recent work has started to question the original findings. Brock, Lakonishok and LeBaron (1992) find strong support for the ability of several widely used technical rules to predict the Dow Jones Industrial Average index. These results are confirmed by Sullivan, Timmermann and White (1999) who make corrections for data snooping bias. In addition there is now convincing evidence that stock prices display short term momentum over periods of six months to a year and longer term mean reversion (De Bondt and Thaler, 1985; Chopra, Lakonishok and Ritter, 1992; Jegadeesh and Titman, 1993). There is also evidence of economically significant price reversals over short time horizons of a week to a month (Jegadeesh, 1990; Jegadeesh and Titman, 1995). This can be interpreted as providing support for a particular class of technical trading rule that is designed to detect trends. Such rules have been shown to perform profitably in foreign exchange markets (Dooley and Shafer, 1983; Sweeney, 1986; Levich and Thomas, 1993; Neely, Weller and Dittmar, 1997).

There have been theoretical arguments advanced to explain these observed patterns of momentum and reversal (Barberis, Shleifer and Vishny, 1997; Daniel, Hirshleifer and Subrahmanyam, 1998). These arguments introduce various departures from fully rational behavior, and carry the implication that investors using trading rules of the trend-following variety may be able to profit from these departures from rationality. Other work has demonstrated that even if all agents are rational there may be a role for technical analysis to play (Treynor and Ferguson, 1985; Brown and Jennings, 1989; Blume, Easley and O'Hara, 1994). But these papers are less specific about the type of technical indicator that will be profitable.

Much less academic attention has been paid to the use of technical signals based on *price patterns*, despite the fact that these are widely used by practitioners. Chang and Osler (1999) examine the profitability of using the "head-and-shoulders" pattern in the foreign exchange market to predict changes of trend, and find evidence of excess returns for some currencies but not others. Lo et al. (2000) develop a pattern detection algorithm based on kernel regression. They apply this methodology to identifying a variety of technical price patterns including "head-and-shoulders" in the U.S. stock market over the period 1962 - 1996. They find statistical evidence that there is potentially useful information contained in most of the patterns they consider.

The aim of this paper is to examine the methodology advocated in Lo et al. (2000), to propose some modifications and to use the modified approach to assess the predictive power of a particular pattern. One of the difficulties that an academic investigator must face in assessing the predictive power of price patterns is that the characterization of the patterns is sometimes ambiguous and there may be disagreement among technical analysts themselves. For this reason, we have chosen to focus on the head-and-shoulders (HS) pattern. There is a very general consensus on the important features of this pattern, and it is also agreed that this is one of the most reliable technical indicators.[1]

The occurrence of a technical price pattern is taken as a signal of either a *continuation or reversal in a price trend*. Therefore we concentrate on determining whether there is any evidence that a pattern can predict the sign or magnitude of stock returns. Lo et al. (2000), in contrast, do not focus on price predictability but rather consider the question of whether there is any informational content in the occurrence of a pattern for the whole conditional distribution of returns. While this is certainly of academic interest, it focuses less directly on a test of the claims of technical analysis. We assess the predictive power of the pattern over considerably longer time windows than do Lo et al. – one month, two months and three months rather than one day. Again, our justification for doing this is that it accords better with the practice of technical analysts.

As mentioned above, books written on technical analysis almost never contain any attempt at statistical analysis. A recent exception to this general rule is the book *Encyclopedia of Chart Patterns* (Bulkowski, 2000). Bulkowski uses a computer algorithm to search for a large number of different types of pattern in a population of 500 stocks over the period 1991 to 1996. As a practicing technical analyst he does not rely exclusively on the results of his computer search, which naturally raises questions about data snooping. However, he reports statistics on the patterns he identifies, including number of occurrences, average length of time from initiation to completion of the pattern, failure rate and frequency distribution of returns after the occurrence of a pattern. He also provides a number of examples of each pattern he considers. Using these examples together with others taken from Bulkowski (1997) allows us to calibrate the pattern detection algorithm implemented by Lo et al. and to supplement it with filters on the assumption that the examples presented by Bulkowski are "typical".

We examine the performance of the kernel smoothing algorithm alone, and supplemented by the filters based on the examples of Bulkowski. We do this separately for S&P 500 stocks and for Russell 2000 stocks over the period 1990-1999. For the first group, which is comprised of stocks with relatively large capitalization, we find no evidence that a trading strategy based on HS patterns alone would be profitable, and for the second group, which consists of smaller stocks, the evidence is weak. Thus our results provide no support for the more extreme claims of some technical analysts, including Bulkowski, that trades based on the occurrence of HS patterns alone are consistently profitable. However we do find significant risk-adjusted excess returns after the occurrence of an HS pattern for both groups. For the second group, we find that risk-adjusted excess returns are five to seven per cent per annum over a three-month window. We augment the three-factor model with the momentum factor and find that some, but not all of the excess return can be attributed to negative momentum. We interpret this as a demonstration of mispricing, and show that it persists even after adjustment for transaction costs.

[1] See Edwards and Magee (1992, pp.63-64) and Bulkowski (2000, p. 290).

I. Methodology and Procedures

In this section, we describe the methodology used to identify HS patterns and the procedure used to calculate the return conditional on detecting an HS pattern. Our methodology is a modification of that employed by Lo et al. (2000). As in Chang and Osler (1999), they use a computer-based algorithm for selecting HS patterns where the patterns are defined by the extrema of the price series. The distinctive contribution of Lo et al. (2000) is that they initially smooth the price series using kernel mean regression. The advantage of this approach is that it provides a plausible analogue to the signal extraction task performed by the human eye when it filters out noise and identifies the occurrence of a price pattern.

A. Methodology of Lo et al.

A.1. Data Generation Process

The data consists of observations on prices of the stock, P_i, at integer values of time, X_i, where X_i is the i-th tick of time, that is, $X_i = i$, $1 < i < T$. This notation distinguishes the counter i from the integer value that time takes on. In effect, they assume that the data are generated by a *fixed design model* with a controlled nonstochastic X variable, which in this setup is time (Härdle (1990)). Hence,

(1) $$P_i = m(X_i) + \varepsilon_i\,,\ 1 < i < T,$$

where $m(X_i)$ is a smooth function of time and the ε_i's are independently and identically distributed zero mean random variables with variance σ^2. The $m(X_i)$ series can be interpreted as the filtered or smoothed price series. Note that this model is observationally indistinguishable from the case where the P_i's are autocorrelated.

A.2. Rolling Windows

In practice, HS patterns identified by technical traders typically occur within a three-month period. The maximum allowable period is called a window and the span of the window is denoted by n. Lo et al. (2000) analyze the data using rolling windows of span n = 38 trading days. That is, the price series is divided into successive, overlapping windows of 38 trading days where the difference between the left limits of two adjacent windows is one trading day. In other words, for any given window, the succeeding window starts and ends one business day later. The motivation for using rolling windows is twofold. One is that it approximately mimics the way in which traders analyze the data. If windows did not overlap then the pattern recognition algorithm would not detect any pattern initiated in one window and completed in the next. However traders in principle use all the historical price data as time unfolds. The other is that it automatically constrains the maximum length of the HS pattern.

A.3. Kernel Mean Regression

The price series within each window of span n is smoothed using a kernel nonparametric regression. The kernel nonparametric estimator of the part of *m*(x) that lies within the i-th window, i =1,…, T-n+1, is

(2) $$m_{i,n}(\mathrm{x}) = \frac{\sum_{j=i}^{i+n-1} \mathrm{P_j} K\left(\frac{\mathrm{x}\text{-}\mathrm{X_j}}{h_{i,n}}\right)}{\sum_{j=i}^{i+n-1} K\left(\frac{\mathrm{x}\text{-}\mathrm{X_j}}{h_{i,n}}\right)}$$

that is often called the *Nadaraya-Watson estimator*. K(•) is the kernel, a function which satisfies certain conditions. In our estimation, K(•) is the standard normal density function. The bandwidth, $h_{i,n}$, can be interpreted as a smoothing parameter. The higher the value $h_{i,n}$, the smoother the $m_{i,n}$(x) function. In practice, the bandwidth parameter has to be chosen, which implies that the bandwidth is generally different for different rolling windows. The method used to select the bandwidth is the so-called "leave-one-out" method, which is also called cross-validation. The details on kernel mean regression and bandwidth selection are found in Härdle (1990).

A.4. Extrema

Given the smoothed price series $m_{i,n}(\mathrm{x})$ within a window, the extrema are identified by a two-step procedure. The first step is to find the extrema of the smoothed price series $m_{i,n}(\mathrm{x})$, and the second is to find the corresponding values of the original P_i series at the extrema in the first step. For the purpose of exposition, we call the latter the *relevant* extrema.

The extrema for the smoothed price series $m_{i,n}(\mathrm{x})$ are defined as follows. The point $m_{i,n}(X_i)$ is a local maximum if $m_{i,n}(X_{i-1}) < m_{i,n}(X_i)$ and $m_{i,n}(X_i) \geq m_{i,n}(X_{i+1})$. The inequalities are reversed if $m_{i,n}(X_i)$ is a local minimum. If $m_{i,n}(X_i)$ is identified as a local extremum, then the relevant extremum is defined on the interval from P_{i-1} to P_{i+1}, which implies that it may differ from P_i.

Let $E_1, E_2, \ldots, E_m$ denote the set of relevant extrema and $X_1^*, X_2^*, \ldots, X_m^*$ the dates at which these extrema occur. In Lo et al. (2000), an HS pattern consists of a set of five consecutive relevant extrema which satisfy the following restrictions.

(R1) E_1 is a maximum.

(R2) $E_3 > E_1$.

(R3) $E_3 > E_5$.

(R4) $\max_i \left| E_i - \bar{E} \right| \leq 0.015 \cdot \bar{E}$, i = 1, 5, where $\bar{E} = (E_1 + E_5)/2$.

(R5) $\max_i \left| E_i - \bar{E} \right| \leq 0.015 \cdot \bar{E}$, i = 2, 4, where $\bar{E} = (E_2 + E_4)/2$.

In the restrictions, E_1 is the left shoulder, E_3 is the head and E_5 is the right shoulder. (R4) and (R5) restrict the distance between the height of the left and right shoulder and the left and right trough; namely, E_1 and E_5 are within 1.5 percent of their average and E_2 and E_4 are within 1.5 percent of their average. Figure 1 illustrates the shape of a typical HS pattern with the extrema labelled.

B. Modifications

Our methodology for identifying HS patterns differs from Lo et al. (2000) in four respects. The first involves the span of the rolling windows. We set the span at n = 63. This is based on the number reported in Bulkowski for the average completion time of an HS pattern.

The second concerns the bandwidth. The HS patterns are selected using four different values of the bandwidth. The values of the bandwidth are multiples of $h_{i,\ n}$, the bandwidth obtained by the cross-validation method. The multiples are 1, 1.5, 2 and 2.5. The number and type of HS patterns selected are sensitive to the magnitude of the bandwidth. The number of HS patterns selected decreases substantially as the bandwidth increases.

The third involves the restrictions on the relevant extrema. Technical trading manuals suggest characteristics that HS patterns have to satisfy. The manuals generally agree on the form of the restrictions (R1) to (R5). On the basis of Bulkowski (2000), we recalibrate the Lo restrictions (R4) and (R5) as follows.

(R4a) $\max_i \left| E_i - \overline{E} \right| \leq 0.04 \cdot \overline{E}$, i = 1, 5, where $\overline{E} = (E_1 + E_5)/2$.

(R5a) $\max_i \left| E_i - \overline{E} \right| \leq 0.04 \cdot \overline{E}$, i = 2, 4, where $\overline{E} = (E_2 + E_4)/2$.

This allows a greater difference between the height of the two shoulders. It also allows the neckline to be more steeply sloped where the neckline is the line joining E_2 and E_4.

Restrictions (R1) to (R5) do not capture all the features of an HS pattern of interest to technical traders. We impose four additional restrictions:

(R6) $$\frac{[(E_1 - E_2) + (E_5 - E_4)]/2}{E_3 - (E_2 + E_4)/2} \leq 0.7\,.$$

(R7) $$\frac{[(E_1 - E_2) + (E_5 - E_4)]/2}{E_3 - (E_2 + E_4)/2} \geq 0.25\,.$$

(R8) $$\frac{[E_3 - (E_2 + E_4)/2]}{E_3} \geq 0.03\,.$$

(R9) $$\max_i \left| (X^*_{i+1} - X^*_i) - \overline{X}^* \right| \leq 1.2 \cdot \overline{X}^*, \text{ i = 1,...,4, where } \overline{X}^* = \sum_{i=1}^{4} (X^*_{i+1} - X^*_i)/4.$$

The restrictions (R6) to (R9) are calibrated using eleven examples of HS patterns reported in Bulkowski ((1997), (2000)) that are completed within 63 trading days. We refer to (R4a), (R5a), (R6), (R7), (R8) and (R9) as the Bulkowski restrictions.

Restrictions (R6) and (R7) specify the average height of the shoulders as a proportion of the height of the head from the neckline. In particular, (R6) and (R7) combined with (R4a) and (R5a) typically rule out cases where the height of the shoulders is a very large or very small proportion of the height of the head. (R8) rules out cases where the height of the head from the neckline is a small proportion of the stock price, and (R9) rules out extreme horizontal asymmetries in the HS patterns.

The fourth involves the imposition of the neckline crossing condition. If we refer to Figure 1, the neckline passes through the local minima at E_2 and E_4. In order for an HS pattern to have been completed, we require that price fall below the neckline on its way to the next extremum, E_6, which, by definition, has to be a minimum. Technical manuals generally emphasize the importance of this condition.

In addition, we impose the requirement that E_6 occur at date n-3. In other words, if the neckline crossing condition is satisfied, investment returns are measured from a date three days after E_6. Strictly speaking, the short position should be opened immediately after the neckline crossing condition is satisfied, but our procedure is considerably simpler to implement and in practice the time span between the two dates is generally quite short. In addition, any bias introduced will unambiguously reduce measured excess returns.

C. Procedure for Calculating Conditional Excess Returns

For each HS pattern detected, Lo et al. (2000) calculate the continuously compounded return over one subsequent trading day. In technical trading manuals, substantially longer horizons are considered, often supplemented with exit conditions dependent on the path of prices. However, there is no clear consensus on the appropriate horizon or exit conditions. Accordingly, we calculate the continuously compounded return over the subsequent 20, 40 and 60 trading days. Bulkowski (2000) reports from his investigation of 431 HS patterns in 500 stocks over the period 1991 to 1996 that the time taken to reach the ultimate low was on average three months.

Suppose that an HS pattern is detected in the i-th window. Then the return conditional on observing an HS pattern is defined as

(3) $$r_{i,c} = \ln\left(\frac{P_{i+n+c}}{P_{i+n}}\right), c = 20, 40, 60.$$

The trading strategy requires that one open a short position in the relevant stock. Since the trader is liable for any dividend payments while the position remains open, we make an adjustment to the conditional return when a dividend is paid during the relevant period. The excess return is then calculated by subtracting the daily three-month Treasury bill rate compounded continuously over the same holding period.

The conditional excess returns are calculated for two cases. The first is with restrictions (R1) to (R3), (R4) and (R5), and the second is with (R1) to (R3), (R4a), (R5a), and (R6) to (R9). In other words, they are calculated without and with the Bulkowski restrictions.

II. Data and Descriptive Findings

This section describes the empirical distributions of the conditional excess returns for all stocks in the S&P 500 and the Russell 2000 indices.

A. Stock Market Data

The data sets are based on the price and dividend series for the companies in the S&P 500 and the Russell 2000 over the period 1990-1999. The daily stock price and dividend data came from the CRSP database and were downloaded from the Wharton Research Data Services (WRDS) website. The CRSP database reports prices as they were traded, not adjusting for stock dividends, stock splits, mergers, and the like. To allow for a comparison of the price and dividend levels across time, WRDS has created a cumulative price and dividend adjustment factor (query item: CFACPR – cumulative factor to adjust prices). To calculate the levels of prices and dividends adjusted for stock splits and stock dividends, price (PRC) and dividend (DIVAMT) series were divided by the cumulative adjustment factor (CFACPR).

The companies in the S&P 500 represent approximately eighty-five percent of the total U.S. market capitalization. The Russell 2000 Index is based on the 2000 smallest companies in the Russell 3000 Index, and represents approximately eight percent of the total U.S. market capitalization. The companies in our S&P 500 and Russell 2000 data sets are the companies listed in the indices for June of 1990. In Figure 2 we present a time series plot of the number of patterns terminating each quarter for the S&P 500, without imposing the Bulkowski restrictions and using a bandwidth multiple of 2.5. No obvious trend or pattern is discernible. The same is true of the plot for the Russell 2000 which is therefore omitted.

B. Distributions of Conditional Excess Returns

B.1. S&P 500

Table I reports the estimated means of the excess returns with and without the Bulkowski restrictions. The results are for 20, 40 and 60 trading days and for the bandwidth multiples 1, 1.5, 2, and 2.5. The null hypothesis that the true means for excess returns with and without the Bulkowski restrictions are zero is rejected by the asymptotic 95% confidence intervals for 20, 40 and 60 trading days as well as for all bandwidth multiples. The convention we have adopted means that a negative excess return corresponds to a profitable trading strategy. Since the mean excess returns are all significantly positive, this indicates that the HS trading strategy considered in isolation would have been unprofitable. However, it is important to remember that the strategy involves taking a short position, and that our sample period coincides with an episode of exceptionally high stock returns that has been characterized as a bubble. A simple strategy of shorting the market portfolio and investing the proceeds in Treasury bills would have earned the negative of the equity premium over the period, or – 11.4 per cent per annum. In such a situation it is particularly important to look at risk-adjusted returns. We consider this in Section II.D.

B.2. Russell 2000

Table II reports the estimated means of the excess returns with and without the Bulkowski restrictions. The results are for 20, 40 and 60 trading days and for the bandwidth multiples 1, 1.5, 2, and 2.5. The null hypothesis that the true means for excess returns without the Bulkowski restrictions are zero is rejected by the asymptotic 95% confidence intervals for 60 trading days at all the bandwidth multiples. Without the Bulkowski restrictions, the annual excess return over 60 days using a bandwidth multiple of 2.5 is 3.3 per cent. Thus we find a marked difference between the unadjusted excess returns for the S&P 500 and Russell 2000. But even the excess returns for the HS strategy applied to Russell 2000 stocks fall substantially short of the equity premium over the period. Again, appropriate risk adjustment is essential.

Confidence intervals reported in the tables for the sample means of S&P 500 and Russell 2000 were not corrected for the presence of autocorrelation and heteroskedasticity. The qualitative results were identical, when autocorrelation- and heteroskedasticity-consistent covariance matrix estimators were used in the construction of the confidence intervals.

C. Impact of Bulkowski Restrictions

Tables I and II give the number of HS patterns detected with and without the Bulkowski restrictions. The results show that the number of patterns detected is substantially less when the Bulkowski restrictions are imposed. For example, using the conditional returns of the Russell 2000 stocks, 1990-1999, for 60 trading days and a bandwidth multiple of 2.5, the number of HS patterns detected drops from 6,406 to 4,452. This is a decrease of 30 per cent. However, it is clear from Tables I and II that imposition of the restrictions has no impact on the excess returns. When we consider risk-adjusted returns, again we find no effect of the restrictions on the magnitude of the intercept in the factor regressions. Thus we find no evidence to suggest that fine tuning the algorithm with the calibrations based on Bulkowski adds anything to its predictive power.

D. Risk-Adjustment of the Excess Returns

To test the hypothesis that HS patterns are able to predict risk-adjusted excess returns, we use the Fama-French three-factor model (Fama and French, 1993).

We regress monthly returns to the HS trading strategy on the three factor returns plus a constant. A detailed description of the construction of the monthly returns is contained in Appendices I and II[2]. The selected results for the S&P500 are presented in Table III. The asymptotic ninety five per cent confidence intervals are given for the intercept and factor loadings. All of the confidence intervals for regression coefficients were calculated using autocorrelation- and heteroskedasticity-consistent covariance matrix estimators. The results show that the intercept is significantly negative at the 60-day horizon for both bandwidth multiples. At that horizon with a bandwidth multiple of 2.5, the intercept is – 0.4 per cent per month, or – 4.8 per cent per annum. This means that the HS strategy in this case earns a positive risk adjusted excess return of 4.8 per cent per annum.

In order to interpret the factor loadings it is useful to have estimates of the factor risk premia for our sample period. Since the factors are zero-wealth portfolios the factor risk premia are equal to the sample means. For the market excess return, size, and book-to-market the annualized means in percent are respectively 11.4, – 2.5, and – 0.6. Only the market risk premium is significantly different from zero at the five per cent level. Given our convention of reporting the profitable return to a short position as negative, the loading on the market excess return reveals that the returns to the HS strategy (with profits measured as positive returns) are *negatively* correlated with the market. At the 60-day horizon the HS strategy has a market beta of – 0.77 for a bandwidth multiple of 2.5. If we focus on the market factor alone and ignore the other two factors, which play only a minor role, then we would expect the HS strategy to earn – 0.77 times the equity premium over the period, or – 8.8 per cent per annum. The fact that the

[2] Data on factor portfolio excess returns was obtained from Ken French's web site. (http://mba.tuck.dartmouth.edu/pages/faculty/ken.french/data_library.html). Note that monthly returns to the HS strategy are not precisely comparable to the 20, 40 and 60 day returns reported in Tables I and II. This is a

strategy, although it earns significantly negative raw excess returns, does much less poorly is the source of its success. So for example the 60-day raw excess return for bandwidth multiple 2.5 is 0.0083 (see Table I) which translates into a raw excess return of – 3.5 per cent per annum (250/60 x 0.83 per cent) and a risk-adjusted excess return of 4.8 per cent as reported above. In other words, the strategy is able to select stocks that will underperform relative to the market. Since the strategy involves taking short positions, this underperformance is advantageous.

In Table IV we present the selected results of the three-factor regression for the Russell 2000. As in Table III, the ninety five per cent confidence intervals are presented for the intercept and factor loadings. We find that the intercept is significantly negative for both bandwidths at the 60-day horizon, and for the narrower bandwidth at the 20-day horizon. Over a three-month horizon with a bandwidth multiple of 1, the intercept is – 0.55 per cent per month or – 6.6 per cent per annum. For the larger multiple of 2.5, the intercept is – 0.45 per cent per month or – 5.4 per cent per annum. The factor loading on the market return is somewhat reduced compared to the S&P500 and the loading on the size factor rises substantially. The latter effect is to be expected since the Russell 2000 contains relatively small firms and the size factor portfolio consists of a long position in small firms and a short position in large firms.

E. The Role of Momentum

If we consider the schematic representation of the HS pattern in Figure 1, it is evident that its evolution from the occurrence of the head of the pattern involves a declining trend. By construction this will occur over a time period of less than three months, in most cases considerably less. This suggests the hypothesis that negative momentum may explain our results.

---

consequence of the fact that monthly returns are calculated by averaging over the number of positions open within the month, which varies from month to month.

To investigate this, we augment the three-factor model with an additional momentum factor. The results of the monthly return regressions for this model are presented in Tables V and VI.

We find strong evidence of a significant loading on the momentum factor. The negative factor loading implies that HS excess returns are positively correlated with momentum. Since the momentum factor is constructed as a zero wealth portfolio consisting of a long position in positive momentum stocks (stocks whose prices have risen unusually fast over the past year) and a short position in negative momentum stocks (stocks whose prices have fallen unusually fast over the past year) this indicates that some of the predictive power of the HS pattern arises from the fact that it signals short positions in negative momentum stocks. However, we find that this is not a complete explanation for our results. For the Russell 2000 at the 60-day horizon and a narrow bandwidth parameter the excess return is still 41 basis points per month or 4.9 per cent per annum. In the case of the S&P500, at the 60-day horizon and a narrow bandwidth parameter the excess return is 18 basis points per month, but this rises to 25 basis points per month with a broad bandwidth.

The finding that momentum can provide a partial explanation for the predictive power of the HS price pattern is interesting in that the use of the pattern by market participants long predated the discovery of the momentum effect in the academic literature. The pattern was described at least as early as 1930 (see Schabacker, 1930 as cited in Chang and Osler, 1999).

F. The Effect of Transaction Costs

There is considerable disagreement over the appropriate adjustment for transaction costs for stock trading strategies. Jegadeesh and Titman (1993) cite a figure for one-way transaction costs of 23 basis points for an institutional trader (Berkowitz, Logue and Noser, 1988) and describe the figure of 0.5% that they use as “conservative”. However, the figure from Berkowitz et al. does not consider short positions, for which costs may be significantly higher. Lesmond, Ogden and Trczinka (1999) use a different methodology and arrive at higher figures for transaction costs, particularly for small stocks. Jones (2002) concludes that the one-way transaction cost for NYSE stocks had fallen to less than 18 basis points in 2000.

To assess the impact of transaction costs on returns we consider the one-way breakeven transaction cost in various cases. We look first at the raw excess returns. For the S&P 500 in Table I there is no nonnegative breakeven figure because the trading strategy is loss-making. For the Russell 2000 in Table II the breakeven cost will be half the raw excess return. For a bandwidth of 2.5 without Bulkowski restrictions we find the figures for 20, 40 and 60-day holding periods are 6, 26 and 39.5 basis points respectively. This indicates that only for the 60-day holding period does it appear that raw excess returns net of transaction costs remain positive. The remaining margin of benefit is slender and would disappear if one were to choose a somewhat more conservative figure for transaction costs. Thus we find very little support for the claim that using the HS pattern as a basis for a stand-alone trading strategy is profitable.

Next we turn to Table III, which reports risk-adjusted returns for the S&P 500. Since we report monthly returns, we need to use the monthly trading frequency to work out the one-way break-even transaction cost. For a 20-day holding period there will be a single one-way transaction per month for each pattern observed, so that the measured excess return is also the break-even transaction cost. For the S&P 500 this is not significantly different from zero. For the Russell 2000 the figure lies between 17 and 30 basis points depending on the bandwidth. For a 60-day holding period, there will be two one-way transactions every four months, so that twice the monthly excess return gives us the break-even transaction cost. For a bandwidth of 2.5 we obtain a breakeven figure of 80 basis points. For the Russell 2000 the corresponding figure is 90 basis points. These numbers are substantially above the figure of 18 reported by Jones (2002) or that of 23 suggested by Jegadeesh and Titman. Alternatively, if we assume a one-way transaction cost of 20 basis points, midway between the two numbers, then the impact of transactions costs on monthly excess returns is to reduce them by 10 basis points for a 60-day holding period. Therefore, for the S&P 500, and a bandwidth of 2.5, the excess return net of transaction costs is 3.6 percent per annum. For the Russell 2000 the comparable figure is 4.2 percent. These are certainly economically significant.

An alternative, rather more transparent way of quantifying the potential benefits generated by risk-adjusted excess returns is to calculate the increase in excess return that a mean-variance investor would have earned if he had split his wealth optimally between the market portfolio and the HS strategy so as to hold a portfolio with the same standard deviation as the market. We also calculate the portfolio weights to determine the extent to which the optimal portfolio was levered. In order to explain the calculations we define the following notation:

(4) $$r = (r_m, r_{hs})$$

where r_m, r_{hs} are the excess returns of the market portfolio and the HS strategy respectively. We denote by r_o the excess return of the optimal portfolio with standard deviation equal to σ_m, the standard deviation of the market portfolio. V is the (2x2) covariance matrix of excess returns of the market portfolio and the HS strategy.

Then the vector of portfolio weights w on the risky assets in the optimal portfolio is given by:

(5) $$w = \frac{r_O}{r'V^{-1}r} V^{-1} r$$

and the total excess return r_o on the optimal portfolio is given by:

(6) $$r_O = \sigma_m \sqrt{r'V^{-1}r}\ .$$

The results for the Russell 2000 are given in Table VII. The HS strategy raw excess return for the 20-day horizon and unit bandwidth multiple is insignificantly different from zero. However the correlation with the market excess return is -0.8. The optimal portfolio is highly levered, with portfolio weights of 1.66 on the market portfolio and 2.18 on the HS strategy. This leads to an increase in excess return over the market of 6.9 percent. The results are sensitive to the size of correlation, but are in excess of 2 percent in all cases. We see that the smallest increase occurs for the 20-day horizon and multiple 2.5, consistent with the fact that the intercept in the 3-factor regression in Table IV is not significantly different from zero. Results for the S&P500 (not presented) are less striking, but at the 60-day horizon the optimal portfolio earns an extra 5.5 percent (unit multiple) and 4.7 percent (2.5 multiple). Even after appropriate adjustment for transaction costs the improvement in tradeoff between risk and return is economically significant.

III Summary and Conclusion

We develop an algorithm for the detection of HS patterns in stock prices. The algorithm, as in Lo et al. (2000) is based on a non-parametric smoothing procedure used to detect particular sequences of extrema in the price series. We augment these restrictions with additional ones sufficient to characterize a set of typical examples of HS patterns identified by a technical analyst (Bulkowski, 1997; Bulkowski, 2000). We also impose an additional "neckline crossing condition" considered by technical analysts to be an important component of the pattern.

We consider the predictive power of the HS pattern for two groups of stocks, the S&P 500 and the Russell 2000, focusing explicitly on the mean excess return conditional on the occurrence of the pattern over the subsequent one, two and three months. We concentrate on evaluating mean return because technical trading manuals are unanimous in interpreting the occurrence of HS patterns as a signal of an imminent decline in the stock price.

We find no evidence that a stand-alone trading strategy based on HS patterns applied to S&P 500 stocks would be profitable and only weak evidence for stocks in the Russell 2000. However, for both groups of stocks, we find that the pattern predicts a decline in price relative to the market. Typical risk-adjusted excess returns are between five and seven per cent per annum. This provides a strong indication that the HS strategy is successfully timing short sales, since a policy of purely random short sales with the probability of selecting a stock weighted by market capitalization would earn the negative of the equity premium, which in our sample period is 11.4 per cent per annum. We find that the predictive ability of the pattern is partly due to its identification of negative momentum stocks.

The large difference between raw and risk-adjusted returns is largely a result of negative correlation with the market. This implies that HS patterns can be used successfully in conjunction with a passive indexing strategy to improve the tradeoff between risk and return. We show that the portfolio optimally combining the HS strategy with the market portfolio to reproduce the volatility of the market can yield an increase in excess return of up to 8 percent per annum.

References

Barberis, Nicholas, Andrei Shleifer and Robert Vishny, 1998, “A Model of Investor Sentiment,” Journal of Financial Economics 49, 307-343.

Berkowitz, Stephen A., Dennis E. Logue and Eugene A. Noser, 1988, “The Total Costs of Transactions on the NYSE,” Journal of Finance 43, 97-112.

Blume, Lawrence, David Easley and Maureen O'Hara, 1994, "Market Statistics and Technical Analysis: The Role of Volume," Journal of Finance 49, 153-181.

Brock, William, Josef Lakonishok and Blake LeBaron, 1992, “Simple Technical Trading Rules and the Stochastic Properties of Stock Returns,” Journal of Finance 47, 1731-1764.

Brown, David P., and Robert H. Jennings, 1989, “On Technical Analysis,” The Review of Financial Studies 2, 527-51.

Bulkowski, Thomas N., 1997, “The Head and Shoulders Formation,” Technical Analysis of Stocks and Commodities 15:8, 366-372.

Bulkowski, Thomas N., 2000, *Encyclopedia of Chart Patterns*, John Wiley and Sons: New York.

Chang, P.H. Kevin and Carol L. Osler, 1999, “Methodical Madness: Technical Analysis and the Irrationality of Exchange Rate Forecasts,” Economic Journal 109, 636-652.

Chopra, Navin, Josef Lakonishok and Jay R. Ritter, 1992, “Measuring Abnormal Performance: Do Stocks Overreact?,” Journal of Financial Economics 31, 235-268.

Daniel, Kent, David Hirshleifer and Avanidhar Subrahmanyam, 1998, “Investor Psychology and Security Market Under- and Overreaction,” Journal of Finance 53, 1839-1885.

DeBondt, Werner F. M. and Richard H. Thaler, 1985, “Does the Stock Market Overreact?”, Journal of Finance 40, 793-808.

Dooley, Michael P. and Jeffrey R. Shafer, 1983, "Analysis of short-run exchange rate behavior: March 1973 to November 1981", in D. Bigman and T. Taya (eds.), Exchange rate and trade instability: causes, consequences, and remedies, Ballinger: Cambridge, Mass, 43-69.

Edwards, Robert D and John Magee, 1992, *Technical Analysis of Stock Trends*, 6th edition, John Magee Inc.: Boston, Massachusetts.

Fama, Eugene F., 1965, "The Behavior of Stock Market Prices," Journal of Business 38, 34–105.

Fama, Eugene F., 1970, "Efficient Capital Markets: A Review of Theory and Empirical Work," Journal of Finance, 383-417.

Fama, Eugene F. and Kenneth R. French, 1993, "Common Risk Factors in the Returns on Stocks and Bonds," Journal of Financial Economics 33, 3-56.

Härdle, Wolfgang, 1990, Applied nonparametric regression, Cambridge University Press.

Jegadeesh, N., 1990, "Evidence of Predictable Behavior of Security Returns," Journal of Finance 45, 881-898.

Jegadeesh, N. and S. Titman, 1993, "Returns to Buying Winners and Selling Losers: Implications for Stock Market Efficiency," Journal of Finance 48, 65-91.

Jegadeesh, N. and S. Titman, 1995, "Overreaction, Delayed Reaction and Contrarian Profits," Review of Financial Studies 8, 973-993.

Jones, Charles M., 2002, "A Century of Stock Market Liquidity and Trading Costs," Working Paper.

Lesmond, David A., Joseph P. Ogden and Charles A. Trczinka, 1999, "A New Estimate of Transaction Costs," Review of Financial Studies 12, 1113-1141.

Levich, R., and L. Thomas, 1993, "The Significance of Technical Trading Rule Profits in The Foreign Exchange Market: A Bootstrap Approach," Journal of International Money and Finance 12, 451-74.

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

Neely, Christopher J., Paul A. Weller and Robert Dittmar, 1997, "Is Technical Analysis in the Foreign Exchange Market Profitable? A Genetic Programming Approach," Journal of Financial and Quantitative Analysis 32, 405-426.

Pring, Martin J., 1991, Technical Analysis Explained, Second Edition, McGraw-Hill, New York.

Schabacker, R.W., 1930, Stock Market Theory and Practice, Forbes Publishing, New York, N.Y.

Sweeney, Richard J., 1986, "Beating the foreign exchange market", Journal of Finance 41, 163-182.

Sullivan, Ryan, Allan Timmermann and Halbert White, 1999, "Data-Snooping, Technical Trading Rule performance, and the Bootstrap," Journal of Finance 54, 1647-1691.

Treynor, Jack L. and Robert Ferguson, 1985, "In Defense of Technical Analysis," Journal of Finance 40, 757-773.

Table I

Means, Variances and Confidence Intervals of Conditional Excess Returns: S&P 500, 1990-1999

The table reports the means, variances and confidence intervals for the excess returns conditional on detecting an HS pattern when the span of the rolling windows is n = 63. The reported returns are 20, 40 and 60 day returns. A negative excess return corresponds to a profitable HS trading strategy.

Trading Days	Bandwidth Multiple	Mean	St. Dev.	95 % Confidence Interval Lower	Upper	Number of patterns
Without Bulkowski restrictions						
20	1	0.0053	0.0987	0.0037	0.0069	14244
	1.5	0.0045	0.0844	0.0027	0.0062	9160
	2	0.0056	0.0935	0.0032	0.0079	5972
	2.5	0.0057	0.0791	0.0031	0.0083	3550
40	1	0.0064	0.1411	0.0041	0.0087	14244
	1.5	0.0035	0.1283	0.0009	0.0061	9160
	2	0.0053	0.1289	0.0020	0.0085	5972
	2.5	0.0053	0.1204	0.0014	0.0093	3550
60	1	0.0097	0.1722	0.0068	0.0125	14244
	1.5	0.0054	0.1560	0.0022	0.0086	9160
	2	0.0078	0.1555	0.0039	0.0118	5972
	2.5	0.0083	0.1483	0.0034	0.0132	3550
With Bulkowski restrictions						
20	1	0.0054	0.0876	0.0036	0.0073	8676
	1.5	0.0041	0.0845	0.0020	0.0062	6165
	2	0.0054	0.1001	0.0023	0.0085	4086
	2.5	0.0058	0.0812	0.0026	0.0090	2463
40	1	0.0053	0.1417	0.0023	0.0083	8676
	1.5	0.0030	0.1332	-0.0003	0.0064	6165
	2	0.0055	0.1340	0.0014	0.0096	4086
	2.5	0.0057	0.1224	0.0009	0.0106	2463
60	1	0.0083	0.1699	0.0047	0.0119	8676
	1.5	0.0048	0.1607	0.0008	0.0088	6165
	2	0.0080	0.1597	0.0031	0.0129	4086
	2.5	0.0083	0.1520	0.0023	0.0143	2463

Table II

Means, Variances and Confidence Intervals of Conditional Excess Returns: Russell 2000, 1990-1999

The table reports the means, variances and confidence intervals for the excess returns conditional on detecting an HS pattern when the span of the rolling windows is n = 63. The reported returns are 20, 40 and 60 day returns. A negative excess return corresponds to a profitable HS trading strategy.

Trading Days	Bandwidth Multiple	Mean	St. Dev.	95 % Confidence Interval Lower	Upper	Number of patterns
Without Bulkowski restrictions						
20	1	-0.0014	0.1661	-0.0032	0.0005	30675
	1.5	-0.0026	0.1619	-0.0049	-0.0002	18566
	2	-0.0023	0.2220	-0.0064	0.0018	11073
	2.5	-0.0012	0.2213	-0.0066	0.0042	6406
40	1	-0.0031	0.2332	-0.0057	-0.0005	30675
	1.5	-0.0048	0.2280	-0.0081	-0.0015	18566
	2	-0.0065	0.2638	-0.0114	-0.0016	11073
	2.5	-0.0052	0.2647	-0.0117	0.0013	6406
60	1	-0.0057	0.2796	-0.0089	-0.0026	30675
	1.5	-0.0074	0.2836	-0.0115	-0.0033	18566
	2	-0.0079	0.3240	-0.0140	-0.0019	11073
	2.5	-0.0079	0.3053	-0.0153	-0.0004	6406
With Bulkowski restrictions						
20	1	-0.0005	0.1813	-0.0030	0.0019	20800
	1.5	-0.0018	0.1696	-0.0047	0.0011	12998
	2	0.0003	0.2375	-0.0049	0.0056	7813
	2.5	0.0001	0.2306	-0.0066	0.0069	4452
40	1	-0.0032	0.2469	-0.0066	0.0001	20800
	1.5	-0.0048	0.2439	-0.0090	-0.0007	12998
	2	-0.0052	0.2853	-0.0115	0.0012	7813
	2.5	-0.0032	0.2758	-0.0113	0.0049	4452
60	1	-0.0067	0.2980	-0.0107	-0.0026	20800
	1.5	-0.0086	0.2843	-0.0135	-0.0037	12998
	2	-0.0095	0.3228	-0.0167	-0.0024	7813
	2.5	-0.0069	0.3172	-0.0162	0.0024	4452

Table III

Regression Coefficients and Confidence Intervals in the Three-Factor Regression: S&P 500, 1990-1999

The table reports the regression coefficients and their 95% confidence intervals in the three-factor linear regression, where the dependent variables consist of monthly excess returns conditional on detecting an HS pattern when the span of the rolling windows is n = 63. The returns are reported for 20 and 60-day windows without applying the Bulkowski restrictions. The bandwidth parameter given in the first column is either 1 or 2.5. Autocorrelation- and heteroskedasticity-consistent standard errors are used to construct the confidence intervals.

	Intercept	Excess Market Return Factor	Size Factor	Book-to-Market Factor
		20 days		
1	-0.00106 (-0.0027;0.0006)	0.50840 (0.4608;0.5560)	0.11381 (0.0529;0.1747)	0.12574 (0.0537;0.1978)
2.5	0.00045 (-0.0025;0.0016)	0.49875 (0.4408;0.5567)	0.05112 (-0.0230;0.1253)	0.12637 (0.0387;0.2141)
		60 days		
1	-0.00360 (-0.0057;-0.0015)	0.76299 (0.7036;0.8224)	0.13902 (0.0661;0.2119)	0.26127 (0.1742;0.3483)
2.5	-0.00400 (-0.0063;-0.0017)	0.77032 (0.7048;0.8358)	0.11605 (0.0357;0.1964)	0.34034 (0.2443;0.4363)

Table IV

Regression Coefficients and Confidence Intervals in the Three-Factor Regression: Russell 2000, 1990-1999

The table reports the regression coefficients and their 95% confidence intervals in the three-factor linear regression, where the dependent variables consist of monthly excess returns conditional on detecting an HS pattern when the span of the rolling windows is n = 63. The returns are reported for 20 and 60-day windows and without applying the Bulkowski restrictions. The bandwidth parameter given in the first column is either 1 or 2.5. Autocorrelation- and heteroskedasticity-consistent standard errors are used to construct the confidence intervals.

	Intercept	Excess Market Return Factor	Size Factor	Book-to-Market Factor
		20 days		
1	-0.00304 (-0.0046;-0.0015)	0.45727 (0.4125;0.5020)	0.42943 (0.3721;0.4867)	0.11371 (0.0459;0.1815)
2.5	-0.00172 (-0.0047;0.0012)	0.4194 (0.3359;0.5028)	0.55467 (0.4479;0.6615)	0.10392 (-0.0224;0.2303)
		60 days		
1	-0.00549 (-0.0075;-0.0034)	0.70630 (0.6483;0.7643)	0.63034 (0.5591;0.7016)	0.27149 (0.1864;0.3566)
2.5	-0.00449 (-0.0071;-0.0018)	0.66399 (0.5892;0.7388)	0.76310 (0.6713;0.8549)	0.28924 (0.1796;0.3989)

Table V

Regression Coefficients and Confidence Intervals in the Four-Factor Regression: S&P 500, 1990-1999

The table reports the regression coefficients and their 95% confidence intervals in the four-factor linear regression, where the dependent variables consist of monthly excess returns conditional on detecting an HS pattern when the span of the rolling windows is n = 63. The returns are reported for 20 and 60-day windows without applying the Bulkowski restrictions. The bandwidth parameter given in the first column is either 1 or 2.5. Autocorrelation- and heteroskedasticity-consistent standard errors are used to construct the confidence intervals.

	Intercept	Excess Market Return Factor	Size Factor	Book-to-Market Factor	Momentum Factor
			20 days		
1	-0.00003 (-0.0017;0.0016)	0.49975 (0.4549;0.5446)	0.08109 (0.0217;0.1405)	0.09552 (0.0262;0.1648)	-0.10784 (-0.1615;-0.0542)
2.5	0.00025 (-0.0019;0.0024)	0.49287 (0.4356;0.5502)	0.02885 (-0.0470;0.1047)	0.10580 (0.0174;0.1942)	-0.07339 (-0.1419;-0.0049)
			60 days		
1	-0.00182 (-0.0038;0.0002)	0.74740 (0.6941;0.8007)	0.08795 (0.0204;0.1555)	0.20255 (0.1221;0.2830)	-0.17513 (-0.2386;-0.1117)
2.5	-0.00249 (-0.0048;-0.0002)	0.75707 (0.6952;0.8189)	0.07263 (-0.0059;0.1511)	0.29042 (0.1969;0.3839)	-0.14888 (-0.2226;-0.0752)

Table VI

Regression Coefficients and Confidence Intervals in the Four-Factor Regression: Russell 2000, 1990-1999

The table reports the regression coefficients and their 95% confidence intervals in the four-factor linear regression, where the dependent variables consist of monthly excess returns conditional on detecting an HS pattern when the span of the rolling windows is n = 63. The returns are reported for 20 and 60-day windows without applying the Bulkowski restrictions. The bandwidth parameter given in the first column is either 1 or 2.5. Autocorrelation- and heteroskedasticity-consistent standard errors are used to construct the confidence intervals.

	Intercept	Excess Market Return Factor	Size Factor	Book-to-Market Factor	Momentum Factor
			20 days		
1	-0.00258 (-0.0042;-0.0009)	0.45333 (0.4088;0.4978)	0.41454 (0.3556;0.4735)	0.09996 (0.0313;0.1687)	-0.04907 (-0.1023;0.0041)
2.5	-0.00117 (-0.0043;0.0019)	0.41472 (0.3310;0.4984)	0.53704 (0.4261;0.6479)	0.08763 (-0.0416;0.2169)	-0.05814 (-0.1582;0.0420)
			60 days		
1	-0.00406 (-0.0061;-0.0021)	0.69377 (0.6395;0.7480)	0.58928 (0.5204;0.6581)	0.22429 (0.1423;0.3063)	-0.14077 (-0.2054;-0.0761)
2.5	-0.00257 (-0.0052;0.0000)	0.64719 (0.5777;0.7167)	0.70808 (0.6199;0.7962)	0.22598 (0.1210;0.3310)	-0.18866 (-0.2715;-0.1059)

Table VII

Characteristics of Combined Portfolio: Russell 2000, 1990-1999

The table reports the monthly mean and standard deviation of the excess return on the market portfolio and the HS strategy for stocks in the Russell 2000, the correlation $C(r_m, r_{hs})$ between the two excess returns, and the optimal weights for a portfolio with the same standard deviation as the market excess return. The excess return for the HS strategy uses the convention that a positive number represents a positive excess return to the short position. The last two columns report the annualized excess return for the optimal portfolio, and the difference in annualized excess returns $(r_o - r_m)$.

		Mean	SD	$C(r_m, r_{hs})$	Optimal weights	Excess Return r_o	Difference r_o - r_m
				20 days			
1	Mkt	0.0089	0.0389	-0.7975	1.657	0.1882	0.0688
	HS	-0.0001	0.0232		2.184		
2.5	Mkt	0.0089	0.0389	-0.6448	1.300	0.1367	0.0220
	HS	-0.0008	0.0276		1.024		
				60 days			
1	Mkt	0.0095	0.0389	-0.7930	1.641	0.2099	0.0806
	HS	0.0002	0.0339		1.513		
2.5	Mkt	0.0095	0.6472	-0.7110	1.422	0.1743	0.0485
	HS	-0.0001	0.0360		1.085		

Notes: the monthly HS excess returns differ from those reported in Table II because of the different averaging procedures used. When HS returns are calculated over calendar months (as here) a position held for part of a month is treated as earning a zero excess return over the remainder of the month, resulting in a lower figure for mean return. The market excess return for 20-day and 60-day horizons differs because the former uses a sample shortened by two calendar months at the end of the period. Figures for excess return less difference do not give exactly the same numbers because of annualization.

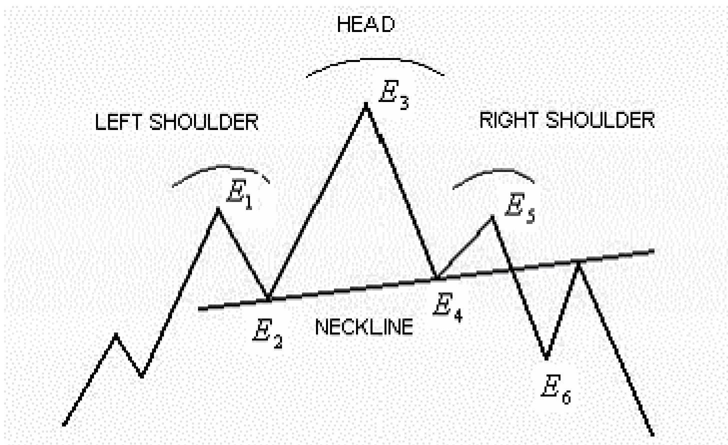

Figure 1 The head-and-shoulders pattern

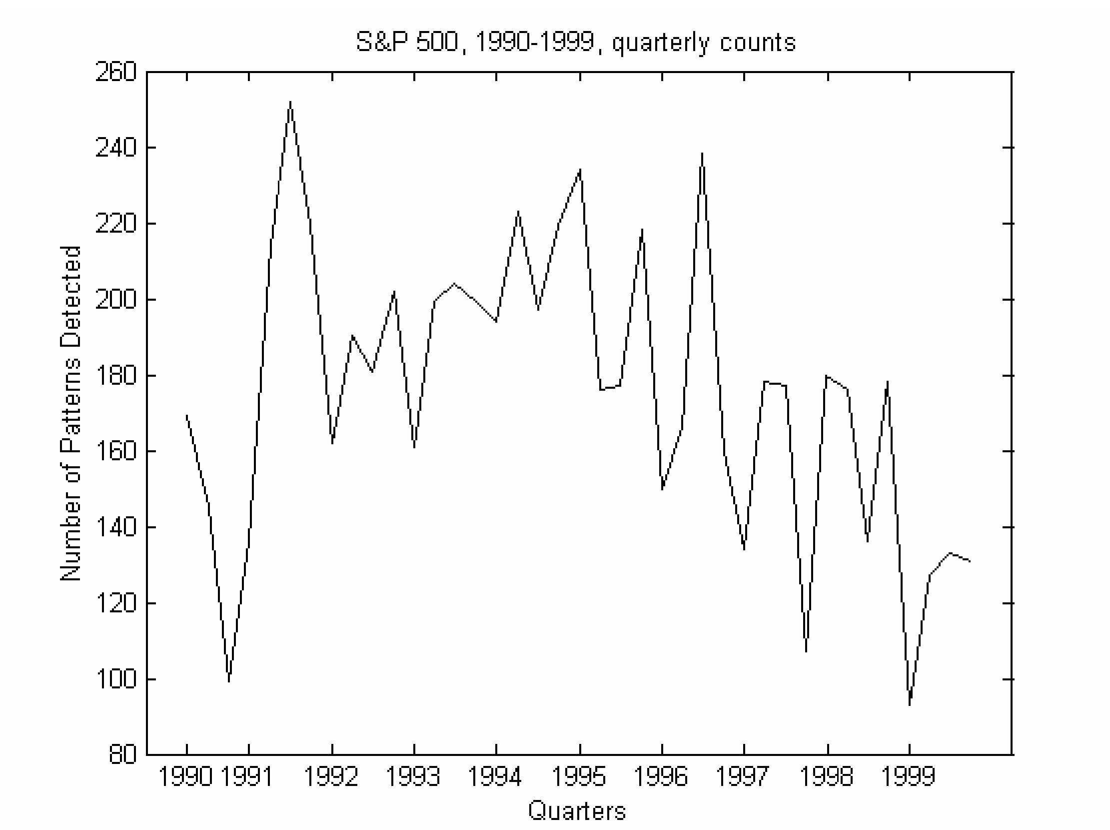

Figure 2 Quarterly occurrences of head-and-shoulders patterns for the S&P 500: 1990-1999

The figure records the number of head-and-shoulders patterns terminating each quarter, without imposing the Bulkowski restrictions and using a bandwidth multiple of 2.5.

Appendix I: Calculation of HS monthly returns

The calculation of monthly returns from the short positions is complicated by the fact that the majority of the positions start on dates that are different from the end or beginning of a particular month. However, the excess returns for the factor portfolios are monthly returns.

The calculation uses the following input data: the index of the business day on which the short position was opened and closed, the price data for the given stock, and the daily 3-month Treasury bill rate. Given the price of the stock at the starting date of the short position $P_{i,j}$, where i denotes the month and j denotes the day of the month, we find the continuously compounded return for that particular month using the following formula:

$$r_{i,j} = \log\left(\frac{P_{i,end}}{P_{i,j}}\right), \tag{A.1}$$

where $P_{i,end}$ denotes the price on the last business day of month i. The excess return $r^e_{i,j}$ is found as follows:

$$r^e_{i,j} = r_{i,j} - \sum_{t=j}^{end} r_t \ , \tag{A.2}$$

where r_t is the daily 3-month Treasury bill rate continuously compounded and *end* denotes the last business day of the month. Appropriate adjustment in compounding the risk-free rate is made for non-business days. Note the convention that a profitable trade is associated with a negative excess return.

If the short position is two or three months long, it will span three or four months. The monthly returns for the months that are fully spanned by the short position are calculated using the following formula:

$$r_{i,i+1} = \log\left(\frac{P_{i+1,end}}{P_{i,end}}\right). \tag{A.3}$$

The excess return is found as follows:

$$r_{i,i+1}^{e} = r_{i,i+1} - \sum_{t=1}^{end} r_t \ . \tag{A.4}$$

For the majority of short positions, the ending date of the short position will be different from the last day of a given month. To find the return from the tail end of the short position, i.e., for the time period from the beginning of the last month until the date when the short position is closed, we use the following formula:

$$r_{i,j} = \log\left(\frac{P_{i,j}}{P_{i-1,end}}\right) \tag{A.5}$$

where $P_{i,j}$ is the price at which the short position is closed.

The excess return is found as follows:

$$r_{i,j}^{e} = r_{i,j} - \sum_{t=1}^{j} r_t \ . \tag{A.6}$$

The following is a derivation of the above formula. The notional price at which the short position is opened is $P_{i-1,end}$. This is therefore the notional cash inflow from opening the short position. A cash outflow of $P_{i,j}$ occurs on day *j*. To calculate the monthly return we need to compound these two positions to the end of the month. Therefore the excess return on these two positions is:

$$r_{i,j}^{e} = \ln\left(P_{i,j} e^{\sum_{t=j+1}^{end} r_t} \Big/ P_{i-1,end} e^{\sum_{t=1}^{end} r_t} \right) = r_{i,j} - \sum_{t=1}^{j} r_t \ . \tag{A.7}$$

Finally, the average excess return is calculated for any given month by summing up all the excess returns from short positions for that given month across all companies and then dividing through by the number of such returns.

Appendix II: Construction of the Risk Factor Portfolios and the Momentum Portfolio

The following description of the construction of the risk factors is based on that in Fama and French (1993).

The Fama-French benchmark factors are the excess return on the market (Rm-Rf), the size factor (Small-Minus-Big, SMB), and the book-to-market (B/M) factor (High-Minus-Low, HML).

Rm-Rf, the excess return on the market, is the value-weighted return on all NYSE, AMEX, and NASDAQ stocks (from CRSP) minus the one-month Treasury bill rate (from Ibbotson Associates).

To construct the size and B/M factors all NYSE, AMEX, and NASDAQ stocks are split into two groups, small (S) and big (B), in June of year t based on the median size of a firm on NYSE. They are separately split into three groups, high (H), medium (M) and low (L), based on book equity/market equity at the end of year $t - 1$. Six size/book-to-market portfolios are formed from these separate groups. Monthly value-weighted returns are calculated on the six portfolios from July of year t to June of year $t + 1$.

SMB is the average return, calculated monthly on the three small portfolios minus the average return on the three big portfolios. HML is the average return, calculated monthly on the two high B/M portfolios minus the average return on the two low B/M portfolios.

The momentum factor UMD (Up-Minus-Down) is constructed using six value-weighted portfolios formed on size and prior returns.[3] The portfolios, which are formed monthly are the intersections of two portfolios formed on size and three portfolios formed on return over the period from twelve months to two months previously. The monthly size breakpoint is the median NYSE market equity. The monthly prior (2-12) return breakpoints are the 30th and 70th NYSE percentiles. UMD is the average return on the two high prior return portfolios minus the average return on the two low prior return portfolios.

[3] This description is taken from Ken French's web site:
http://mba.tuck.dartmouth.edu/pages/faculty/ken.french/data_library.html