

The Predictive Power of “Head-and-Shoulders” Price Patterns in the U.S. Stock Market

Gene Savin†

Paul Weller*

Janis Zvingelis**

January 13, 2003

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 7-9 percent per year.

Keywords: technical analysis, kernel regression, stock prices

JEL subject numbers: G1, G11, G14

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

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

** Department of Finance
Henry B. Tippie College of
Business
University of Iowa
Iowa City, IA 52242
(319) 335-0980, (319) 335-
3690 (fax)
janis-zvingelis@uiowa.edu

Acknowledgements: We are grateful to the editor, Bruce Lehmann and a referee for very helpful comments, and to Ken French for making available the data on factor risk portfolios.

The Predictive Power of “Head-and-Shoulders” Price Patterns in the U.S. Stock Market

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. 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 (Treyner 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. Osler and Chang (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 a *change 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. We do not address the question of whether a profitable trading strategy net of transaction costs exists. Rather, we take the view expressed by Daniel, Hirshleifer and Teoh (2001, p.14) in analyzing patterns of return predictability. They observe that some patterns seem to be profitable net of transaction costs and some do not, and comment “In either case these patterns present a challenge for scientific explanation, and are relevant for policy.”

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*

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

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 Russell 2000 stocks and for S&P 500 stocks over the period 1990-1999. For the first group, which is comprised of smaller stocks, we find evidence of significant excess returns after the occurrence of a head-and-shoulder pattern over one, two and three-month windows. For the second group, which consists of stocks with relatively large capitalization, our findings are similar. Using the Fama-French three factor model augmented by the addition of a momentum factor, we find that risk-adjusted excess returns are economically and statistically significant over all time horizons, and seven to nine per cent per annum over a three-month window.

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 Osler and Chang (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.

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,

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

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. In contrast 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, $i=1, \dots, T-n+1$, is

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

that is often called the *Nadaraya-Watson estimator*. $K(\bullet)$ is the kernel, a function which satisfies certain conditions. In our estimation, $K(\bullet)$ 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 $m_{i,n}(x)$ 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 $m_{i,n}(x)$, 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 $m_{i,n}(x)$ 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, \dots, E_m denote the set of relevant extrema and $X_1^*, X_2^*, \dots, X_m^*$ 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) $\max_i |E_i - \bar{E}| \leq 0.015 \cdot \bar{E}$, $i = 1, 5$, where $\bar{E} = (E_1 + E_5)/2$.

(R5) $\max_i |E_i - \bar{E}| \leq 0.015 \cdot \bar{E}$, $i = 2, 4$, 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 three 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.

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

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

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:

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

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

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

$$(R9) \quad \max_i \left| (X_{i+1}^* - X_i^*) - \bar{X}^* \right| \leq 1.2 \cdot \bar{X}^*, \quad i = 1, \dots, 4, \quad \text{where } \bar{X}^* = \sum_{i=1}^4 (X_{i+1}^* - X_i^*)/4.$$

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.

In addition, we impose the requirement that E_5 occur at date $n-3$. This is a simple way to capture the fact that the final extremum must be identified before concluding that an HS pattern has been observed.

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. However, there is no clear consensus on the appropriate horizon. Accordingly, we calculate the continuously compounded return over the subsequent 20, 40 and 60 trading days.

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

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

The usual assumption according to the claims of technical analysts is that the sign of the conditional expected return is negative, and thus that a short sale is on average profitable. 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 series for the companies in the S&P 500 and the Russell 2000 over the period 1990-1999. 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

The main result for the 60-day excess returns is the following:

The mode of the empirical distribution is slightly positive. The mean is statistically significantly negative, which is explained by the long left tail of the empirical distribution.

This is illustrated in Figure 3. The figure presents the histogram and the kernel estimate of the density of the excess returns with the Bulkowski restrictions for a bandwidth multiple of 2.5. The histogram and the kernel estimate of the density are overlaid with a normal distribution that has the mean and variance of the empirical distribution. The histograms and kernel estimates of the density of the excess returns are very similar for 20 and 40 trading days and for smaller bandwidth multiples. The same is true without the Bulkowski restrictions.

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. The annual excess returns over the three horizons are quite similar. Without the Bulkowski restrictions, annual excess returns over 20, 40 and 60 days using a bandwidth multiple of 2.5 are 6.2, 6.2 and 5.8 per cent respectively. With the restrictions imposed, the corresponding figures are 5.8, 6.0 and 5.7 per cent.

B.2. Russell 2000

The main result for the 60-day excess returns is the following:

The mode of the empirical distribution is zero. The mean is statistically significantly negative, which is explained by the slightly larger probability weight on small negative values.

Unlike the rather skewed distributions of the excess returns for the S&P 500, the excess returns for the Russell 2000 have much more symmetrical distributions. The sample variances are slightly larger for the Russell 2000 than for the S&P 500. Figure 4 presents the histogram and the kernel estimate of the density of the excess returns with the Bulkowski restrictions for a bandwidth multiple of 2.5. The histogram and the kernel estimate of the density are overlaid with a normal distribution that has the mean and variance of the empirical distribution. The histograms and kernel estimates of the density of the excess returns are very similar for 20 and 40 trading days and for smaller bandwidth multiples. The same is true without the Bulkowski restrictions.

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 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 the bandwidth multiples. Again, the annual excess returns over the three horizons are quite similar. Without the Bulkowski restrictions, annual excess returns over 20, 40 and 60 days using a bandwidth multiple of 2.5 are 5.7, 6.4 and 6.2 per cent respectively. With the restrictions imposed, the corresponding figures are 5.4, 5.7 and 5.4 per cent.

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.

To provide some interpretation for these numbers, consider a strategy in which stocks are randomly selected for short selling. The expected excess return would be the *negative* of the equity premium, which over our sample period is 11.4 per cent per annum. Thus the difference in return between random short selling and the HS strategy is of the order of seventeen per cent per annum. This gives a strong initial indication that HS price patterns are indeed successfully predicting changes in price trends.

C. Risk-Adjustment of the Excess Returns

The fact that we have demonstrated the existence of significant excess returns to an HS trading strategy in itself provides no conclusive evidence that these returns are not simply compensation for the risk involved. To test the hypothesis that HS patterns are able to predict excess risk-adjusted returns, we use the Fama-French three-factor model. Although there is no evidence to suggest that price momentum can be interpreted as a risk factor, we also include a momentum factor for diagnostic purposes.

We regress monthly returns to the HS trading strategy on the four factor returns plus a constant. A detailed description of the construction of the monthly returns is contained in Appendices I and II². 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 for all three investment horizons and all values of 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)

bandwidth multiple. With a trading horizon of three months, a bandwidth multiple of 2.5 and the Bulkowski restrictions imposed, the intercept is -0.78 per cent per month, or -9.4 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 risk premia are equal to the sample means. For the market excess return, size, book-to-market and momentum the annualized means in percent are respectively 11.4, -2.5 , -0.6 and 12.1. Only the market and momentum risk premia are 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 i.e. the HS strategy on average generates profits when holding the market portfolio produces losses. This in turn indicates that the HS strategy is successfully timing the market as a whole in the sense that profits from short sales occur on average when the market is declining.

The loading on the size factor is much smaller and sometimes insignificant, and does not play an important role. The loading on the book-to-market factor is larger and always significant. But given that the factor risk premium is not significantly different from zero, again we conclude that the factor does not play an important role in our results. The picture changes when we look at the momentum factor. There is a significant loading on this factor which increases in absolute magnitude with the trading horizon. The loading ranges from -0.11 to -0.21 . The sign indicates that momentum returns are *positively* correlated with HS returns (with profits measured as positive returns). The momentum factor portfolio generates profits by taking long positions in firms that have done well over the previous months two to twelve, and short positions in those

that have done poorly over the same period. Since HS returns are generated by short sales, this suggests that they can in part be explained by negative momentum i.e. taking short positions in firms that have recently done poorly.

In Table IV we present the selected results of the four-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. Again we find that the intercept is significantly negative for all horizons, all bandwidths and with and without the Bulkowski restrictions imposed. Over a three-month horizon with a bandwidth multiple of 2.5 and with the Bulkowski restrictions imposed, the intercept is -0.58 per cent per month or -6.9 per cent per annum. The factor loading on the market return is somewhat reduced 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. The loading on the momentum factor is significantly negative in all cases. Thus we find that all the qualitative findings for the S&P500 carry over to the Russell 2000. All of the confidence intervals for regression coefficients were calculated using autocorrelation- and heteroskedasticity-consistent covariance matrix estimators.

We perform two further analyses. In the first we regress the HS excess returns on the market factor alone. This may give a better measure of risk-adjusted returns given the insignificant role played by the size and book-to-market factors, and the absence of a risk-based explanation for the return to the momentum portfolio. The results are presented in Table V. For the S&P 500 with a trading horizon of three months, a bandwidth of 2.5 and the Bulkowski restrictions imposed the intercept is -0.93 per cent per month or -11.2 per cent per year. For the Russell 2000 under the same conditions the intercept is -0.85 per cent per month or -10.2 per

cent per year. The increase in excess return is largely accounted for by the contribution of the momentum factor.

If we interpret the previous results within the framework of the CAPM rather than as an application of an ad hoc factor model, it is possible that the obvious departure of the excess returns from normality in Figures 3 and 4 may lead to unreliable inferences. In particular the skewness in the distribution of returns for the S&P 500 indicates a pattern of a large number of small losses outweighed by a small number of big gains. This in turn suggests that the strategy generates a payoff similar to that of a put option. It is well-known that the standard CAPM is unable to price options. However, Rubinstein (1976) demonstrates that an asset pricing model derived under the assumption of constant relative risk aversion preferences may be used to price arbitrary portfolios of state-contingent claims, and in particular options. Brown and Gibbons (1986) use the model to estimate the coefficient of relative risk aversion and are unable to reject the hypothesis that preferences are logarithmic. Since the stochastic discount factor for the log CAPM is the inverse of the gross return on the market portfolio, $1/r_m$, a simple prediction of the model is that the HS excess return r^e must satisfy $E[r^e/r_m] = 0$. On the assumption that this ratio is i.i.d. we test the hypothesis that its expected value is zero. The results of this test are reported in the last column of Table V. In all cases the ratio is statistically significantly negative. Our conclusions are thus shown to be robust to more general assumptions about the distribution of security returns than those required for the standard CAPM.

One respect in which we must qualify our results stems from the fact that we ignore the impact of dividends on the return to a short sale. In practice a short-seller is liable for any dividend payments made during the period of the short sale. However, the dividend yield on the S&P 500 averaged only 2.4 per cent per year over our sample period. This suggests that even

after accounting for dividends the risk-adjusted return to HS trading remains economically significant.

D. 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 13,569 to 9,483. 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.

E. Distributions of Pattern Characteristics

We examined the empirical distributions of the pattern characteristics to evaluate the effects of the restrictions (R4) to (R9) on number and type of HS patterns detected. The distributions are illustrated in Figures 5, 6 and 7 for the Russell 2000 and a bandwidth multiple of 2.5. Figure 5 shows that restrictions imposed by Lo et al. (2000), namely, (R4) and (R5) are much more stringent than our restrictions, (R4a) and (R5a). By contrast, Figures 6 and 7 indicate that restrictions (R6), (R7), (R8) and (R9) do not have substantial effects in eliminating potential patterns. When the bandwidth is one, however, the effects of (R6) and (R7) are much more pronounced. The effects of the restrictions are similar for the S&P 500 stocks. Patterns satisfying restrictions (R4) and (R5) lie to the left of the verticals at 0.015 in Figure 5. Patterns satisfying restrictions (R4a) and (R5a) lie to the left of the verticals at 0.04 in Figure 5. Patterns satisfying restrictions (R6) and (R7) lie between the two verticals in the upper panel of Figure 6. Patterns

satisfying (R8) lie to the right of the vertical in the lower panel of Figure 6. Finally, patterns satisfying (R9) lie to the left of the vertical in Figure 7.

III. Conditional versus Unconditional Returns

So far we have concentrated attention on the conditional mean of the distribution of excess returns, both unadjusted and risk-adjusted. To provide additional information on the performance of the HS trading strategy we examine the complete conditional distribution of excess returns and compare it to the unconditional distribution. This comparison on its own is not greatly informative. For example, there is an extensive literature documenting the fact that price volatility is predictable. So a significant change in conditional distribution might simply reflect a change in volatility of conditional returns. But coupled with our previous evidence that the HS trading strategy produces significant risk-adjusted returns, it can give us useful information for understanding how the HS strategy works.

Lo et al. (2000) also compared the conditional versus the unconditional distributions of returns. They made a comparison between the distribution of the post-HS pattern *one-day* returns and the distribution of all of the *one-day* returns within the relevant five-year period. Also, they normalized returns by subtracting the unconditional return and dividing by the unconditional standard deviation. The two distributions were compared using the Kolmogorov-Smirnov (KS) test and were found to be significantly different. However, if we focus on the means it is worth noting that for Nasdaq stocks, only two out of their five-quintile groupings and three out of seven of their five-year periods show a negative post-HS return. The signs are more consistently negative for NYSE/AMEX stocks, but the mean one-day return for this group is only 3.8 per cent of the daily standard deviation.

Our comparison of conditional and unconditional returns uses a different methodology. The conditional returns are based on 20, 40 and 60-day post-HS pattern windows where the patterns were those detected using a bandwidth of 2.5. To obtain the unconditional 20, 40 and 60-day returns with bandwidth 2.5 for each five-year period, we sample each stock for which an HS pattern was detected, constraining the number of windows chosen to be equal to the number of HS patterns detected for the stock. The sampling of the returns is done to make the size of unconditional sample equal to that of conditional sample. Also, restricting the unconditional returns to the stocks where the HS patterns were detected allows us to eliminate the across-the-stocks variation in returns as one of the sources of differences between unconditional and conditional distributions of the returns.

Figure 8 shows the quantile-quantile (QQ) plots for the Russell 2000 and S&P 500. The quantiles of the conditional distribution are plotted against the quantiles of the unconditional distribution. If the two distributions are the same, then the plot is a straight line with an angle of 45 degrees. As can be seen from the figures, the plots clearly depart from a straight line. The difference is most striking in the case of the S&P 500. The QQ plot starts substantially below the straight line, coincides with the straight line and then again eventually falls below it. This is explained by the fact that the conditional distribution has more probability in the left tail and less in the right tail than the unconditional distribution. The difference is less pronounced for the Russell 2000, and the shift in probability weight does not seem to occur in the extreme left tail of the distribution.

The QQ plots are functions that are subject to random variation due to sampling. In particular, the unconditional distribution varies from replication to replication whereas the conditional distribution remains fixed. However, the main features of the QQ plots remain the

same across replications. The shape of the QQ plots reflects the difference of the skewness coefficients of the conditional and unconditional distributions. For example, for the conditional and unconditional distributions employed in Figure 8, the skewness coefficients are -1.16507 and -0.84453 for the Russell 2000 and -2.87042 and 0.84214 for the S&P 500. As is revealed by the plots, the difference in skewness is much more marked for the S&P 500. This suggests that, at least for large stocks the HS pattern works by successfully predicting unusually steep price declines for a relatively small number of stocks.

IV 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 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. There is less agreement on the length of time over which the decline will occur. However 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.

For both groups of stocks, we find strong evidence that the pattern predicts a decline in price relative to the market. The distribution of excess returns conditional on the occurrence of an HS pattern tends to be skewed to the left and has a mean that is significantly negative. This result is rather robust. It is not sensitive either to the horizon over which the return is calculated or to the degree of data smoothing (governed by the magnitude of the bandwidth parameter). Typical excess returns are between five and six 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 would earn the negative of the equity premium, which in our sample period is 11.4 per cent per annum.

We use the Fama-French three-factor model augmented with a momentum factor to determine whether these results survive adjustment for factor risk. We find strong evidence of an economically and statistically significant intercept of between seven and nine per cent per year for both S&P 500 and Russell 2000 stocks. This rises to between ten and eleven per cent per year when the effect of the momentum factor is ignored.

This is the first time that the HS pattern has been shown to have economically significant predictive value for stock market returns. The evidence is consistent with the use of technical price patterns as a guide to trading decisions. The skewness in the conditional returns provides an explanation for the exaggerated claims made by the advocates of technical trading. Bulkowski, for example, reports that only seven per cent of the patterns he identified were not followed by a price decline, and that the average price decline conditional on a successful prediction was twenty three per cent. There is no reason to doubt that claims of this kind can be supported by judicious data mining. Our results suggest that technical analysis based on HS patterns may produce more modest, but still substantial risk-adjusted excess returns in the range of 7–9

percent per annum, leaving aside transaction costs. The finding that such price patterns have predictive power for future returns is one for which at present there is no satisfactory theoretical explanation.

References

- Barberis, Nicholas, Andrei Shleifer and Robert Vishny, 1998, "A Model of Investor Sentiment," *Journal of Financial Economics* 49, 307-343.
- Blume, Lawrence, David Easley and Maureen O'Hara, 1994, "Market Statistics and Technical Analysis: The Role of Volume," *Journal of Finance* 49, 153-181.
- Brown, David P., and Michael R. Gibbons, 1986, "A Simple Econometric Approach for Utility-Based Asset Pricing Models," *Journal of Finance* 40, 359-381.
- 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.
- 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.
- Daniel, Kent, David Hirshleifer and Siew Hong Teoh, 2001, "Investor Psychology in Capital Markets: Evidence and Policy Implications," Working Paper.
- 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.

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.

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

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

Rubinstein, Mark, 1976, "The Strong Case for the Generalized Logarithmic Utility Model as the Premier Model of Financial Markets," *Journal of Finance* 31, 551-571.

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

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	Variance	95 % Confidence Interval		Number of patterns
				Lower	Upper	
Without Bulkowski restrictions						
20	1	-0.00187	0.01211	-0.00338	-0.00035	20405
	1.5	-0.00182	0.01161	-0.00357	-0.00007	14558
	2	-0.00389	0.01260	-0.00608	-0.00170	10096
	2.5	-0.00515	0.01410	-0.00796	-0.00235	6860
40	1	-0.00698	0.02432	-0.00912	-0.00484	20405
	1.5	-0.00752	0.02350	-0.01001	-0.00502	14558
	2	-0.00927	0.02445	-0.01232	-0.00623	10096
	2.5	-0.01027	0.02619	-0.01410	-0.00644	6860
60	1	-0.01363	0.03702	-0.01627	-0.01101	20405
	1.5	-0.01473	0.03572	-0.01780	-0.01165	14558
	2	-0.01483	0.03656	-0.01856	-0.01110	10096
	2.5	-0.01460	0.03779	-0.01920	-0.01001	6860
With Bulkowski restrictions						
20	1	0.00022	0.01159	-0.00165	0.00208	12732
	1.5	-0.00045	0.01111	-0.00254	0.00164	9770
	2	-0.00397	0.01232	-0.00659	-0.00136	6894
	2.5	-0.00487	0.01471	-0.00834	-0.00141	4693
40	1	-0.00532	0.02434	-0.00803	-0.00261	12732
	1.5	-0.00576	0.02320	-0.00878	-0.00275	9770
	2	-0.00909	0.02483	-0.01281	-0.00538	6894
	2.5	-0.01004	0.02699	-0.01474	-0.00535	4693
60	1	-0.01187	0.03653	-0.01519	-0.00856	12732
	1.5	-0.01382	0.03595	-0.01758	-0.01007	9770
	2	-0.01457	0.03732	-0.01913	-0.01002	6894
	2.5	-0.01434	0.03804	-0.01992	-0.00875	4693

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	Variance	95 % Confidence Interval		Number of patterns
				Lower	Upper	
Without Bulkowski restrictions						
20	1	-0.00556	0.01873	-0.00679	-0.00432	47563
	1.5	-0.00566	0.01803	-0.00714	-0.00418	31624
	2	-0.00491	0.01732	-0.00671	-0.00312	20534
	2.5	-0.00477	0.01772	-0.00701	-0.00253	13569
40	1	-0.01043	0.03835	-0.01219	-0.00868	47563
	1.5	-0.01171	0.03735	-0.01384	-0.00958	31624
	2	-0.01014	0.03530	-0.01271	-0.00756	20534
	2.5	-0.01060	0.03572	-0.01378	-0.00744	13569
60	1	-0.01618	0.05992	-0.01838	-0.01398	47563
	1.5	-0.01846	0.05781	-0.02111	-0.01582	31624
	2	-0.01529	0.05508	-0.01850	-0.01207	20534
	2.5	-0.01544	0.05623	-0.01943	-0.01146	13569
With Bulkowski restrictions						
20	1	-0.00479	0.02036	-0.00633	-0.00326	32986
	1.5	-0.00473	0.01812	-0.00650	-0.00296	22223
	2	-0.00449	0.01779	-0.00667	-0.00231	14382
	2.5	-0.00454	0.01826	-0.00726	-0.00182	9483
40	1	-0.01001	0.04043	-0.01218	-0.00784	32986
	1.5	-0.01079	0.03762	-0.01334	-0.00825	22223
	2	-0.00958	0.03644	-0.01270	-0.00647	14382
	2.5	-0.00949	0.03659	-0.01334	-0.00563	9483
60	1	-0.01550	0.06353	-0.01822	-0.01278	32986
	1.5	-0.01707	0.05850	-0.02025	-0.01388	22223
	2	-0.01433	0.05665	-0.01822	-0.01044	14382
	2.5	-0.01338	0.05831	-0.01824	-0.00852	9483

Table III**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. 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
<i>20 days</i>					
Without Bulkowski Restrictions					
1	-0.00325 (-0.0051;-0.0014)	0.43841 (0.3771;0.4805)	0.06831 (0.0128;0.1191)	0.15082 (0.0559;0.2421)	-0.11761 (-0.1949;-0.0362)
2.5	-0.00515 (-0.0086;-0.0017)	0.48119 (0.4187;0.5295)	0.11605 (0.0466;0.1880)	0.21678 (0.0847;0.3468)	-0.1181 (-0.1930;-0.0337)
With Bulkowski Restrictions					
1	-0.00208 (-0.0041;-0.0010)	0.42697 (0.3522;0.4837)	0.05045 (-0.0053;0.1033)	0.12846 (0.0324;0.2205)	-0.12218 (-0.2017;-0.0389)
2.5	-0.00545 (-0.0089;-0.0021)	0.49243 (0.4225;0.5497)	0.14128 (0.0434;0.2397)	0.25114 (0.0960;0.4004)	-0.11205 (-0.1985;-0.0238)
<i>60 days</i>					
Without Bulkowski Restrictions					
1	-0.00687 (-0.0084;-0.0054)	0.70817 (0.6186;0.7880)	0.05958 (-0.0096;0.1268)	0.24777 (0.1321;0.3614)	-0.21417 (-0.2892;-0.1334)
2.5	-0.00724 (-0.0092;-0.0052)	0.7111 (0.5999;0.8123)	0.09378 (0.0175;0.1679)	0.23747 (0.1239;0.3510)	-0.17089 (-0.2497;-0.0895)
With Bulkowski Restrictions					
1	-0.00617 (-0.0078;-0.0045)	0.70654 (0.6103;0.7941)	0.05619 (-0.0122;0.1226)	0.23302 (0.1090;0.3551)	-0.2146 (-0.2913;-0.1324)
2.5	-0.00782 (-0.0098;-0.0058)	0.7325 (0.6141;0.8422)	0.12152 (0.0096;0.2313)	0.24927 (0.1048;0.3934)	-0.15985 (-0.2337;-0.0829)

Table IV

**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. The bandwidth parameter given in the first column is either 1 or 2.5. The autocorrelation- and heteroskedasticity-consistent errors are used to construct the confidence intervals.

	Intercept	Excess Market Return Factor	Size Factor	Book-to-Market Factor	Momentum Factor
<i>20 days</i>					
Without Bulkowski Restrictions					
1	-0.0042 (-0.0052;-0.0030)	0.40425 (0.3301;0.4577)	0.39591 (0.3233;0.4481)	0.11151 (0.0558;0.1609)	-0.08485 (-0.1311;-0.0382)
2.5	-0.00421 (-0.0061;-0.0022)	0.40099 (0.3377;0.4487)	0.41233 (0.3460;0.4646)	0.1744 (0.1013;0.2377)	-0.08396 (-0.1370;-0.0278)
With Bulkowski Restrictions					
1	-0.00362 (-0.0047;-0.0024)	0.40765 (0.3279;0.4667)	0.41604 (0.3394;0.4721)	0.10774 (0.0490;0.1604)	-0.10899 (-0.1571;-0.0590)
2.5	-0.00419 (-0.0063;-0.0019)	0.4109 (0.3335;0.4708)	0.40921 (0.3434;0.4563)	0.16852 (0.0780;0.2480)	-0.08878 (-0.1618;-0.0102)
<i>60 days</i>					
Without Bulkowski Restrictions					
1	-0.0066 (-0.0080;-0.0051)	0.64402 (0.5377;0.7348)	0.5647 (0.4644;0.6504)	0.24114 (0.1516;0.3244)	-0.17242 (-0.2540;-0.0832)
2.5	-0.00653 (-0.0085;-0.0046)	0.62365 (0.5489;0.6873)	0.6018 (0.5199;0.6728)	0.31138 (0.2046;0.4148)	-0.1545 (-0.2538;-0.0487)
With Bulkowski Restrictions					
1	-0.00644 (-0.0079;-0.0050)	0.65627 (0.5500;0.7479)	0.59368 (0.4926;0.6785)	0.23483 (0.1424;0.3220)	-0.18831 (-0.2649;-0.1056)
2.5	-0.00576 (-0.0079;-0.0037)	0.63018 (0.5517;0.6969)	0.6197 (0.5548;0.6728)	0.28328 (0.1841;0.3803)	-0.14637 (-0.2376;-0.0472)

Table V

**Regression Coefficients and Confidence Intervals in the One-Factor Regression:
S&P 500 and Russell 2000, 1990-1999**

The table reports the regression coefficients and their 95% confidence intervals in the one-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. The bandwidth parameter given in the first column is either 1 or 2.5. The ratio test reports the 95% confidence interval for $E[r^e/r_m]$ which is zero for the log CAPM. Autocorrelation- and heteroskedasticity-consistent standard errors are used to construct the confidence intervals.

	Intercept	Excess Market Return Factor	Ratio Test	Intercept	Excess Market Return Factor	Ratio Test
S&P 500			Russell 2000			
<i>20 days</i>			<i>20 days</i>			
Without Bulkowski Restrictions						
1	-0.0042 (-0.0059;-0.0024)	0.3901 (0.3216;0.4586)	(-0.0011;-0.0011)	-0.0061 (-0.0089;-0.0032)	0.403 (0.3125;0.4935)	(-0.0029;-0.0029)
2.5	-0.0061 (-0.0089;-0.0032)	0.4195 (0.3408;0.4981)	(-0.0028;-0.0028)	-0.006 (-0.0089;-0.0030)	0.3847 (0.2982;0.4712)	(-0.0029;-0.0029)
With Bulkowski Restrictions						
1	-0.0031 (-0.0048;-0.0014)	0.3847 (0.3047;0.4647)	(-0.0001;-0.0001)	-0.0058 (-0.0087;-0.0030)	0.4102 (0.3097;0.5107)	(-0.0026;-0.0026)
2.5	-0.0063 (-0.0094;-0.0033)	0.4236 (0.3279;0.5192)	(-0.0030;-0.0030)	-0.006 (-0.0088;-0.0032)	0.3952 (0.2924;0.4979)	(-0.0029;-0.0029)
<i>60 days</i>			<i>60 days</i>			
Without Bulkowski Restrictions						
1	-0.0088 (-0.0117;-0.0058)	0.6167 (0.4894;0.7440)	(-0.0034;-0.0034)	-0.0095 (-0.0139;-0.0050)	0.6251 (0.4633;0.7869)	(-0.0041;-0.0041)
2.5	-0.0088 (-0.0113;-0.0062)	0.6289 (0.4885;0.7693)	(-0.0034;-0.0033)	-0.0092 (-0.0140;-0.0044)	0.5872 (0.4435;0.7310)	(-0.0041;-0.0041)
With Bulkowski Restrictions						
1	-0.0081 (-0.0111;-0.0051)	0.6204 (0.4897;0.7510)	(-0.0028;-0.0027)	-0.0096 (-0.0142;-0.0050)	0.6434 (0.4786;0.8081)	(-0.0037;-0.0037)
2.5	-0.0093 (-0.0123;-0.0063)	0.6511 (0.4956;0.8065)	(-0.0037;-0.0037)	-0.0085 (-0.0130;-0.0039)	0.606 (0.4679;0.7442)	(-0.0033;-0.0032)




Figure 1 The head-and-shoulders pattern

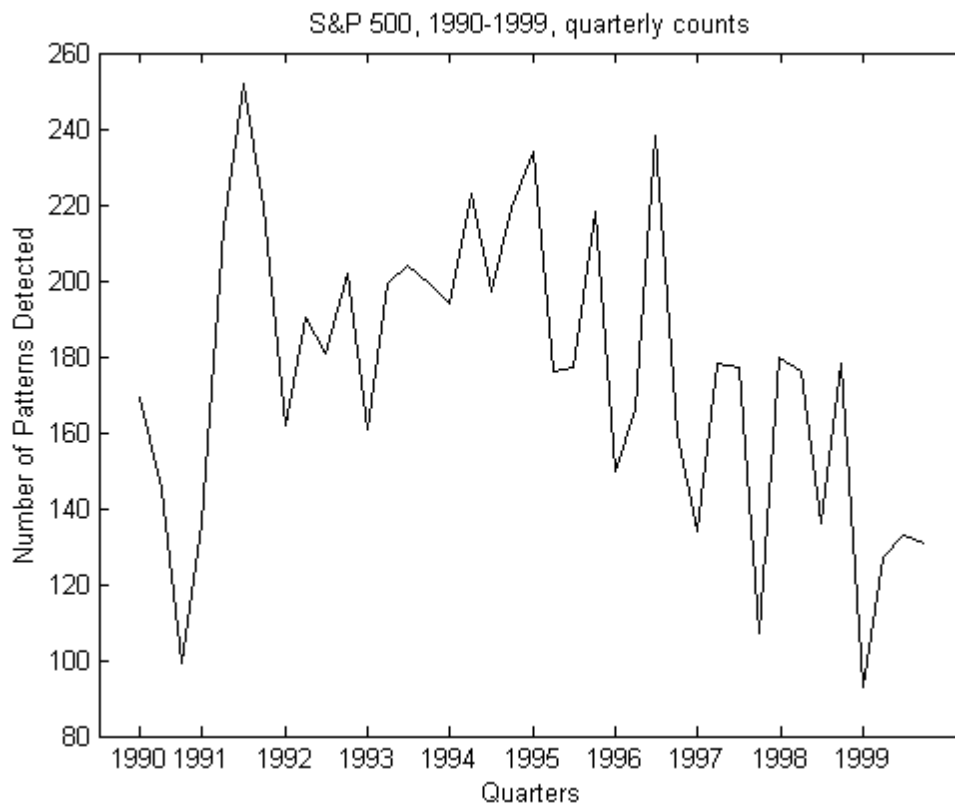


Figure 2 Quarterly occurrences of head-and-shoulders patterns for the S&P 500: 1990-1999

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.

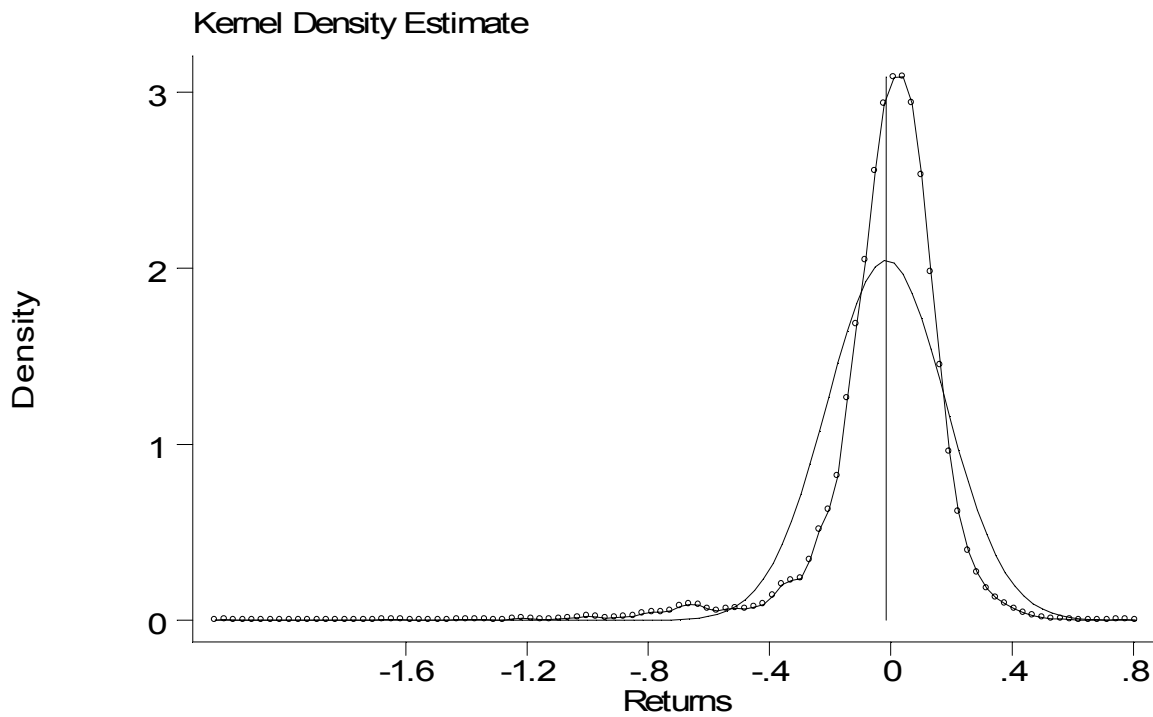



Figure 3. Conditional excess returns for S&P 500, 1990-1999, with the Bulkowski restrictions. The returns are for 60 trading days and a bandwidth multiple of 2.5. The vertical line in the kernel density represents the mean of the empirical distribution. The graphs are overlaid with a normal (solid line) that has the mean and variance of the empirical distribution.

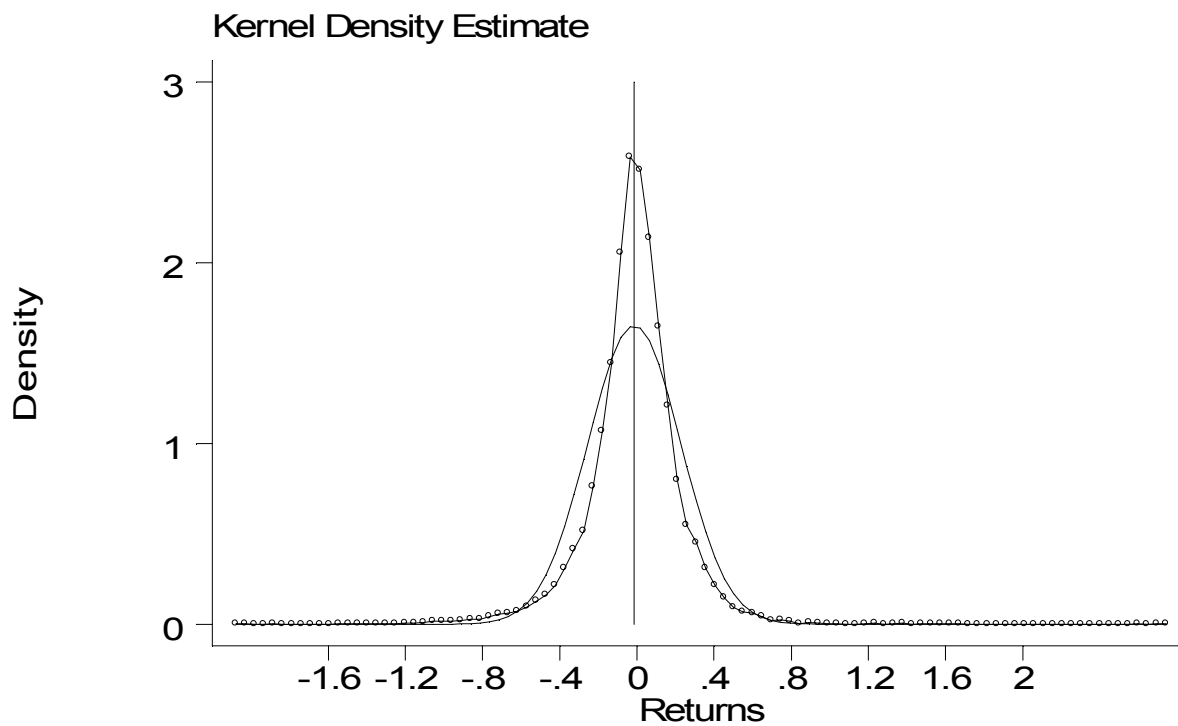



Figure 4. Conditional excess returns for the Russell 2000, 1990-1999, with the Bulkowski restrictions. The returns are for 60 trading days and a bandwidth multiple of 2.5. The vertical line in the kernel density represents the mean of the empirical distribution. The graphs are overlaid with a normal (solid line) that has the mean and variance of the empirical distribution.





Figure 5. Empirical distributions of the pattern characteristics restricted in (R4) and (R4a) and in (R5) and (R5a). These restrict the distance between the height of the left and right shoulder and the left and right trough. The left-hand vertical is at 1.5 percent, the value in (R4) and (R5), and the right-hand vertical is at 4 percent, the value in (R4a) and (R5a). The distributions are based on the patterns detected for the Russell 2000, 1990-1999, with a bandwidth multiple of 2.5.





Figure 6. Empirical distributions of the pattern characteristics restricted in (R6), (R7) and (R8). The top figure combines restrictions (R6) and (R7), which 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. The distributions are based on the patterns detected for the Russell 2000, 1990-1999, with a bandwidth multiple of 2.5.


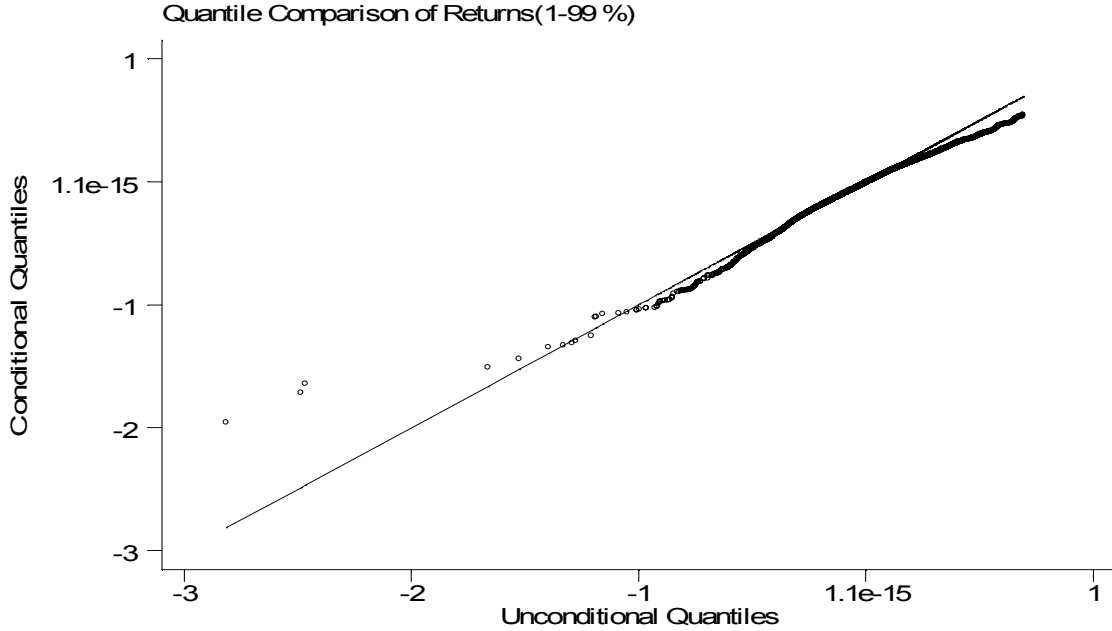


Figure 7. Empirical distributions of the pattern characteristics restricted in (R9). (R9) rules out extreme horizontal asymmetries in the HS patterns. The distributions are based on the patterns detected for the Russell 2000, 1990-1999, with a bandwidth multiple of 2.5.

Russell 2000



S&P 500




Figure 8. Quantile-Quantile Plot of Conditional and Unconditional Returns, 1990-1999. The conditional returns are calculated over a 60-day window after an HS pattern is observed. The HS patterns are selected using the Bulkowski restrictions and a bandwidth multiple of 2.5. The unconditional returns are calculated from a stratified sample of 60-day returns. The quantiles are plotted in one percent increments. The straight line represents the case where the two distributions are identical.

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:

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

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

$$r_{i,j}^e = r_{i,j} - \sum_{t=j}^{end} r_t, \quad (\text{A.2})$$

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:

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

The excess return is found as follows:

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

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:

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

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

The excess return is found as follows:

$$r_{i,j}^e = r_{i,j} - \sum_{t=1}^j r_t. \quad (\text{A.6})$$

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:

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

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 ($R_m - R_f$), the size factor (Small-Minus-Big, SMB), and the book-to-market (B/M) factor (High-Minus-Low, HML).

$R_m - R_f$, 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