

An Econometric Evaluation of Pattern Trading

Mieszko Dudziuk ^{α} and Jonas Olsson ^{β}

Abstract

This thesis implements the method of Lo, Mamaysky and Wang (2000) and analyze its econometric performance. We find no support for significant returns or excess returns generated by the method when using holding periods that range between 20 and 60 days. However, we find that when smoothing is high and the holding period is set to one day, the algorithm can pick returns from S&P 500 stocks during the years 1990-2008 averaging up to 30% per year. This estimate is verified to be statistically greater than the average return of the whole sample when looking at both returns and excess returns. The method is also verified to perform better than a portfolio that is constructed by randomly picking stocks.

Contents

1	Introduction	2
2	Previous research	3
2.1	The efficient market hypothesis	3
2.2	Technical analysis	4
2.3	A pattern detection algorithm	4
3	On technical analysis	5
3.1	Head-and-shoulders	6
3.2	Broadening tops and bottoms	6
3.3	Double tops and bottoms	7
4	Methodology	7
4.1	Smoothing estimators	7
4.2	Kernel regression	8
4.3	Bandwidth selection	9
4.4	The pattern recognition algorithm	11
4.5	Portfolio construction	13
4.6	Econometrical analysis	15
4.7	Over-performance	15
5	Data and descriptive findings	16
5.1	Pattern distribution	17
5.2	Benchmark portfolio	17
5.3	Long term holding period	19
5.4	Short term holding period	22
5.5	Case study of a combined portfolio	25
5.6	Accounting for trading costs	27
5.7	Is technical analysis profitable?	28
6	Conclusions	28
6.1	Suggestions for further research	29
A	Appendix	33

1 Introduction

Technical analysis is the analysis of trading data generated inside stock markets in order to predict future prices. Not uncommonly the analysis is presented in charts, hence the name “charting”. Technical analysis has traditionally not been given much attention in the academic world of financial economics and is rarely taught in business schools. Books about technical analysis are often written by practitioners, often presenting anecdotes rather than statistical evidence in support of their beliefs.

One of the reasons why technical analysis is a controversial topic is that it is hard to quantify, and hence to verify. Lo, Mamaysky and Wang (2000) present a method that overcomes this problem. The authors present an algorithm that detects ten patterns in price data in a systematic way, and the authors conclude that there is a distributional difference in returns detected by their algorithm and unconditional returns. This result is interesting but somewhat abstract, as no evaluation of the method’s economical performance is done. Savin, Weller and Zvingelis (2007) use a modified version of the algorithm to look at the performance of the “head-and-shoulders” pattern. They find that the method identifies trading opportunities that give a risk adjusted excess return of 5-7% per year, when using holding periods that range between 20 and 60 days and using S&P 500 stocks. On the other hand, the raw returns in this study are not confirmed to surpass the returns of the used sample.

We believe that a study of the economical performance of the algorithm is interesting as it is a test of claims by technical analysts. The econometrical evaluation of the method’s returns gives results that are common to economists and comparable with other studies. Our sample includes several years of data after the publishing of the method, which also makes it an out of box test.

We employ the method of Lo, Mamaysky and Wang (2000) and analyze its economical performance by methods similar to Savin, Weller and Zvingelis (2007). We find no support for significant returns or excess returns generated by the method when using holding periods that range between 20 and 60 days. However, when we use a one day holding period we are able to pick returns from S&P 500 stocks during the years 1990-2008 averaging up to 30.7% per year. This estimate is verified to be statistically greater than the average returns of the whole sample when looking at both returns and excess returns. The method is also verified to perform better than a portfolio that is randomly constructed with a p-value equal to 0.5%. The method relies on daily trading and the break even trading cost is 10 basis points for a two way transaction, which makes it hard to implement in practice.

2 Previous research

Technical analysis hinges on the assumption that returns are predictable. Therefore, we find it interesting to discuss the efficient market hypothesis in this context. Further, we present research done on technical analysis, and end this section looking at work that is closely related to our own.

2.1 The efficient market hypothesis

The efficient market hypothesis is associated with the idea of a “random walk”, suggesting that future stock prices cannot be projected using historical data. The idea behind the random walk hypothesis is that if the information flow is unrestricted and information is directly reflected in stock prices, then tomorrow’s price change will reflect only tomorrow’s news and will be independent of the price changes today. However, since news by definition are unpredictable, the resulting price changes must be unpredictable and random. An efficient market implies that when information arises, this news spreads very quickly and is being incorporated into the prices of securities without any delay.

Early academic work testing the efficient market hypothesis concluded that stock market prices are not predictable, see Fama (1965, 1970). However, more recent findings suggest that the efficient market hypothesis do not hold, giving indirect support for the possibility of technical analysis. Lo and MacKinlay (1988, 1999) show that past prices may be used to predict future returns, and strongly reject the random walk model. Treynor and Ferguson (1985) show that past prices combined with other valuable information can predict unusual returns.

Lo and MacKinlay (1999) find that short-run serial correlations are non-zero. The existence of too many successive moves in the same direction enables the authors to reject the hypothesis that stock prices behave as true random walks. The authors therefore conclude that there does seem to be some momentum in short-run stock prices.

In the field of behavioral finance, short-run momentum is consistent with psychological feedback mechanisms. That is, when investors observe a rising stock price they are drawn into the market in a kind of “bandwagon effect”. Shiller (2000) describes this phenomenon and discusses the bullish U.S. stock market during the late 1990s (dot-com boom) as the result of psychological contagion leading to irrational exuberance. Another explanation for patterns of short-run momentum, put forward by behaviorists, is the tendency for investors to underreact to new information.

Continuing on the discussion of the efficiency of the market, there is fur-

ther evidence that stock returns are not fully explained by common risk factors. Fama and French (1993) examines two easily measured variables, size and book-to-market equity, and finds that these two combine to capture the cross-sectional variation in average stock returns associated with market beta, size, leverage, book-to-market equity, and earnings-price ratios. Jegadeesh and Titman (1993) examines strategies which involves buying stocks that have performed well in the past and sell stocks that have performed poorly in the past. These strategies generate significant positive returns over three to twelve months holding periods. Further, the authors find that the profitability of these strategies are not due to their systematic risk or to delayed stock price reactions to common factors. A similar pattern of returns around the earnings announcements of past winners and losers is also documented.

The findings discussed above give room for technical analysis. Even if the market can be treated as an information processor, what can be disputed is the degree of its efficiency.

2.2 Technical analysis

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, Tittermann and White (1999).

There is convincing evidence that stock prices display short-term momentum over six months to a year and longer term mean reversion [DeBondt 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 trading rules that can be 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)].

Despite the fact that practitioners are trading on price patterns, there is little academic attention on this topic. Chang and Osler (1999) examine the profitability of the “head-and-shoulders” pattern in the foreign exchange market and find evidence of excess returns for some currencies.

2.3 A pattern detection algorithm

Lo, Mamaysky and Wang (2000) propose a new approach to evaluate the efficiency of technical analysis. The idea of their method is to use a pattern

recognition algorithm that finds patterns similar to those that technical analysts find compelling. Using smoothing techniques such as nonparametric kernel regression, their approach incorporates the essence of technical analysis - to identify regularities in the time series of prices by extracting nonlinear patterns from noisy data. The data used in the analysis are daily returns of several hundred stocks between 1962 and 1996. As comparison they perform the analysis on data generated by geometric Brownian motion. The pattern density is different for the simulated data compared to the actual stock data. The authors' main conclusion is that the statistical properties of the returns differ when conditioned on an identified pattern compared to unconditional returns. The analysis of the returns is performed with two statistical methods: the *Goodness-of-Fit* test and *Kolmogorov-Smirnov* test. One interesting feature of these two tests is that they test the difference between the conditional and unconditional return distribution, and hence test for a difference in all moments of the returns, whereas regression based models focus on the first moment (the mean) and the second moment (the variance). This means, given that there is a distributional difference between conditional and unconditional returns, the conditional returns do not automatically have to surpass the unconditional returns in an economical sense. This is one of the reasons why we want to perform the economical analysis of the method.

Mossberg and Undin (2007) use the same methodology on a dataset consisting of daily returns for 65 Swedish stocks during 1982-2005. They find similar results.

Savin, Weller and Zvingelis (2007) use the pattern detection algorithm to search for the head-and-shoulders in S&P 500 and Russel 2000 stocks between 1990 and 1999. They modify the algorithm in order to make it more comparable to descriptions made by technical analysts of how the head-and-shoulders pattern behaves. One of the differences is that they use a holding period equal to 20, 40 and 60 trading days, instead of the one day used in the original paper. A second difference is that they set the maximal pattern length to 60 days instead of 35 days. When looking at the S&P 500 stocks, the authors find that the strategy is not profitable on a stand-alone basis. However they find that returns following the patterns give a risk adjusted excess return of 5-7% per year. The analysis is performed using the three and four factor models.

3 On technical analysis

Technical analysis is “the science of recording, usually in graphic form, the actual history of trading (price changes, volume of transactions, etc.) in a

certain stock and then deducting that pictured history the probable future trend”, (Edwards and Magee, 1997). The patterns detected by the Lo, Mamaysky and Wang (2000) algorithm are difficult to quantify. The authors look at five pairs of patterns. Since three of them are distinctive buy or sell patterns, we will focus at them. The patterns are presented below, and graphical examples of patterns detected by the algorithm can be found in the appendix.

3.1 Head-and-shoulders

The head-and-shoulders pattern is believed to be one of the most reliable trend-reversal patterns and is one of the most famous patterns in technical analysis. The head-and-shoulders formation consists of three distinct parts: a left shoulder, a head and a right shoulder. The left shoulder consists of a peak followed by a higher peak (the head) and then a lower peak (the right shoulder) with a break below the neckline. The neckline is drawn through the lowest points of the two intervening troughs and may slope upward or downward. A downward sloping neckline is more reliable as a signal. The confirmation of the break out (i.e. when the stock price crosses the neckline) is a strong signal to sell the stock. The inverse head-and-shoulders is a vertically mirrored image of the head-and-shoulders. With inverted head-and-shoulders the neckline is drawn through the highest points of the two intervening peaks. A downward sloping neckline signals continuing weakness and is less reliable as a reversal signal. The break out in the inverted head-and-shoulders pattern is a strong signal to buy the stock.

3.2 Broadening tops and bottoms

The broadening formation is labeled a broadening bottom when the trend preceding the broadening formation is down and a broadening top when the trend preceding the pattern formation is up. Broadening tops are recognized by initially small fluctuations in price and then large drops/increases, forming a broadening top.

The broadening bottom is different from broadening tops because price enters the pattern from the top. In both patterns, price touches each trend line at least two times and swings in a progressively wider range. Broadening tops and bottoms occurs when price volatility increases since market participants becomes uncertain what the correct valuation (price) should be. The appearance of these two patterns indicates that the stock is entering a dangerous stage. Therefore, investors should sell the stock and consequently the broadenings are sell signals.

3.3 Double tops and bottoms

Double tops are generated when a top or bottom price fails to exceed the price of a prior significant top or bottom. The distance between these two peaks should be a couple of weeks. The double top is a sell signal and the double bottom is a buy signal. When trying to understand what is happening in a double top or bottom, it is important to understand the concept of support and resistance. Support is a price area where buying interest is such that there are more willing buyers than sellers, and this buying will drive prices back up. Resistance is a price area where selling interest is strong, and sellers will overwhelm buyers and drive the price down.

4 Methodology

We will employ the algorithm presented in Lo, Mamaysky and Wang (2000). Their method consists of several steps. The authors build upon ideas presented in Neftci (1991) and combine them with non-parametrical regression methods presented in Härdel (1990).

When we have the set of patterns found by the Lo, Mamaysky and Wang (2000) algorithm, we will construct portfolios and analyze them with conventional regression methods, as performed by Savin, Weller and Zvingelis (2007).

The underlying assumption of this method is that a stock price series P_t evolve according to

$$P_t = m(X_t) + \epsilon_t, \quad t = 1, \dots, T$$

where $m(X_t)$ is an arbitrary fixed but unknown nonlinear function of a state variable X_t and ϵ_t is white noise. In this context X_t can be interpreted as a time, $X_t = t$, but we keep it in the formulas, as it is a common notation in literature on kernel regression. Since the function m is unknown, the idea is to estimate it non-parametrically, that is, without making any assumptions of how this function looks like.

4.1 Smoothing estimators

Smoothing is a common procedure to estimate a non-linear relation (such as the equation above), which aims to reduce the observational errors by averaging data in sophisticated ways. Suppose we are able to sample the stock price at the same point in time n times, that is for a fixed $X_t = x_0$.

Then we would estimate m as

$$\hat{m}(x_0) = \frac{1}{n} \sum_{i=1}^n p_i = \frac{1}{n} \sum_{i=1}^n (m(x_0) + \epsilon_t^i) = m(x_0) + \frac{1}{n} \sum_{i=1}^n \epsilon_t^i.$$

As n increases the last term vanishes by the Law of Large Numbers and we can get an arbitrarily good estimate of m .

Since stock prices P_t are time series, there is no access to repeated observations at the same point in time. Instead, we assume that the function m is sufficiently smooth such that in a small neighborhood around x_0 , $m(X_t)$ will be constant and may be estimated by the average of P_t 's in this neighborhood. This argues for a weighted average of P_t 's with weights that decline as the distance from x_0 increase.

Quantitatively, a smoothing estimator of m can be expressed as

$$\hat{m}(x) = \sum_t \omega_t(x) \cdot P_t, \quad \sum_t \omega_t(x) = 1$$

where the weights $\omega_t(x)$ are chosen such that they are large for P_t 's with X_t close to x , and decrease as the distance from x increase. Here the weights have to be chosen carefully. If they are large even for P_t 's far away there is a risk for smoothing out true nonlinearities of m . On the other hand, if they decrease too fast too much noise may be captured.

4.2 Kernel regression

For a kernel regression estimator, the weight function $\omega_t(x)$ is constructed from a kernel $K(x)$ which displays the same properties as a probability density function:

$$K(x) \geq 0, \quad \int K(x) \, dx = 1.$$

These two properties remain after we rescale the kernel with a parameter h :

$$K_h(x) = \frac{1}{h} K(x/h).$$

Next we express the weight function in terms of the kernel:

$$\omega_{t,h}(x) = \frac{K_h(x - X_t)}{\sum_t K_h(x - X_t)}.$$

For small h the averaging will be done in a rather narrow neighborhood of x , whereas large h will use a large neighborhood. The parameter h is

known as the bandwidth, and controls the degree of smoothing. By using the weight function in the smoothing estimator we get the Nadaraya-Watson kernel estimator $\hat{m}$ of m :

$$\hat{m}_h(x) = \frac{\sum_t K_h(x - X_t) \cdot P_t}{\sum_t K_h(x - X_t)}.$$

If the kernel K fulfills certain regularity conditions, $\hat{m}_h$ converges into m asymptotically. The convergence property holds for a number of different kernels. We will be using the Gaussian kernel

$$K_h(x) = \frac{1}{h\sqrt{2\pi}} e^{-x^2/2h^2}$$

which is a popular choice of kernel and for which the regularity conditions hold. We substitute the Gaussian kernel into the Nadaraya-Watson formula and get the closed form formula of our estimator:

$$\hat{m}_h(x) = \frac{\sum_t e^{-(x-X_t)^2/2h^2} \cdot P_t}{\sum_t e^{-(x-X_t)^2/2h^2}}.$$

Please refer to Härdle (1990) for more information on kernel regression and its statistical properties.

4.3 Bandwidth selection

As we discussed earlier, the selection of bandwidth h is important as it determines how much smoothing of the price data we want to perform. We will use a method called cross-validation, in which h is chosen to minimize the cross-validation function:

$$CV(h) = \sum_t (P_t - \hat{m}_{h,t})^2$$

where

$$\hat{m}_{h,t} = \frac{\sum_{i \neq t} e^{-(0-X_i)^2/2h^2} \cdot P_i}{\sum_{i \neq t} e^{-(0-X_i)^2/2h^2}}.$$

For a given bandwidth parameter h , the cross-validation function is a measure of the ability of the kernel regression estimator to fit each observation P_t when that observation is left out. By selecting the bandwidth that minimizes the cross-validation function,

$$h^* = \operatorname{argmin} CV(h),$$




Figure 1 Comparison of $CV(h)$ vs. h with a distinctive minimum (left) and a typical look without a distinctive minimum (right).

we obtain a kernel estimator that satisfies certain optimality properties. One example is minimum asymptotic mean-squared error.

When implementing the selection of the optimal bandwidth h^* we run into some ambiguities. Lo, Mamaysky and Wang (2000) conclude h^* is too large and use the bandwidth $0.3 \cdot h^*$ in their analysis. The authors used trial and error and approached professional technical analysts when they choose their bandwidth multiplier. Savin, Weller and Zvingelis (2007), on the other, hand multiply h^* with the factors 1, 1.5, 2 and 2.5. They do not motivate this choice.

None of the two sources describes in what way they arrive upon numerical values of h^* , or present typical values found in their datasets. We used a numerical solver, but unfortunately we found out that the cross-validation function often did not have a distinctive minimum. Instead $CV(h)$ tended to be flat for values of h between 0 and 0.4, and then rise. Two examples of how $CV(h)$ looks like are presented in Figure 1. It should also be noted that when $h = 0.2$, then $\omega_{t,h}(X_t) = 0.99999$. Qualitatively, this means that virtually no smoothing is done.

Instead, we let the numerical solver find a minimum for $CV(h)$ for h between 0.416 and 10. In most cases the numerical solver will yield the value $h^* = 0.416$. The upper limit 10 is set so high that it will not affect the results. The numerical value of 0.416 was chosen simply because $\omega_{t,h=0.416}(X_t) = 0.90$, which means that we at most will use 90% of P_t 's value when estimating $m(X_t)$. Another reason for choosing the lower boundary was to enable the analysis of increased smoothing. For example, suppose we do not use the lower boundary condition and at one occasion get $h^* = 0.10$. Then, even with a bandwidth multiplier set to 2.5, $\omega_{t,h=0.1 \cdot 2.5}(X_t)$ is still very close to 1, and there is very little smoothing performed. On the other

hand $\omega_{t,h=0.416 \cdot 2.5}(X_t) = 0.44$, which gives a smooth curve.

When looking at patterns detected by our implementation of the algorithm we conclude that 0.3 is a too small bandwidth multiplier. Instead, we will use the multipliers $\varphi = 1, 1.5, 2$ and 2.5 , which enable us to look at the effects of increased smoothing. Please note that we decided how to select bandwidth before the econometrical analysis was performed and without any revising. We are able to study the effect of further increased smoothing by selecting patterns with $h^* > 0.416$. We call this applying the *filter*.

4.4 The pattern recognition algorithm

Now that we have a mathematical representation m of P_t , we are ready to quantify the technical patterns and describe how they are searched for in the price series.

Definitions of technical patterns

Lo, Mamaysky and Wang (2000) present five pairs of popular technical patterns. As we described earlier, we will use three of those pairs as they are distinctive buy or sell patterns.

Consider the systematic component m of a price history P_t and suppose we have identified n local extrema in it. Denote them by $E_1, E_2, \dots, E_n$ and the dates they occur by $t_1^*, t_2^*, \dots, t_n^*$. Note that if E_i is a maximum then E_{i+1} is a minimum, as extrema must alternate between maxima and minima for a continuous function. Then we have the following definitions:

Definition 1: Head-and-Shoulders

Head-and-shoulders (HS) and inverted-head-and-shoulders (IHS) are characterized by a sequence of five consecutive local extrema $E_1, \dots, E_5$ such that

$$\begin{aligned} \text{HS} &= \begin{cases} E_1 \text{ is a maximum} \\ E_3 > E_1, E_3 > E_5 \\ E_1 \text{ and } E_5 \text{ are within 1.5 percent of their average} \\ E_2 \text{ and } E_4 \text{ are within 1.5 percent of their average} \end{cases} \\ \text{IHS} &= \begin{cases} E_1 \text{ is a minimum} \\ E_3 < E_1, E_3 < E_5 \\ E_1 \text{ and } E_5 \text{ are within 1.5 percent of their average} \\ E_2 \text{ and } E_4 \text{ are within 1.5 percent of their average} \end{cases} \end{aligned}$$

Definition 2: Broadening

Broadening tops (BTOP) and bottoms (BBOT) are characterized by a se-

quence of five consecutive local extrema $E_1, \dots, E_5$ such that

$$\text{BTOP} = \begin{cases} E_1 \text{ is a maximum} \\ E_1 < E_3 < E_5 \\ E_2 > E_4 \end{cases}$$

$$\text{BBOT} = \begin{cases} E_1 \text{ is a minimum} \\ E_1 > E_3 > E_5 \\ E_2 < E_4 \end{cases}$$

Defintion 3: Double top and Bottom

Double tops (DTOP) and bottoms (DBOT) are characterized by a sequence of extrema $E_1, E_2, \dots, E_n$ such that

$$\text{DTOP} = \begin{cases} E_1 \text{ is a maximum} \\ E_1 \text{ and } E_n \text{ are within 1.5 percent of their average} \\ E_n > E_i, \ i = 2, \dots, n-1 \\ t_n^* - t_1^* > 22 \end{cases}$$

$$\text{DBOT} = \begin{cases} E_1 \text{ is a minimum} \\ E_1 \text{ and } E_n \text{ are within 1.5 percent of their average} \\ E_n < E_i, \ i = 2, \dots, n-1 \\ t_n^* - t_1^* > 22 \end{cases}$$

The identification algorithm

The algorithm takes a sample of prices $P_1, \dots, P_T$ and look for patterns in subsamples or windows. Each window begins at t and ends at $t + L - 1$, where t varies from 1 to $T - L + 1$. Lo, Mamaysky and Wang (2000) use a window length of $L = 38$ trading days. Savin, Weller and Zvingelis (2007) use $L = 63$, arguing that the head and shoulder pattern typically spans over this period of time. We use $L = 38$, as we are looking at six different patterns and have no presumption of typical pattern lengths. It should also be noted that the window length only sets an upper limit on how wide detected patterns can be. The length of a particular pattern may as well be shorter.

Further we restrict each pattern to end at day 35 in each window. This is done to avoid ambiguous cases where more than one pattern occurs in the window. It also speeds up the algorithm. Suppose that we drop this restriction and consider, for example, a pattern that is 15 days wide. When the window rolls by this pattern it can be detected in 20 different windows. This is not necessary as we only need to detect the pattern once. Therefore, the restriction should not affect the pattern detection algorithm significantly. The

reason to choose day number 35 instead of, for example, day number 38 (the last observation in the window) is that each pattern ends at an extremum, but the last observation in the window cannot be confirmed as local minimum or maximum as the following observation is unknown. Lo, Mamaysky and Wang (2000) use this distance of 3 days, but they do not explicitly state the restriction that the pattern should end at day 35. However, they present graphs where all patterns end at day 35.

For each window the optimal bandwidth is selected and multiplied by the bandwidth multiplier φ (in order to get h) and the kernel regression is fitted.

The extrema in $\hat{m}$ are defined as follows. The point $\hat{m}(X_i)$ is a local maximum if $\hat{m}(X_{i-1}) < \hat{m}(X_i)$ and $\hat{m}(X_i) \geq \hat{m}(X_{i+1})$. The inequalities are reversed if $\hat{m}(X_i)$ is a local minimum. If $\hat{m}(X_i)$ is identified as a local extremum, then the relevant extremum is chosen from P_{i-1} , P_i and P_{i+1} . When the set of relevant extrema $E_1, E_2, \dots, E_n$ and their corresponding dates are determined in the window, the patterns defined in the previous section are searched for.

Lo, Mamaysky and Wang (2000) argue that when using the kernel regression approach to identify patterns, local extrema that are “too local” are ignored. For example, consider the simpler case where extrema are searched for in the raw price data. If we identify P_t as a local maximum if $P_t > P_{t-1}$ and $P_t > P_{t+1}$ and vice versa for a local minimum, we would identify a large number of extrema and get patterns that are not visually consistent with the kind of patterns that technical analysts find compelling. This example also further motivates our own introduction of a lower boundary of the optimal bandwidth, as $\hat{m}(X_t) \approx P_t$ for small h .

4.5 Portfolio construction

After the pattern recognition algorithm has found all patterns in the dataset we are ready to aggregate the results into portfolios. Let n be the holding period for each stock, after a pattern has been identified. We will look at both long term and short term holding periods. The long term holding period will refer to a holding period equal to $n = 20, 40$ and 60 days, whereas the short term holding period refers to $n = 1$. Lo, Mamaysky and Wang (2000) use $n = 1$. Savin, Weller and Zvingelis (2007) use the long term holding period, arguing that it is a better approximation to how technical analysts trade.

Our starting point is the one day return of stock i in the dataset, which typically is calculated by the following formula:

$$R_{i,t} = P_{i,t}/P_{i,t-1} - 1.$$

Please note that we do not calculate those ourselves, as stock splits, dividends and other events must be accounted for. Instead we download $R_{i,t}$ along with the stock prices $P_{i,t}$. Next step is to introduce an indicator $I_{i,t}$ that indicates if stock i is included in the portfolio at time t .

$$I_{i,t} = \begin{cases} S_i(t_i^*), & 0 \leq t - t_i^* < n \\ 0, & \text{otherwise} \end{cases}$$

where $t_i^* \leq t$ is the time of the most recent observation following a window containing a pattern for stock i . The function $S_i(t^*) = 1$ if the pattern is a buy signal, and -1 for a sell signal. The number of stocks included in the portfolio at day t is denoted by N_t :

$$N_t = \sum_i |I_{i,t}|.$$

The portfolio return RP is calculated by weighting the returns equally each day. If no stocks are included in the portfolio, it earns the risk free rate.

$$RP_t = \frac{1}{N_t} \sum_i I_{i,t} \cdot R_{i,t}, \quad N_t \neq 0,$$

$$RP_t = RF_t, \quad N_t = 0,$$

where RF is the risk free one day return. When calculating the excess returns RPX we average the excess returns of the included stocks:

$$RPX_t = \frac{1}{N_t} \sum_i I_{i,t} (R_{i,t} - RF_t) = RP_t - RF_t \frac{1}{N_t} \sum_i I_{i,t}, \quad N_t \neq 0,$$

$$RPX_t = 0, \quad N_t = 0.$$

The continuously compounded one day returns and excess returns are calculated simply using

$$rp_t = \ln(1 + RP_t),$$

$$rpx_t = \ln(1 + RPX_t).$$

Finally we calculate the continuously compounded return and excess return of month m by summing the daily returns in each month:

$$rp_m = \sum_{t=t_m^{\min}}^{t_m^{\max}} rp_t, \quad rpx_m = \sum_{t=t_m^{\min}}^{t_m^{\max}} rpx_t,$$

where $t_m^{\min}$ and $t_m^{\max}$ is the time of the first and last observation of month m , respectively.

4.6 Econometrical analysis

When the monthly portfolio excess returns are calculated, the risk adjusted excess returns can be calculated. We will use the risk factors presented in Fama and French (1993).¹ We will also use the fourth momentum factor developed by Carhart (1997), as we want to check if profits (if any) from technical analysis can be explained by loadings on the momentum factor. When referring to a three factor regression we mean

$$rpx_t = \alpha^{rpx} + \beta^m \cdot m_t + \beta^{smb} \cdot smb_t + \beta^{hml} \cdot hml_t,$$

whereas the four factor model refers to

$$rpx_t = \alpha^{rpx} + \beta^m \cdot m_t + \beta^{smb} \cdot smb_t + \beta^{hml} \cdot hml_t + \beta^{umd} \cdot umd_t.$$

The factors are denoted with lower case as they are continuously compounded.

There are several reasons why we choose to perform the regressions on monthly returns, rather than on daily. Firstly we get values that are directly comparable to Savin, Weller and Zvingelis (2007) as the authors also perform analysis on monthly returns. Secondly, regression estimates of daily stock returns may be biased by the presence of autocorrelation in a larger extent than that of monthly returns. The regressions are performed with robust standard errors, as the stocks in the portfolio are changing and most probably introducing heteroscedasticity.

4.7 Over-performance

As a last step in our analysis we will check for something that we define as *over-performance*. In order to be over-performing a portfolio must meet the following criteria:

1. The mean of the returns rp must have positive lower boundary in its 95% confidence interval.

In other words, the raw returns must be significantly larger than zero. The next criterion looks at the risk adjusted excess return α^{rpx} . When regressing the four factor model we want a significantly positive α :

2. $\alpha^{rpx} \neq 0$ at the 5% significance level.
3. $\alpha^{rpx} > 0$

¹Data on factor portfolio excess returns was obtained from Kenneth French's web site: <http://mba.tuck.dartmouth.edu/pages/faculty/ken.french/>

As a final step we want the risk adjusted excess return of the portfolio to be larger than that of all the underlying stocks. Therefore, we regress the following two equations simultaneously

$$\begin{aligned} rpx_t &= \alpha^{rpx} + \beta_1^m \cdot m_t + \beta_1^{smb} \cdot smb_t + \beta_1^{hml} \cdot hml_t + \beta_1^{umd} \cdot umd_t, \\ snpx_t &= \alpha^{snpx} + \beta_2^m \cdot m_t + \beta_2^{smb} \cdot smb_t + \beta_2^{hml} \cdot hml_t + \beta_2^{umd} \cdot umd_t. \end{aligned}$$

and check that the following two criteria are fulfilled:

4. $\alpha^{rpx} \neq \alpha^{snpx}$ at the 5% significance level.
5. $\alpha^{rpx} > \alpha^{snpx}$.

In the equation above $snpx$ is the equally weighted excess return of all stocks in the dataset. The reason to use this two last criteria is to check whether a portfolio is better than the stocks it is created from. For example, suppose we are sampling returns from a set of stocks that all fulfill criteria [1-3]. Then, if a randomly constructed portfolio also meets criteria [1-3], this does not have to be the result of a stock picking ability, but rather a property of the underlying stocks. Therefore criteria [4-5] also should be checked. Since our statistics program only can perform double-sided tests, we want to stress that when using double sided tests at the 5% significance level, and only looking at one side of the test, we are actually using a 2.5% significance level. For example, combining criteria [2] and [3] is equivalent to the test: " $\alpha^{rpx} > 0$ at the 2.5% significance level."

5 Data and descriptive findings

Our dataset consists of the stocks in the S&P 500 index over the period 1990-2008. The daily stock price and return data is from the Center for Research in Security Prices (CRSP) downloaded from the Wharton Research Data Services (WRDS) website. To account for splits, mergers and a like we use the cumulative price and dividend adjustment factor created by CRSP to be able to compare stock prices across time.

The dataset spans over 4,791 days and contains 978 stocks. In total we have 2.4 million price observations. In case of a missing closing price CRSP reports the average bid/ask price (35 occasions in our sample). In case of a missing closing price and average bid/ask we replace the price with the previous observation (10 occasions).

Prices are divided by the cumulative price and dividend adjustment factor so that the pattern detection algorithm is unaffected by stock splits. We also

Table 1 The distribution of each pattern in the dataset.

Pattern	No filter				Filter			
	$\varphi = 1$	1.5	2	2.5	1	1.5	2	2.5
HS	45,845	29,427	19,780	14,307	11,661	6,077	3,524	1,996
IHS	45,478	28,954	19,378	13,803	11,426	6,012	3,401	1,962
BTOP	11,445	8,443	6,278	4,830	3,482	2,085	1,201	654
BBOT	10,476	7,650	5,707	4,373	3,041	1,814	1,093	610
DTOP	34,105	26,815	21,459	17,727	11,668	8,375	6,176	4,627
DBOT	30,984	24,460	19,579	16,264	10,787	7,668	5,632	4,329
Total	244,254	171,204	122,340	93,134	70,139	41,145	26,052	16,932

download the one day holding period returns that are adjusted for dividends and other distributional events.

A final comment about the sample period is that the econometrical analysis of all portfolios, benchmarks and simulations is done in the period March 1990 - December 2008. In other words, we have excluded the first two months in the dataset. The reason behind this is that the algorithm can detect the first trading opportunity at day 39 of the dataset. Earlier patterns cannot be identified as we use a window of 38 days. As the econometrical analysis is performed on monthly data, we have 226 observations for each constructed portfolio.

5.1 Pattern distribution

The first step is to look at the distribution of patterns in the dataset, which is presented in Table 1. We note that the number of patterns is reduced when the bandwidth multiplier increases. Applying the filter also decreases the number of patterns. The patterns found in each column may not be unique. The same pattern can be discovered with different bandwidth multipliers. All patterns found when the filter is applied are found without the filter as well.

5.2 Benchmark portfolio

Before we start to look at the performance of portfolios created on the basis of the patterns found by the algorithm, we create a benchmark portfolio. This portfolio is created by including all the stocks in our dataset. We will refer to the time series with returns of this portfolio as *snp* and to its excess returns as *snp_x*.

As our dataset consists of the stocks in the S&P 500 index, *snp* should be

Table 2 Mean and confidence interval for the benchmark returns.

	Mean	low95	high95	Annual(%)
<i>snp</i>	0.0077	0.0015	0.0140	9.74
<i>snp_x</i>	0.0045	-0.0018	0.0107	5.52
<i>negsnp</i>	-0.0106	-0.0165	-0.0046	-11.92
<i>negsnp_x</i>	-0.0073	-0.0133	-0.0014	-8.40
<i>rf</i>	0.0033	0.0031	0.0035	4.00

identical to the equally weighted with dividend return of the S&P 500 index. We have compared the returns and they are very similar, but not identical. The largest error is 4.7 basis points. Out of the 4,791 days the difference is larger than 1 bps at 21 occasions. The signs of the error are random, and are distributed as 2,406/2,385. We believe that these errors stem from numerical calculations and that they are nonsystematic.

Some of the patterns are sell patterns and we are allowed to take short positions. Therefore we construct the *negsnp* and *negsnp_x* time series:

$$\begin{aligned} negsnp_t &= \ln(1 - SNP_t), \\ negsnp_x_t &= \ln(1 + RF_t - SNP_t). \end{aligned}$$

The analysis of the aggregated returns helps us to understand how a portfolio that is opening and closing a short position in all stocks every day performs. We also look at the time series of the continuously compounded risk free rate *rf*, which was downloaded together with the four factors. The confidence intervals show that all stocks in the dataset, when equally weighted, increase 9.74% per year during the sample period. Holding the risk free rate would give 4.00% per year, whereas a strategy short selling all the stocks would give a negative return, averaging -11.92% per year. All the time series, except *snp_x*, have averages that are significantly different from zero.

We also regress the excess return time series *snp_x* and *negsnp_x* on the three and four factors. Table 3 shows that the dataset itself has positive alphas in three out of four cases (with *snp_x* and three factors being the exception).

Performing this analysis may seem superfluous as our dataset is known to most financial economists. Our point here is to show how the dataset performs when using the same analysis that is later performed on our portfolios, which are subsamples of the dataset. If we choose returns from our dataset randomly to construct a portfolio, its returns should converge into the returns of *snp* or *negsnp* depending on if we are going long or short.

Table 3 Regression of the benchmark returns. P-values are indicated in parenthesis below their corresponding point estimate.

Variable	Factors	mx	smb	hml	umd	alpha
<i>snp</i>	3	1.12	0.01	0.41		-0.0001
		(0)	(0.817)	(0)		(0.917)
	4	1.06	0.03	0.37	-0.18	0.0018
		(0)	(0.277)	(0)	(0)	(0.017)
<i>negsnp</i>	3	-1.05	0.02	-0.33		-0.0032
		(0)	(0.710)	(0)		(0.007)
	4	-0.98	0	-0.29	0.20	-0.0053
		(0)	(0.937)	(0)	(0)	(0)

5.3 Long term holding period

As Savin, Weller and Zvingelis (2007) found significant econometrical results for the HS pattern in their dataset, we thought it would be interesting to look if we can reproduce their results, and if those results can be found for the other patterns. For each pattern, bandwidth multiplier and holding period we construct a portfolio. This gives a total of 72 portfolios.

We begin by looking at the returns of the constructed portfolios, presented in Table 4. The buy patterns, IHS and DBOT, have returns between 0.0060 and 0.0084, significantly larger than zero in all but one case. This is similar to the average return of *snp*, which is 0.0077 and also significant. The other patterns are sell patterns and their mean return ranges between -0.0070 and -0.0114, significant in all but one case. These figures should be compared with the significant return of *negsnp* which is -0.0106. As the holding period n increases (each stock is held for a longer time after a pattern) and the bandwidth multiplier φ decreases (more patterns are found) the average number of stocks included in the portfolio each day increases. Under the assumption that the pattern selection is random, a larger number of stocks in the portfolio should result in a return that converges faster into the return of *snp* or *negsnp*, depending on the type of position. The mean returns are suggesting that the portfolios are behaving like a portfolio that is constructed by randomly selecting stocks. The mean returns are qualitatively the same if we apply the bandwidth filter.

As a second step in the analysis we look at the alphas of the portfolios from a four factor regression, presented in Table 5. When we apply the filter, the results are qualitatively the same. When using the three factor model instead of the four factor model to estimate the alphas, the point estimates have the same signs but more of them are insignificant, thus behaving as the alpha of *snp* (results not presented). As comparison, the alphas of a four

Table 4 Mean return of portfolios with long term holding period.

Pattern	n	Mean rp				p-value $\leq 5\%$			
		$\varphi = 1$	1.5	2	2.5	1	1.5	2	2.5
HS	20	-0.0113	-0.0112	-0.0108	-0.0105	x	x	x	x
	40	-0.0105	-0.0107	-0.0103	-0.0108	x	x	x	x
	60	-0.0109	-0.0109	-0.0109	-0.0114	x	x	x	x
IHS	20	0.0078	0.0080	0.0078	0.0084	x	x	x	x
	40	0.0078	0.0078	0.0077	0.0078	x	x	x	x
	60	0.0079	0.0077	0.0076	0.0076	x	x	x	x
BTOP	20	-0.0098	-0.0105	-0.0093	-0.0103	x	x	x	x
	40	-0.0091	-0.0089	-0.0094	-0.0096	x	x	x	x
	60	-0.0099	-0.0098	-0.0101	-0.0107	x	x	x	x
BBOT	20	-0.0091	-0.0090	-0.0091	-0.0070	x	x	x	-
	40	-0.0100	-0.0095	-0.0099	-0.0094	x	x	x	x
	60	-0.0103	-0.0104	-0.0099	-0.0108	x	x	x	x
DTOP	20	-0.0089	-0.0095	-0.0094	-0.0097	x	x	x	x
	40	-0.0092	-0.0093	-0.0095	-0.0094	x	x	x	x
	60	-0.0098	-0.0099	-0.0103	-0.0101	x	x	x	x
DBOT	20	0.0060	0.0067	0.0067	0.0075	-	x	x	x
	40	0.0079	0.0075	0.0076	0.0076	x	x	x	x
	60	0.0082	0.0078	0.0079	0.0079	x	x	x	x

factor model for snp and $negsnp$ are 0.0018 and -0.0053 respectively, both significant. Here again, it seems like the portfolios are performing similarly to the dataset.

Our conclusion is that the portfolios seem to converge into the underlying dataset, and that there is no evidence that the pattern recognition algorithm is able to pick stocks that perform better than the aggregated dataset for long term holding periods. Before moving on, we will try to explain why Savin, Weller and Zvingelis (2007) found statistical evidence of risk adjusted excess returns. The difference between their and our approach is the window length and years that are analyzed. The authors' window length is 63 days, and they look at S&P 500 stocks in the years 1990-1999, and they only look at the HS pattern. They do not use short selling as we do - instead they are looking for negative performance of a buying strategy.

First of all, the authors do not find that the mean returns are significantly different from zero. That is, their portfolio cannot be profitable on a standalone basis. Instead, they find a negative alpha for their HS portfolios. We believe that this alpha should not be interpreted as a result of the algorithms' stock picking ability. Instead, we believe that it is a result of regressing returns without dividends on the market portfolio. We have reproduced the

Table 5 Four factor risk adjusted excess returns (α^{rpx}) of portfolios with long term holding period.

Pattern	n	α^{rpx}				p-value $\leq 5\%$			
		$\varphi = 1$	1.5	2	2.5	1	1.5	2	2.5
HS	20	-0.0056	-0.0051	-0.0046	-0.0043	x	x	x	x
	40	-0.0045	-0.0046	-0.0041	-0.0048	x	x	x	x
	60	-0.0048	-0.0049	-0.0049	-0.0055	x	x	x	x
IHS	20	0.0011	0.0015	0.0016	0.0021	-	-	-	-
	40	0.0012	0.0011	0.0009	0.0008	-	-	-	-
	60	0.0014	0.0011	0.0009	0.0007	-	-	-	-
BTOP	20	-0.0040	-0.0043	-0.0026	-0.0039	x	x	-	x
	40	-0.0032	-0.0025	-0.0027	-0.0032	x	-	-	x
	60	-0.0041	-0.0034	-0.0037	-0.0046	x	x	x	x
BBOT	20	-0.0046	-0.0042	-0.0045	-0.0030	x	x	x	-
	40	-0.0052	-0.0048	-0.0054	-0.0049	x	x	x	x
	60	-0.0058	-0.0057	-0.0054	-0.0064	x	x	x	x
DTOP	20	-0.0032	-0.0038	-0.0037	-0.0038	x	x	x	x
	40	-0.0033	-0.0034	-0.0036	-0.0032	x	x	x	x
	60	-0.0041	-0.0043	-0.0047	-0.0043	x	x	x	x
DBOT	20	-0.0014	-0.0001	0.0000	0.0012	-	-	-	-
	40	0.0013	0.0009	0.0008	0.0007	-	-	-	-
	60	0.0017	0.0014	0.0014	0.0011	x	-	-	-

results of Savin, Weller and Zvingelis (2007) by looking at the ex-dividend results of the portfolios. When regressing the ex-dividend returns of our portfolios (as well as the aggregated dataset) between the years 1990-1999, there are negative alphas (for a buy strategy). The negative alphas exists for portfolios that are constructed by randomly picking returns as well. Interestingly, these alphas disappear when looking at the period 2000-2008. Our guess is that this is a result that can be derived from the sample rather than the performance of the HS pattern. The period 1990-1999 includes the dot-com boom. During this period IT stocks experienced a tremendous growth, often without paying dividends (notice that many of those companies never were profitable). The S&P 500 stocks, on the other hand, are mature corporations with rather high payout ratios. Therefore, regressing ex-dividend returns of the S&P 500 companies on the market factor during this period yields a negative alpha. We conclude that portfolios created using a long term holding period does not give returns or risk adjusted excess returns that beat the market. The previously found performance of the HS pattern in the years 1990-1999 in S&P 500 stocks is a sample specific result.

Table 6 Mean return of portfolios with short term holding period.

Pattern	Filter	Mean rp (bps)				p-value $\leq 5\%$			
		$\varphi = 1$	1.5	2	2.5	1	1.5	2	2.5
HS	-	-135	-87	-37	-47	-	-	-	-
IHS	-	34	139	132	121	-	x	x	x
BTOP	-	-59	-91	25	-22	-	-	-	-
BBOT	-	-25	-90	-27	37	-	-	-	-
DTOP	-	-81	-69	-78	-80	-	-	-	-
DBOT	-	36	-21	6	31	-	-	-	-
HS	x	-124	-44	49	25	-	-	-	-
IHS	x	53	137	99	51	-	x	x	x
BTOP	x	39	9	55	52	-	-	x	x
BBOT	x	35	-28	-10	15	-	-	-	-
DTOP	x	-101	-87	-82	-43	-	-	-	-
DBOT	x	-12	-12	-13	29	-	-	-	-

5.4 Short term holding period

The next step is to look at the short term holding period returns. The motivation behind this is that Lo, Mamysky and Wang (2000) found that the one day returns following a window with a pattern are different from returns not following such a window. They used two statistical tests, the *Goodness-of-Fit* test and *Kolmogorov-Smirnov* test, to show this. However, the results do not say anything about the economical performance of their algorithm. Therefore, we find it interesting to analyze the algorithms' performance with econometrical methods.

We begin by looking at the mean returns of the portfolios, presented in Table 6. The first thing we notice is that the returns are not as uniform as they were when we looked at long term holding periods. The next interesting thing is to note that BTOP, which is a sell pattern, is able to give returns that are significantly larger than zero. The HS based portfolios, also a short selling pattern, is also able to produce positive mean returns (using the filter and bandwidth multiplier $\varphi = 2$ and 2.5), although they are not significant.

As the underlying dataset is increasing 9.74% per year on average, a strategy that is randomly short selling stocks would on average give negative results, converging into the performance of *negsnp*, which is significantly negative. On the other hand, a strategy randomly going long would yield an average return that converges into the significantly positive results of *snp*. Therefore, the positive mean returns of the HS and BTOP portfolios are, in a way, more interesting than those of IHS at this stage. However, these results cannot be interpreted alone as the portfolios are allowed to earn the

Table 7 Four factor risk adjusted excess returns (α^{rpx}) of portfolios with short term holding period.

Pattern	Filter	α^{rpx} (bps)				p-value $\leq 5\%$			
		$\varphi = 1$	1.5	2	2.5	1	1.5	2	2.5
HS	-	-83	-24	14	3	x	-	-	-
IHS	-	-42	90	76	55	-	-	-	-
BTOP	-	-14	-42	65	28	-	-	-	-
BBOT	-	-7	-96	-21	54	-	-	-	-
DTOP	-	-15	-7	-11	-15	-	-	-	-
DBOT	-	-33	-92	-57	-21	-	x	-	-
HS	x	-85	-9	65	25	x	-	x	-
IHS	x	-1	79	47	17	-	x	-	-
BTOP	x	41	13	46	45	-	-	-	-
BBOT	x	25	-68	-38	2	-	-	-	-
DTOP	x	-49	-49	-44	-14	-	-	-	-
DBOT	x	-81	-80	-66	-22	x	x	-	-

risk free rate, which also is significantly positive. We note that the HS, IHS and BTOP seems to show promising results when using the bandwidth multipliers $\varphi = 2$ and 2.5 and the bandwidth filter.

When we look at the alphas of the portfolios, we see similarities with the mean returns. HS, IHS and BTOP have positive alphas when we look at bandwidth multiplier $\varphi = 2$ and 2.5. However, those alphas are not significant (with one exception). The results presented in the Table 7 are for the four factor model, but they are similar when using the three factor model. None of the portfolios have an alpha that is significantly larger than the alpha of *snp*, i.e. meeting criteria [4-5]. Therefore, we do not have any portfolio that is over-performing. Another difference between the short and long term holding period is the number of stocks that are included in the portfolio every day. As we argued earlier, the number of stocks in the long term holding period portfolios is large enough to converge into the performance of the underlying dataset. The number of stocks drops significantly when we hold the stocks only for one day after an identified pattern (compared to a holding period of 20, 40 and 60 days). The fact is that the drop is so large that the portfolios are earning the risk free rate for a large number of days. In Table 8 we show the number of days each portfolio contains zero stocks, and hence is earning the risk free rate.

For example, if we look at BTOP with bandwidth multiplier $\varphi = 2.5$, the portfolio is on average earning 6.5% per year ($e^{0.0052301 \cdot 12} - 1$) while it is earning the risk free rate (which is 4.0% in the sample period) during 4,196 days out of the 4,750 days. This increase is done by taking short positions

Table 8 The number of passive days for the short term holding period portfolios.

Pattern	No filter				Filter			
	$\varphi = 1$	1.5	2	2.5	1	1.5	2	2.5
HS	251	446	730	1,062	1,221	2,017	2,778	3,390
IHS	223	441	800	1,111	1,216	2,097	2,833	3,445
BTOP	1,265	1,666	2,024	2,432	2,760	3,343	3,810	4,196
BBOT	1,489	1,982	2,400	2,742	3,061	3,621	3,952	4,253
DTOP	536	702	852	1,026	1,399	1,792	2,127	2,502
DBOT	623	821	1,033	1,190	1,616	1,973	2,331	2,607

from a set of stocks that are increasing on average.

In order to cope with the problem with few active positions in the portfolios, we construct portfolios by combining the HS, IHS and BTOP patterns with bandwidth multiplier $\varphi = 2$ and 2.5 and the bandwidth filter applied. These three patterns have shown signs of being able to give positive returns when the filter is applied.

Before we look at the results from the combined portfolios, we compare our findings to what Lo, Mamaysky and Wang (2000) find. The authors analyze stocks from NYSE and NASDAQ separately. If we look at the signs of their mean returns, only HS, IHS and BTOP have the correct signs (according to our assumptions of what each pattern indicate) in the two different exchanges. This similarity is interesting to note, as their sample range over the years 1962-1996. Mossberg and Undin (2007) also have correct signs for these three patterns in their sample.

The performance of the portfolios constructed from the three patterns is interesting. We find that we can construct six different combinations that are meeting the over-performance criteria. Those are presented in the Table 9. There are several things to note. First of all, the difference between the alpha and mean return is small. This suggests that the portfolios have low factor loadings. We also note that the momentum factor is insignificant in all cases. This suggests that the positive returns in the portfolios do not stem from loading on the momentum factor.

The portfolio returns average between 0.0111 and 0.0136 per month with continuous compounding. This corresponds to 14.3% - 17.8% per year, figures larger than the average return of the aggregated dataset which is 9.6% and the risk free rate which is 4.0%. Even if the number of passive days is smaller when we combine the patterns, the portfolios are still earning the risk free rate for between 1,093 and 2,193 days. We have also calculated the active performance, which is the annualized return when we look at mean returns of

Table 9 Performance of portfolios constructed by including investment opportunities from different patterns and bandwidth multipliers.

HS	IHS	BTOP	$\varphi = 2$	$\varphi = 2.5$	rp_x (bps)	α^{rp_x} (bps)	pass days	tot pos	p-val umd	D-P (bps)	A-P (%)
-	x	x	x	-	136	106	2,193	4,599	0.35	10.72	30.7
-	x	x	x	x	133	116	1,996	5,245	0.60	9.81	27.8
x	x	-	x	-	115	114	1,537	6,922	0.70	7.36	20.2
x	x	-	x	x	112	115	1,366	7,813	0.80	6.88	18.7
x	x	x	x	-	114	120	1,266	8,122	0.61	6.82	18.6
x	x	x	x	x	111	131	1,093	9,232	0.75	6.41	17.4

By combining three patterns (HS, IHS and BTOP) with two bandwidth multipliers ($\varphi = 2$ and 2.5) we are able to construct six portfolios that are meeting the over-performance criteria. Each row represents one of the six portfolios. The first five columns indicate what combination of patterns and multiplier the corresponding portfolio has. The remaining columns are portfolio descriptives: rp_x is the mean excess return, α^{rp_x} is the four factor risk adjusted excess return, “pass days” is the number of days the portfolio is earning the risk free rate, “tot pos” is the number of positions the portfolio takes in stocks, “p-val umd” is the p-value of the momentum factor from the four factor regression, D-P (daily performance) is the average continuously compounded daily return of the portfolio during the days when the portfolio holds one or more stocks, and A-P (active performance) is the corresponding annualized simple return.

those days that are not earning the risk free rate. These annualized returns range between 17.4% and 30.7%. We have performed the analysis with ex-dividend data as well. The qualitative results are the same, and the numerical performance is very similar (the active performance ranges between 18.9% and 34.2% per year).

5.5 Case study of a combined portfolio

In order to get some insights of the performance of the combined portfolios, we look at the performance of the portfolio combined from HS and IHS patterns with bandwidth multiplier $\varphi = 2$ and 2.5 (this corresponds to the fourth row in Table 9). The reason for choosing this portfolio is that the two patterns are probably the most popular.

The portfolio earns an average return of 0.0112 per month, which corresponds to 14.4% per year. It earns the risk free rate for 1,366 out of the 4,750 days and takes 7,813 positions (3,823 long and 3,990 short). When only looking at the days when the portfolio takes positions in stocks, the portfolio earns an average of 18.7% per year.

When we perform a three and four factor regression we find that the factor coefficients are insignificant (results not presented). The alpha (the risk adjusted excess return) is significant and similar to the mean return. The insignificant factor coefficients are probably due to the mixture of both short and long positions in stocks with different characteristics. Secondly, the

Table 10 Annual return and gross return of the combined HS and IHS portfolio.

Year	total pos	Annual return (%)			Cum. gross return (%)		
		rpx	mx	snp	rpx	mx	snp
1990	266	-15	-6.6	-12	85	93	88
1991	356	47	27	30	126	118	115
1992	427	11	5.4	12	140	125	129
1993	453	-1.1	8.4	13	138	135	145
1994	498	-0.5	-4.4	-2.2	137	129	141
1995	434	4.1	29	26	143	166	178
1996	447	17	15	14	168	191	204
1997	430	26	24	23	211	237	250
1998	324	16	17	8.6	246	277	272
1999	303	53	20	7.3	375	332	292
2000	221	12	-16	4.7	421	278	305
2001	264	64	-14	-2.1	693	239	299
2002	320	10	-22	-18	762	186	246
2003	479	16	32	41	883	244	346
2004	487	1.8	12	16	899	273	402
2005	549	-18	4.2	4.8	734	284	421
2006	528	11	11	11	816	316	468
2007	654	6	2.5	-3.6	865	323	451
2008	373	19	-39	-39	1,034	197	275

portfolio has an excess return equal to zero for 1,366 days, and consequently not loading anything on the factors these days.

The conclusion is that the factor models fit the returns of the portfolio badly. We continue our study by looking at the confidence intervals for the portfolio return rp and excess return rpx . As our portfolio is meeting the over-performance criteria, the portfolio return rp has a 95% confidence interval that ranges over positive values. rpx is also significantly larger than zero, which is not the case for the market factor mx and the excess returns of the aggregated dataset snp . The Sharpe ratio is 2.56 for our portfolio, which should be compared to 1.02 for mx and 1.42 for snp .

An interesting quality of our portfolio is the low correlation with the market. We saw that the market coefficient was insignificant in the regression. To enlighten this, we calculate the correlations between the excess returns. We find that $Cov(rpx, mx) = -0.056$ and $Cov(rpx, snp) = -0.067$. The result is that the portfolio returns have a negative but small correlation with the market and the aggregated dataset.

In order to get a feeling for how the portfolio performs over the 19 years, we look at the excess returns rpx , and compare them with mx and snp . The low correlations are exemplified here - the portfolio is performing well during

the dot com bubble burst and during the year of the financial crisis. Over the whole period, the portfolio's gross excess return is 1,034%, compared to 197% for *mx* and 275% for *snp*. The over-performance criteria also hold if year 2008 is excluded from the analysis. Table 10 presents the annual excess returns.

To give a hint of how strong the over-performance criteria are, we apply them on all the stocks in our dataset, with the restriction that they have at least 12 whole months of observations. This leaves us with 925 stocks and out of them only three are over-performing. Another way of testing this is to check how well a portfolio with the same number of positions (3,823 long and 3,990 short) randomly selected from the dataset perform. We construct 1,000 such portfolios and conclude that only five out of them are over-performing. The interpretation of this is that the p-value of the over-performance of our portfolio is 0.5%.

5.6 Accounting for trading costs

An important step of evaluating a trading strategy is to account for trading costs. This is where the success of our portfolios ends. Since the holding period is one day, the portfolio is opening and closing its positions every day with active positions. On average, the combined portfolios earn between 6.4 and 10.7 basis points per day with active positions.

Jegadeesh and Titman (1993) cite a figure for one-way transaction costs of 23 basis points for an institutional trader and describe the figure of 0.5% as “conservative”. Lesmond, Ogden and Trzcinka (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 costs for NYSE stocks had fallen to less than 18 basis points in 2000. We also checked the transactional costs at Swedish online banks. We found that an active trader can set up a trading account with transactional fees between 3.9 and 8.5 basis points, depending on bank. However, all these figures are for long positions. Short positions should have even larger transactional costs. We can conclude that even if we use a very generous figure of 5 basis points per one-way transaction, this would wipe out all the returns generated by the algorithm. Therefore we must conclude that the portfolios are not able to beat a buy and hold strategy of the aggregated dataset (the S&P 500 index) or even the risk free rate.

5.7 Is technical analysis profitable?

We have landed at the classical dilemma of many trading strategies - they are profitable as long as you neglect the trading costs. Our findings suggest that our trading strategy cannot be profitable on a stand-alone basis. However, we can draw some valuable insights from our analysis.

A possible explanation to our results is that they simply are the result of short-term autocorrelation in stock returns. We have shown that the returns increase when we apply the bandwidth filter and use high bandwidth multipliers. Qualitatively, this introduces more smoothing. Secondly we notice that HS and BTOP (sell patterns) end with a maximum, whereas IHS (buy pattern) ends with a minimum. Suppose that this ending extremum is at time t . As the amount of smoothing increases, the prices at $t + 1$, $t + 2$ and $t + 3$ must be increasingly smaller than the price at t in order to identify t as a local maximum. The opposite holds for identifying a local minimum. Then the algorithm goes long or short in the stock at time $t + 4$, without looking forward in time. If there is a positive autocorrelation in returns, then negative or positive returns at time $t + 1$, $t + 2$ and $t + 3$ will on average be followed by a negative or positive return at $t + 4$.

The results of our analysis can be interpreted in a more general context, beyond the topic of technical analysis. We have shown that a set of relatively simple rules applied on historical prices is able to pick out *future* returns that are above or below average. This must be interpreted as a rejection of the efficient market hypothesis.

6 Conclusions

By implementing the algorithm of Lo, Mamaysky and Wang (2000) we have repeated the authors' results by using a different sample and a different statistical evaluation method. We have also shown that the algorithm is able to pick returns that average over 30% per year, however heavy trading most probably cancel out this high performance. We found that the momentum factor do not explain the above average returns that the algorithm generates, as well as the other conventional risk factors. We have also shown that the algorithm performs better when the degree of smoothing is high.

Our results contradict the results presented in Savin, Weller and Zvingelis (2007) and we conclude that the algorithm should not be used with holding periods that are 20 days or longer.

6.1 Suggestions for further research

It would be interesting to use a much larger dataset, for example stock prices from all major stocks exchanges. We saw that one of the problems with the patterns that performed well was that they were too few to create investment opportunities every day. A larger set of stocks would make it possible to take active positions every day, and reveal if returns over the whole sample could reach the levels of “active performance” we found.

Another approach is to test if the performance of the patterns we found stem from trading in autocorrelation. For example, by constructing a zero-cost portfolio out of returns following three days of returns that are exceptionally high or low, a short-term momentum factor can be created. By regressing pattern portfolios on such a risk-factor, it would be possible to see if the results from the pattern recognition algorithm stem from autocorrelation.

References

- Brock, W., J. Lakonishok, and B. LeBaron. (1992). "Simple Technical Trading Rules and the Stochastic Properties of Stock Returns." *Journal of Finance* 47, 1731-1764.
- Carhart, M. M. (1997). "On persistence in mutual fund performance." *Journal of Finance* 52, 57-82.
- Chang, P. H. K., and C. L. Osler. (1999). "Methodical Madness: Technical Analysis and the Irrationality of Exchange Rate Forecasts." *Economic Journal* 109, 636-652.
- Chopra, N., J. Lakonishok, and J. R. Ritter. (1992). "Measuring Abnormal Performance: Do Stocks Overreact?" *Journal of Financial Economics* 31, 235-268.
- DeBondt, W. F. M., and R. H. Thaler. (1985). "Does the Stock Market Overreact?" *Journal of Finance* 40, 793-808.
- Dooley, M. P., and J. 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*. Cambridge, MA: Ballinger, 43-69.
- Edwards, R. D., and J. Magee. (1997). *Technical Analysis of Stock Trends*, 7th Edition, Chicago, Ill: John Magee Inc.
- Fama, E. F. (1965). "The Behavior of StockMarket Prices." *Journal of Business* 38, 34-105.
- Fama, E. F. (1970). "Efficient Capital Markets: A Review of Theory and Empirical Work." *Journal of Finance* 25, 383-417.
- Fama, E. F., and K. R. French. (1993). "Common Risk Factors in the Returns on Stocks and Bonds." *Journal of Financial Economics* 33, 3-56.
- Härdle, W. (1990). *Applied Nonparametric Regression*, Cambridge, UK: 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, C. M. (2002). "A Century of Stock Market Liquidity and Trading Costs." Working Paper.

Lesmond, D. A., J. P. Ogden, and C. A. Trzcinka. (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, 451-474.

Lo, A. W., and A. C. MacKinlay. (1988). "Stock market prices do not follow random walks: Evidence from a simple specification test." *Review of Financial Studies* 1, 41-66.

Lo, A. W., and A. C. MacKinlay. (1999). *A Non-Random Walk down Wall Street*, Princeton, NJ: Princeton University Press.

Lo, A. W., H. Mamaysky, and J. Wang. (2000). "Foundations of Technical Analysis: Computational Algorithms, Statistical Inference and Empirical Implementation." *Journal of Finance* 55, 1705-1765.

Mossberg, J., and J. Undin. (2007). "On the Informative Content of Technical Indicators - a Systematic Approach to Pattern Recognition." Thesis within Finance, Stockholm School of Economics.

Neely, C. J., P. A. Weller, and R. Dittmar. (1997). "Is Technical Analysis in the Foreign Exchange Market Profitable? A Genetic Programming Approach." *Journal of Financial and Quantitative Analysis* 32, 405-426.

- Neftci, S. (1991). "Naive trading rules in financial markets and WienerKolmogorov prediction theory: A study of technical analysis." *Journal of Business* 64, 549-571.
- Savin, G., P. Weller and J. Zvingelis. (2007). "The Predictive Power of "Head-and-Shoulders" Price Patterns in the U.S. Stock Market." *Journal of Financial Econometrics* 5, 243-265.
- Shiller, R. J. (2000). *Irrational Exuberance*, Princeton, NJ: Princeton University Press.
- Sullivan, R., A. Timmermann, and H. White. (1999). "Data-Snooping, Technical Trading Rule Performance, and the Bootstrap." *Journal of Finance* 54, 1647-1691.
- Sweeney, R. J. (1986). "Beating the Foreign Exchange Market." *Journal of Finance* 41, 163-182.
- Treynor, J. L., and R. Ferguson. (1985). "In Defense of Technical Analysis." *Journal of Finance* 40, 757-773.

A Appendix




Figure 1 Head-and-Shoulders.




Figure 2 Inverted Head-and-Shoulders.




Figure 3 Broadening top.




Figure 4 Broadening bottom.




Figure 5 Double top.




Figure 6 Double bottom.