Journal of Business Finance & Accounting, 30(1) & (2), January/March 2003, 0306-686X

On the Existence of Visual Technical Patterns in the UK Stock Market

Edward R. Dawson and James M. Steeley*

1. INTRODUCTION

Technical trading patterns are used by analysts to forecast stock market prices using past prices and a range of summary statistics about trading activity. These patterns that rely on visual inspection of graphs and the identification of certain geometric patterns in the data are viewed by many as the original form of security analysis, dating back to a time before the regular disclosures of financial information, and probably to the earliest days of the oldest stock markets. Academic interest in the value of forecasting stock prices followed the publication of two papers, Roberts (1959) and Osborne (1959), that both suggested that stock market prices were indistinguishable from a series of cumulated random numbers. If stock prices are random, then there would be no value in attempting to forecast them, by any means. The market would be, in current language, efficient.

There followed a flurry of academic papers that tested either the random behaviour of prices, the value of technical trading rules, or both. In the US stock market, Alexander (1961 and 1964) and Fama and Blume (1966) for example, suggested that

* The authors are, respectively, Quantitative Analyst, Park Place Capital; and Professor of Finance at Aston Business School. Any views expressed are those of the authors. They are grateful to participants at the ACCA/CBP/JBFA Capital Markets Conference 2002, the referee and an editor of this special issue for helpful comments.

Address for correspondence: James M. Steeley, Aston Business School, Aston University, Aston Triangle, Birmingham B4 7ET, UK.

filter rules, where investors trade in response to past price movements of certain sizes (filters), could not earn excess returns. In the UK, papers by Dryden (1970) and Cunningham (1973) reached similar conclusions. In general, these early studies supported market efficiency.

More recently, the efficient markets hypothesis has been challenged by an increasing number of studies that suggest that stock returns are not fully explained by common measure of risk. Various return 'anomalies' have been identified relating to calendar periods, such as those across weekends and turns of the year.[1] Variables such as market-to-book ratio and size have been shown to explain expected returns,[2] while other studies have documented predictability in returns measured across a variety of horizons.[3]

Such evidence of return predictability led to a renewed interest in examining the value of technical analysis. Sweeney (1988) re-examined the data used by Fama and Blume (1966). They had found filter rules applied to 15 of the 30 Dow Jones stocks earned excess returns over buy-and-hold alternatives. Sweeney found that 14 of these 15 stocks still produced excess returns for a number of years following the end of the Fama and Blume sample. Moreover, these returns would exceed the commission levels of some traders (floor brokers). Brock et al. (1992) examined moving average and trading range break-out trading rules using US data from 1897 to 1986 and found that these could produce excess returns. However, the results in the paper have recently been challenged by Sullivan et al. (1999). They argue that trading rules are subject to survivorship bias as only those that have been perceived to perform well continue to be examined. If these rules are a small subset of all possible rules, and by chance some rules will always appear to outperform, the results of such trading rule tests will be biased in favour of the rules. They show that the results of Brock et al. are substantially weakened when this bias is corrected, and also when the tests are repeated out-of-sample. Moreover, adjustments for transactions costs, not included by Brock et al., weaken the results still further.

A similar conclusion was reached by Hudson et al. (1994), who replicated the Brock et al. analysis on the UK stock market between 1935 and 1994, and found that although the trading rules do have predictive ability, obtaining returns in excess of a buy and hold strategy was unlikely. Recent studies of filter rules in the UK stock market by Sauer and Chen (1996) and Chelley-Steeley and Steeley (1997) also suggest that filter rules could not earn profits in excess of transactions costs, despite the often strong predictability in observed returns. Goodacre et al. (1999) examined the CRISMA trading system, which is a package of trading rules (Cumulative volume, Relative Strength, Moving Average), on UK data over the period 1987 to 1996 and again found little evidence of excess returns after accounting for risk and transactions costs.

In recent years, the examination of technical trading rules has moved away from conducting tests of the existence of excess returns from applying relatively simple rules to (i) choosing an optimal trading rule, and (ii) testing the predictive ability of the more visual (and so less easily mathematically formulated) rules. To search for optimal trading rules, use has been made of genetic algorithms, which allow the trading system to 'learn' the rules to apply rather than have them exogenously specified. Allan and Karjalainen (1999) applied genetic algorithms to trading rule learning for the S&P 500 index. They found that their endogenous rules could not outperform buy-and-hold alternatives when transactions costs were factored in. Skouras (2001) has shown that learned trading rules can outperform those considered by, for example, Brock et al. (1992).

Alongside the mechanical trading rules such as moving averages and filter rules, technical analysts (or Chartists as they are often known) make wider use of charts of price and volume data. In addition to the information provided by mechanical rules, they search out a variety of geometric patterns in the data. Perhaps the best known of these patterns is the 'head-and-shoulders' pattern that comprises three successive peaks in the price history where the middle peak ('head') is above the first and last peak ('shoulders'). This pattern is interpreted as a forecast of a subsequent fall in prices. Osler and Chang (1995) find head-and-shoulders patterns in exchange rate data using a computerised method that finds local extrema by a 'zigzagging' technique. Once one extremum was found neighbouring extrema were classified as such if a minimum percentage price difference existed between the two. They found that significant profits, even after adjustment for interest rate differentials, risk and transaction costs, could be made on dollar-yen and dollar-mark transactions. This was not the case for dollar-sterling, dollar-franc, dollar-Swiss Franc and dollar-Canadian dollar.

This idea of formally identifying technical trading patterns has been recently applied to the US stock market by Lo et al. (2000). Their study is different from that of Osler and Chang in three main ways. First, they examine a set of 10 visual patterns commonly used by technical analysts, including the head-and-shoulders pattern. Second, they search for the patterns in smoothed data rather than the raw data. Third, although they compare the distributions of returns conditioned on patterns with unconditional returns, they do not examine whether these differences could produce economic profits.

Our aim in this paper is to replicate and extend the work of Lo et al. (2000) using data on the UK market. The Lo et al. (2000) study is, in fact, similar to an earlier UK study by Girmes and Damant (1975). Girmes and Damant identified head-and shoulders patterns in stock price series that had been smoothed using a gradient smoothing technique. They found five times as many patterns in their actual data than in corresponding random data. By replicating the Lo et al. (2000) study, which considers a wider range of technical patterns, and an alternative smoothing technique, it is possible to both extend and update that earlier UK study. In addition, it is interesting to know whether such patterns as were found for the US market in recent years also exist in the UK market, and whether returns distributions are influenced by them. This paper also provides further validation of the specific smoothing techniques used by Lo et al. (2000) to 'clean' the data prior to searching for patterns, and as a replication study can inform the issue of data-snooping biases, which can affect empirical work. Finally, as the simulations required to provide control samples take several days of computer time each, even on the fastest available machines, the additional evidence regarding the value of technical patterns may caution against too hasty conclusions being drawn.

We find evidence in the UK market of the technical patterns of the kind that analysts seek in their charts. The different patterns occur with different frequencies to each other and in different relativities to the frequencies found in the US market. However, the frequency of patterns overall in both markets is very similar, with the frequency in the US market being slightly higher. The pattern frequencies are reasonably stable across time with only the smallest size quintile of firms showing systematically fewer patterns of certain kinds. We also find that the signs of returns conditioned on patterns respond as Chartists would predict, and that conditional returns are in many cases distributed differently from unconditional returns. However, in an extension of the work of Lo et al. (2000), we examine market-adjusted conditional returns and find that the evidence that technical patterns are conditioning returns is much weaker. Overall, our results show less evidence that returns are being influenced by technical patterns than for the US. Both our study and that of Lo et al. (2000) indicate, however, that economic profits arising from the predictive ability of the technical patterns are unlikely to materialise.

The rest of the paper is structured as follows. Section 2 defines the 10 technical patterns that we aim to detect in the data. Section 3 describes the detection method. This first uses a smoothing technique to extract the nonlinear trend in the price history and then searches for the patterns in this smoothed function. Section 4 describes the data set of UK stocks used, the results of applying the pattern detection algorithm, and the statistical analysis of these results. Section 5 provides some further discussion of the results overall.

2. DEFINING TECHNICAL PATTERNS

We define technical patterns in terms of their basic geometrical properties. These are typically the occurrence of local extrema (minima and maxima). We focus on five pairs of technical patterns that are popular with technical analysts (see, for example, Edwards and Magee (1966, Chaps. VII-X): head and shoulders (HS) and inverse head and shoulders (IHS), broadening tops (BTOP) and bottoms (BBOT), triangle tops (TTOP) and bottoms (TBOT), rectangle tops (RTOP) and bottoms (RBOT), and double tops (DTOP) and bottoms (DBOT). While there are other technical patterns and these patterns may be easier to detect, for example, moving averages and support and resistance levels, we choose the patterns that are likely to be more difficult to detect to demonstrate the power of smoothing as a tool to identify technical patterns.

Suppose that n local extrema (maxima and minima) have been identified in a price history. We denote the n extrema by $E_1, E_2, \ldots E_n$ and the dates on which they occur by $t_1^*, t_2^*, \ldots, t_n^*$.

We can then define the following patterns:

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, $E_1, E_2, \ldots, E_5$, such that:

$$\mathrm{HS} \equiv \begin{cases} E_1 \text{ is a maximum} \\ E_3 > E_1, E_3 > E_5 \\ E_1 \text{ and } E_5 \text{ are within 1.5 percent of their average} \\ E_2 \text{ and } E_4 \text{ are within 1.5 percent of their average} \end{cases}$$

$$\mathrm{IHS} \equiv \begin{cases} E_1 \text{ is a minimum} \\ E_3 < E_1, E_3 < E_5 \\ E_1 \text{ and } E_5 \text{ are within 1.5 percent of their average} \\ E_2 \text{ and } E_4 \text{ are within 1.5 percent of their average} \end{cases}$$

Definition 2 (Broadening) Broadening tops (BTOP) and bottoms (BBOT) are characterized by a sequence of five consecutive local extrema, $E_1, E_2, \ldots, E_5$, such that:

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

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

Definition 3 (Triangle) Triangle tops (TTOP) and bottoms (TBOT) are characterized by a sequence of five consecutive local extrema, $E_1, E_2, \ldots, E_5$, such that:

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

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

Definition 4 (Rectangle) Rectangle tops (RTOP) and bottoms (RBOT) are characterized by a sequence of five consecutive local extrema, $E_1, E_2, \ldots, E_5$, such that:

$$\text{RTOP} = \begin{cases} E_1 \text{ is a maximum} \\ \text{tops are within 0.75 percent of their average} \\ \text{bottoms are within 0.75 of their average} \\ \text{lowest top} > \text{highest bottom} \end{cases}$$

$$\text{RBOT} = \begin{cases} E_1 \text{ is a minimum} \\ \text{tops are within 0.75 percent of their average} \\ \text{bottoms are within 0.75 of their average} \\ \text{lowest top} > \text{highest bottom} \end{cases}$$

Introducing the requirement that events happen at least a period apart complicates the definition of double tops and bottoms. The double top starts at a local maximum, E_1, then the highest local maximum E_a occurring after E_1 in the set of all local extrema is located. The two tops are required to be within 1.5 percent of their average and a month apart (22 trading days). Double bottoms are inverted double tops. Formally:

Definition 5 (Double Top and Bottom) Double tops (DTOP) and bottoms (DBOT) are characterized by an initial local extremum, E_1, and subsequent local extrema E_a and E_b such that:

$$E_a \equiv \sup\{P^*_{t_k} : t^*_k > t^*_1, k = 2, \ldots, n\}$$

and

$$E_b \equiv \inf\{P^*_{t_k} : t^*_k > t^*_1, k = 2, \ldots, n\}$$

$$\text{DTOP} \equiv \begin{cases} E_1 \text{ is a maximum} \\ E_1 \text{ and } E_a \text{ are within 1.5 percent of their average} \\ t_a^* - t_1^* > 22 \end{cases}$$

$$\text{DBOT} \equiv \begin{cases} E_1 \text{ is a minimum} \\ E_1 \text{ and } E_b \text{ are within 1.5 percent of their average} \\ t_b^* - t_1^* > 22 \end{cases}$$

These formal definitions characterise the visual features of the patterns that are well established among technical analysts. The head-and-shoulders pattern has E_3 as the 'head' and E_1 and E_5 as the 'shoulders'. Broadening patterns appear as a diverging sequence of extrema, while triangle patterns appear as a converging sequence of extrema. Rectangle formations appear as a sequence of extrema that could be bounded by horizontal parallel lines. Double tops (bottoms) are formed when stock prices increase (decrease) to a certain level, show a considerable decline (rise), and then rebound to the previous high (low) level.

3. THE IDENTIFICATION OF TECHNICAL PATTERNS

The basis of technical analysis is the visual recognition of patterns in the nonlinear evolution of share prices. To capture these patterns, we begin by assuming that share prices satisfy the following relation:

$$P_t = m(X_t) + \varepsilon_t, \quad t = 1, \ldots, T \tag{1}$$

where $m(X_t)$ is some fixed but unknown nonlinear function of a state variable X_t and ϵ_t is white noise.

As the patterns are to be detected in the time series of prices, we set the state variable equal to time, $X_t = t$. Thus, equation (1) permits the time series of prices to be divided into a nonlinear function of time and noise. The nonlinear function, $m(X_t)$, to be estimated, will then be used to detect patterns in the price series.

To estimate the nonlinear function $m(X_t)$, we used a smoothing estimator. This has the effect of removing some of the finer structure of movements ('noise') from the price history, to concentrate on the underlying nonlinear trend. This procedure is designed to represent the rather more informal manner in which technical analysts proceed to identify patterns in their charts. Smoothing recovers nonlinear relations by a process of sequential local averaging of the data.

Imagine that it is possible to observe repeated samples of a share price on a particular date, $X_{t0} = x_0$. If that is the case, a natural estimate of the function $m(x_0)$ would be the average of those repeated sample observations. In reality, this is not possible, but given sufficient smoothness of the function, averaging observed prices in the neighbourhood of that date can effectively achieve the same goal. If prices nearest to the particular date are given greater weight, then the estimator will be closer still to $m(x_0)$.

Formally, for any given x, a smoothing estimator of $m(x)$ can be expressed as:

$$\hat{m}(x) \equiv \frac{1}{T}\sum_{t=1}^{T} \omega_t(x)P_t \tag{2}$$

where the weights $\omega_t(x)$ are large for those prices on dates near date x, and small for those prices on dates far from x. To implement this procedure it is necessary, therefore, to choose a weighting function and define the meanings of 'near' and 'far'.

Following Lo et al. (2000), we used a weighting function constructed from a Gaussian probability density function.[4] Specifically, this means that function $m(x)$ is estimated by:

$$\hat{m}(x) \equiv \frac{1}{T}\sum_{t=1}^{T} \omega_t(x)P_t = \frac{\sum_{t=1}^{T} K_h(x - X_t)P_t}{\sum_{t=1}^{T} K_h(x - X_t)} \tag{3}$$

where $K_h(x)$ is given by the Gaussian probability density:

$$K_h(x) = \frac{1}{h\sqrt{2\pi}} e^{-x^2/2h^2}. \tag{4}$$

Analogous to the role of the variance parameter in probabilistic applications of the gaussian density, so the parameter h controls the reach of the $K_h(x)$. For larger the value of h, the more averaging takes place over a larger neighbourhood around x, and so the smoother will be the estimated function. Using probability density functions as weighting functions in this way is known as kernel regression. The function $K_h(x)$ is called a kernel, and the parameter h is known as the bandwidth.

Selecting the bandwidth h is central the success of $\hat{m}(x)$ in approximating $m(x)$. Too much averaging will lead to a function that is too smooth, while too little averaging will lead to a function that is too choppy.[5] There are several methods for selecting bandwidth in kernel regression.[6] The most heavily used is the least squares cross validation method, that selects the bandwidth that minimizes the cross validation function:

$$\mathrm{CV}(h) = \frac{1}{T}\sum_{t=1}^{T}(P_t - \hat{m}_{h,t})^2 \tag{5}$$

where:

$$\hat{m}_{h,t} \equiv \frac{1}{T}\sum_{\tau \neq t}^{T}\omega_{\tau,h}P_\tau \tag{6}$$

where $\hat{m}_{h,t}$ is the kernel regression estimator applied to the time series of prices, and the summands in (6) are the squared errors of the $\hat{m}_{h,t}$ each evaluated at the omitted observation. Thus the cross validation function measures the ability of the kernel regression to fit each observation when that observation is not itself used. We use the cross validation method in this study.

The search for technical patterns starts with the sample of prices $\{P_t, \ldots, P_T\}$. Rather than search across the entire sample, which may produce many patterns of various durations, we require the pattern to be completed within a subsample of $d+l$ days.[7] Thus, the first subsample is $\{P_1, \ldots, P_{d+l}\}$ and subsamples move on one trading day at a time until the final subsample $\{P_{T-d+1}, \ldots, P_T\}$ is reached. The separation of the subsample length into two components, l and d, allows the time taken for pattern completion, l, to be distinguished from the subsequent time taken for the pattern to be recognised, d. Lo et al. (2000) use $l = 35$ and $d = 3$ and we use the same here.

For each subsample, we estimate a kernel regression using the prices in that subsample, hence:

$$\hat{m}_h(\tau) = \frac{\sum_{s=t}^{t+l+d-1} K_h(\tau - s)P_s}{\sum_{s=t}^{t+l+d-1} K_h(\tau - s)}, \quad t = 1, \ldots, T - (l + d - 1), \qquad (7)$$

where $K_h(z)$ is given in equation (4) and h is the bandwidth parameter obtained using least squares cross validation and then multiplied by 0.3.

As $\hat{m}_h(\tau)$ is a differentiable function of τ, the identification of extrema is straightforward. Specifically, extrema occur when $\hat{m}'_h(\tau)\times\hat{m}'_h(\tau+1) < 0$ (if this product was positive, then the gradients of the function would be in the same direction).[8] In cases, where the closing price remains constant, and $\hat{m}'_h(\tau)\times\hat{m}'_h(\tau+1) = 0$, we need to determine whether the flat spot that we have identified is an inflection or an extrema. Thus, we locate the first instance when $\hat{m}'_h(\tau) \times \hat{m}'_h(\tau + x) \neq 0$. If $\hat{m}'_h(\tau) \times \hat{m}'_h(\tau + x) > 0$, then an inflection has been found, but if $\hat{m}'_h(\tau) \times \hat{m}'_h(\tau + x) < 0$, then an extrema is deemed to have been found, at the midpoint of the flat section.

Figure 1 shows an example of a head-and-shoulders pattern found for the price history of the BOC Group using data between 26 October 1986 and 30 May 2001. The extrema searching program would have identified extrema at days 5, 11, 15, 17, 23, 32 and 35 within the illustrated subsample. The program attempts to match these extrema to the pattern definitions described in Section 3. In this case, the extrema at days 15, 17, 23, 32 and 35 match to a head-and-shoulders pattern.

4. DATA AND RESULTS

Our dataset comprises companies that were contained within the FTSE100 and FTSE250 indices over the period 26 October, 1986 to 30 May, 2001. To ameliorate any effects due to possible nonstationarities, we divide this time period in three subsamples of around 5 years (1,230 trading days). The influence of company size is captured by grouping the companies into quintiles based upon their average capitalisation during the

Figure 1

An Example of a Head and Shoulders Pattern in BOC Group

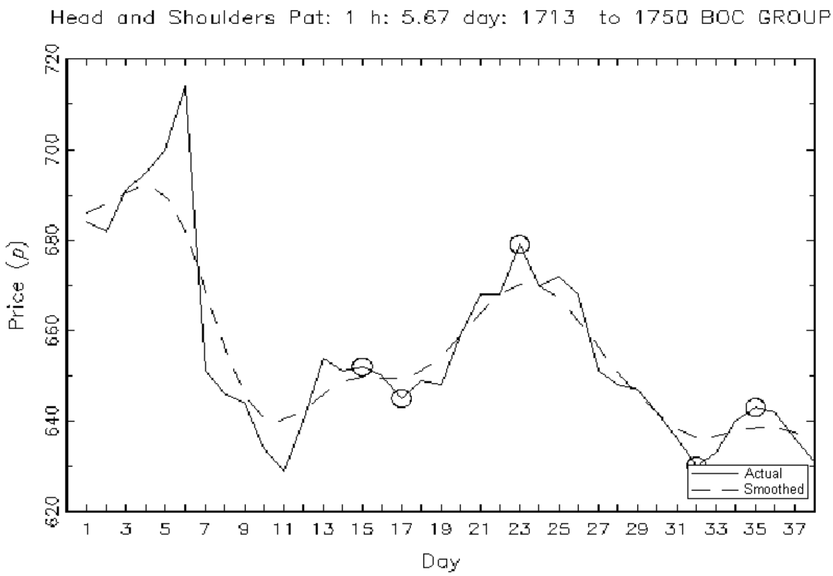

particular sample or subsample period. Daily closing prices and market values were obtained from Datastream. The combined computational complexity of the smoothing and identification procedures means that some sampling is inevitable to produce results from a large population of securities within a reasonable time frame. To obtain a broad cross section of securities, in each 5 year period, we randomly select 15 companies from each quintile, providing a sample of 75 companies within each 5 year period, and a total of 225 across the three subperiods.[9] By comparison, Lo et al. (2000) sampled 10 companies from each quintile of the population of NYSE/AMEX and Nasdaq stocks, but considered a sequence of seven 5 year sub-periods from 1962 to 1996.

Table 1 shows the frequency counts for the number of patterns detected in the smoothed price histories of the 225 companies across 1986 to 2001, and also broken down by size quintile and subsample. The commonest patterns are the head-and-shoulders formations, both normal and inverted, with over 1,480 occurrences of each. These patterns remain

Table 1

Frequency Counts for Each (Marked 'Actual') of the Technical Pattern Detected 1986 to 2001, in 3 Sub-periods, by Quintiles and in a Sample of Simulated Geometric Brownian Motion (SIM. GBM)

Sample	*Totals*	*HS*	*HIS*	*BTOP*	*BBOT*	*TTOP*	*TBOT*	*RTOP*	*RBOT*	*DTOP*	*DBOT*
All Stock 1986–2001											
Actual	8421	1485	1488	370	274	425	471	1174	1148	826	760
Sim. GBM	10495	1797	1717	312	331	338	298	1710	1715	1144	1133
Smallest Quintile 1986–2001											
Actual	851	89	121	20	10	27	33	185	179	91	96
Sim. GBM	1532	171	167	29	22	34	24	361	359	192	173
2nd Quintile 1986–2001											
Actual	1684	271	270	63	48	69	76	278	257	183	169
Sim. GBM	2271	362	348	64	75	68	51	422	433	225	223
3rd Quintile 1986–2001											
Actual	1958	358	336	83	63	117	126	267	254	183	171
Sim. GBM	2429	451	387	52	57	59	49	411	404	255	304
4th Quintile 1986–2001											
Actual	2018	378	389	99	67	116	110	223	244	203	189
Sim. GBM	1998	374	373	81	84	87	88	245	256	215	189

Table 1 (Continued)

Sample	*Totals*	*HS*	*HIS*	*BTOP*	*BBOT*	*TTOP*	*TBOT*	*RTOP*	*RBOT*	*DTOP*	*DBOT*
Largest Quintile 1986–2001											
Actual	1910	389	372	105	86	96	126	221	214	166	135
Sim. GBM	2271	439	442	86	93	90	86	271	263	257	244
All Stock 1986–1991											
Actual	3412	589	562	113	95	150	161	549	552	342	299
Sim. GBM	3750	692	665	125	125	136	119	519	551	421	397
All Stock 1991–1996											
Actual	3140	573	592	145	107	139	173	439	419	273	280
Sim. GBM	3901	615	604	75	99	104	89	776	729	389	421
All Stock 1996–2001											
Actual	1869	323	334	112	72	136	137	186	177	211	181
Sim. GBM	2844	490	448	112	107	98	90	415	435	334	315

Notes:
In each of the 3 time periods 15 companies are selected at random and technical patterns are searched for. They are Head and Shoulders (HS), Inverse Head and Shoulders (IHS), Broadening Top (BTOP), Broadening Bottom (BBOT), Triangle Top (TTOP), Triangle Bottom (TBOT), Rectangle Top (RTOP), Rectangle Bottom (RBOT), Double Top (DTOP) and Double Bottom (DBOT).

the most frequently occurring when the data is divided into subsamples, and also into size quintiles with the exception of the small firm quintile where the rectangle patterns (RTOP and RBOT) are the most frequent. Rectangle patterns are the second most frequently occurring pattern in all other cases. These results are different to those found for the US market. There, the most frequently occurring patterns were double top and bottoms, with head-and-shoulders patterns, occurring with second highest frequency. The least frequently occurring patterns in the UK market were found to be the broadening patterns, these are the diverging sequences of extrema, BTOP and BBOT. This result was also found in the US market. The results in Table 1 also show that size is positively associated with increasing frequency of patterns, with more patterns detected in the smoothed price histories of relatively larger firms. This is likely to be a trading volume effect. Smaller firms, which trade less frequently, are more likely to experience periods of zero returns. The greater the number of zero returns within a 38 day pattern detection window, the less the chance of a pattern being completed within that window.

Most pairs of patterns show a similar number of occurrences for each part of the pair. For example, RTOP and RBOT have 183 and 169 occurrences in the second quintile. For the broadening patterns, however, there is a marked difference on all occasions in the real data. There are always a lot less BBOT than BTOP. Interestingly, this result is consistent with the observation of Robert Edwards in Edwards and Magee that:

> It has been assumed in the past that the Broadening Bottoms must exist, but the writer [Edwards] has never found a good one in his examination of thousands of individual stocks over many years and only one or two patterns which bore resemblance to it (1966, p. 148).

Two further interpretations can also be given to this result, which would also apply for the US market. It could mean that the computerised pattern detection algorithm is finding patterns that, from a Chartist's viewpoint, are not really there. Or, it could mean that Chartists are not as good as a computer at detecting patterns.[10] This latter interpretation is the one made implicitly by Lo et al. (2000).[11]

By way of comparison, Table 1 also reports the frequency counts of patterns detected for a sample of simulated geometric Brownian motion calibrated to match the mean and standard deviation of each company in each of the three 5 year sub-periods.[12] This yields a different set of frequencies. Overall, we find that there are more patterns detected in the simulated data than in the actual data. There are more head-and-shoulders, rectangle and double top and bottom and broadening top patterns, but fewer triangle and broadening bottom patterns. Head-and-shoulders patterns remain the most often detected, with rectangle patterns the second most common. In the simulated data, however, the frequency counts of these two types of patterns are much closer to each other than in the actual data. Our results are in complete contrast to those found for the US market by Lo et al. (2000). They found fewer patterns in their simulated data overall, with only the broadening patterns showing an increase. This difference could however, be an artifact of the simulations, since in both cases only one simulation is carried out.[13] While it is difficult to draw general conclusions from only one simulation, the results do point to differences between the actual and IID lognormal returns. Moreover, our contrasting results suggest that whether there are more or less patterns in actual or simulated data is not possible to gauge without a huge expenditure of computer time.

Figure 2 shows each occurrence of the head-and-shoulders pattern, providing a visual representation of the results for that pattern in Table 1.[14] Each occurrence is plotted with the day of pattern completion along the x-axis and the market capitalization along the y-axis. Each different shade of the plot represents one of the three sub-periods, while the y-axis marks the quintile divisions. While there seems little to indicate any clustering of patterns within any sub-period, there does appear to be relatively few occurences among the smallest quintile. A similar picture emerged for triangle and broadening patterns, while rectangle and double tops and bottoms appeared with no less frequency in the smallest quintile. The size effect observed here for some of the patterns was not so obvious among US stocks and, as suggested earlier, most likely reflects the relatively thin trading among these smaller UK companies.

Figure 2

An Example of Pattern Occurrences Sorted by Date and Firm Size

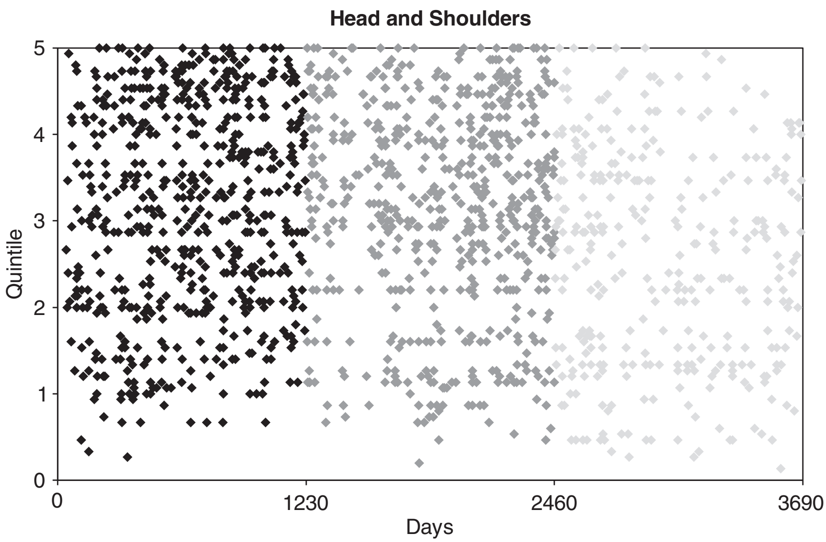

Associated with each technical pattern is a predicted market direction in subsequent days. Price falls (rises) are predicted by (inverted) head-and-shoulders, broadening tops (bottoms), triangle tops (bottoms), rectangle tops (bottoms), and double tops (bottoms). We examine the predictive power of the technical patterns in a number of ways. If technical patterns have no predictive power, then returns conditioned on one of the technical patterns occurring should be no different to unconditional returns. So, we compare the moments of the distribution of returns conditional on a technical pattern occurring to the distribution of unconditional returns. Then, we formally test whether the distribution of returns conditional on a technical pattern occurring is different to the unconditional distribution of returns. In addition, we examine whether the returns conditional on a pattern occurring exceed a passive market return benchmark. If technical patterns have predictive power, then excess returns should be positive for patterns predicting rises, and negative for patterns predicting price falls.[15]

Conditional returns are defined as the one day continuously compounded returns d days after a pattern has completed. Thus, if the sliding window is across days t to $t+l+d-1$ and a pattern is completed at $t+l-1$, then the conditional return for stock i $R_{i,t}$ is:

$$R_{i,t} = \log\left(\frac{P_{i,t+l+d+1}}{P_{i,t+k+d}}\right). \tag{8}$$

For each stock a sample of unconditional returns is obtained by sampling at random from the sub-period price history 100 one-day continuously compounded returns. In order for the unconditional and conditional returns to be compared, they are standardized by subtracting their mean and dividing by their standard deviation.[16] Market adjusted returns are defined as the conditional return minus the contemporaneous market return.

Table 2 contains summary statistics for the conditional and unconditional standardized returns. So far as the means of the conditional returns are concerned, their sign is as expected if the patterns have value, in all cases except for the double patterns. When the data are partitioned by sub-period and size quintile, however, this result is less strong, with rather more variation in sign observed. What is apparent across all the results is the considerable variation among the results for different patterns. For example, the mean, standard deviation, skewness and kurtosis for the RBOT for the full sample are {0.0051, 0.9347, −2.9299, 421.4563}, while for the TTOP they are {−0.0085, 1.0475, −0.1965, 2.4221}. When the conditional return moments are compared to those for the unconditional returns, there is some suggestion that conditioning of returns is taking place.

Table 3 contains summary statistics for the market-adjusted returns. Across all technical patterns and sample periods, the average market adjusted return is negative suggesting that the patterns do not have predictive power in general. Moreover, when the results are partitioned by pattern, only the TTOP and RBOT patterns display the expected sign. The results for the larger firm quintiles, 4 and 5, are more in line with expectations, with patterns HS, HIS, BTOP, BBOT, TTOP and TBOT all displaying the expected sign. But, here again, with the other patterns there is contrary evidence.

Table 2

Summary Statistics (Mean, Standard Deviation, Skewness and Excess Kurtosis of Raw and Conditional One-day Normalised Returns for FTSE 100 and FTSE 250 Companies 1986 to 2001, by Quintiles

Moment	*Raw*	*HS*	*HIS*	*BTOP*	*BBOT*	*TTOP*	*TBOT*	*RTOP*	*RBOT*	*DTOP*	*DBOT*
					All Stock 1986–2001						
Mean	0.0000	−0.0269	0.0214	−0.0076	0.0238	−0.0085	0.0361	−0.0086	0.0051	0.0251	−0.0336
S.D.	0.9950	0.9536	0.9780	1.0996	1.2482	1.0475	1.1588	0.9520	0.9347	0.9904	1.0217
Skew	−0.1849	−0.3664	0.6403	−0.4478	−0.2933	−0.1965	−0.8067	−1.1265	−2.9299	0.2991	−0.2497
Kurt	6.0973	6.3934	7.7613	5.2142	9.7009	2.4221	10.5075	10.7979	42.4563	2.7753	6.0476
					Smallest Quintile 1986–2001						
Mean	0.0000	−0.0798	0.0597	−0.5275	0.0696	0.1479	0.2823	0.1434	−0.1222	−0.0503	0.0875
S.D.	0.9950	0.8919	1.0150	1.0240	0.5262	0.8789	1.2996	0.9598	0.9313	1.1799	1.1506
Skew	−0.6186	0.9824	−0.0519	−0.8965	0.2535	0.7773	0.6204	0.0329	−0.6143	0.1030	1.8538
Kurt	11.2883	3.3748	2.6929	0.8387	1.0304	1.5972	1.9109	5.5206	7.3110	1.2748	9.7991
					2nd Quintile 1986–2001						
Mean	0.0000	−0.0331	0.1461	−0.0796	−0.1366	−0.2257	0.1972	−0.0863	0.1095	0.0596	−0.0489
S.D.	0.9950	0.7796	0.8815	0.9267	1.4896	1.2833	0.9665	0.9025	0.8501	0.9224	0.9477
Skew	−0.3121	−0.5397	2.8798	0.1928	−2.2560	−0.5596	1.0612	−1.1837	−0.8336	1.0666	−0.5571
Kurt	5.2270	2.2256	35.1836	6.7459	12.0785	3.5113	5.3601	5.7346	4.3023	6.0482	3.2786
					3rd Quintile 1986–2001						
Mean	0.0000	−0.0382	−0.0117	0.0083	0.2159	−0.0565	0.1337	−0.0473	−0.0744	0.0341	0.0306
S.D.	0.9950	1.0088	0.9872	0.7928	1.4912	0.9685	1.1008	1.0244	1.1439	0.9009	0.9542
Skew	−0.0992	−1.0430	−0.2839	0.1678	1.7858	−0.2142	−1.1836	−2.2858	−6.7188	0.9751	0.0956
Kurt	5.9418	10.9920	6.8299	1.6578	7.1316	1.8480	8.3561	15.4708	79.6712	4.5426	1.7400

Table 2 (Continued)

Moment	*Raw*	*HS*	*HIS*	*BTOP*	*BBOT*	*TTOP*	*TBOT*	*RTOP*	*RBOT*	*DTOP*	*DBOT*
4th Quintile 1986–2001											
Mean	0.0000	−0.0152	−0.0070	0.0077	0.0108	0.1688	−0.1619	−0.0060	0.0134	0.0001	−0.1318
S.D.	0.9950	1.0486	1.0329	1.0842	1.1661	0.9766	1.3420	0.8820	0.9052	1.0591	1.0586
Skew	0.0366	0.1875	0.6595	0.0572	−1.3467	0.4759	−2.4818	0.6075	0.0602	−0.1926	−1.1410
Kurt	3.6609	1.9150	3.3515	1.1993	6.6760	0.3749	15.7919	5.6388	2.4720	1.4965	6.4425
Largest Quintile 1986–2001											
Mean	0.0000	−0.0113	−0.0220	0.1077	−0.0226	−0.0522	−0.0502	0.0062	0.0709	0.0490	−0.0445
S.D.	0.9950	0.9272	0.9569	1.3682	0.9805	1.0381	1.0710	0.9677	0.7528	0.9524	1.0290
Skew	0.0642	−0.5391	0.6036	−0.9570	−0.6080	−0.2162	1.1690	−1.7465	0.6220	0.0281	−1.0444
Kurt	4.3142	8.5560	2.3736	6.0985	2.4967	1.7715	5.7008	16.3670	1.8086	1.5655	6.7490
All Stock 1986–1991											
Mean	0.0000	−0.0362	0.0249	0.1722	−0.0989	0.0911	0.0577	−0.0459	0.0043	0.0272	−0.0735
S.D.	0.9950	0.9821	1.0147	1.0609	1.3338	1.1238	1.0513	0.9352	0.8542	0.9427	1.1029
Skew	0.0746	−1.0452	0.1532	0.2008	−2.3487	−0.2709	−1.0241	−2.1300	−0.2536	0.1940	0.1201
Kurt	6.6562	8.7368	3.7497	−0.0212	11.5321	3.2620	7.0676	17.4999	4.0412	1.5302	6.9575
All Stock 1991–1996											
Mean	0.0000	0.0073	0.0301	−0.1093	0.2420	−0.1600	0.0194	0.0162	−0.0081	0.0050	−0.0775
S.D.	0.9950	0.8853	0.9958	1.0193	1.1734	0.9874	1.0985	0.9305	0.7992	1.0545	0.9669
Skew	−0.5868	0.3268	1.5334	0.2856	2.4668	−0.3302	0.6377	−0.6950	0.1734	0.3801	−1.2122
Kurt	7.5863	3.9527	12.5483	4.7231	10.8726	2.3821	5.1285	4.6590	1.1757	4.1904	6.3217
All Stock 1996–2001											
Mean	0.0000	−0.0704	0.0001	−0.0573	−0.1386	0.0364	0.0319	0.0432	0.0387	0.0606	0.1002
S.D.	0.9950	1.0134	0.8780	1.2103	1.1912	1.0016	1.3391	1.0425	1.3700	0.9788	0.9505
Skew	−0.0693	−0.0177	−0.4368	−1.4710	−0.3356	−0.0590	−1.6673	0.2555	−6.0154	0.3416	0.4357
Kurt	4.1047	4.8329	5.1091	8.6882	2.2180	1.1908	14.4824	6.6825	58.4191	2.0571	2.1819

Notes:
The conditioning technical patterns are: Head and Shoulders (HS), Inverse Head and Shoulders (IHS), Broadening Top (BTOP), Broadening Bottom (BBOT), Triangle Top (TTOP), Triangle Bottom (TBOT), Rectangle Top (RTOP), Rectangle Bottom (RBOT), Double Top (DTOP) and Double Bottom (DBOT). All returns have been normalised by subtracting their mean and division by their standard deviation of the time period.

Table 3

Summary Statistics (Mean, Standard Deviation, Skewness and Excess Kurtosis of Conditional One-day Market-adjusted Returns for FTSE 100 and FTSE 250 Companies 1986 to 2001, by Quintiles

Moment	*Combined*	*HS*	*HIS*	*BTOP*	*BBOT*	*TTOP*	*TBOT*	*RTOP*	*RBOT*	*DTOP*	*DBOT*
					All Stock 1986–2001						
Mean	−0.0003	−0.0001	−0.0001	−0.0002	−0.0008	0.0000	0.0002	−0.0004	−0.0004	−0.0009	0.0000
S.D.	0.0112	0.0110	0.0099	0.0109	0.0139	0.0107	0.0124	0.0109	0.0128	0.0103	0.0116
Skew	−1.6752	−2.1001	0.1457	0.1688	−2.0205	0.3248	0.2156	−0.5341	−6.3501	0.0023	1.0042
Kurt	38.4143	28.9103	5.2986	1.5902	19.7007	0.9301	5.6893	23.5899	116.125	1.6871	10.2383
					Smallest Quintile 1986–2000						
Mean	−0.0006	−0.0006	0.0000	0.0007	−0.0005	−0.0018	0.0001	0.0001	−0.0009	−0.0032	0.0003
S.D.	0.0134	0.0094	0.0120	0.0112	0.0106	0.0119	0.0174	0.0124	0.0162	0.0136	0.0135
Skew	−1.1991	0.0263	−0.2647	−0.5890	0.3236	0.7046	−0.8714	−0.4497	−3.4076	−0.2250	2.2593
Kurt	16.2863	0.7075	1.1406	3.3430	0.3353	−0.4413	7.9225	12.7547	25.8887	1.1416	14.5608
					2nd Quintile 1986–2001						
Mean	−0.0005	−0.0013	0.0008	−0.0011	−0.0034	−0.0027	−0.0005	−0.0009	0.0005	−0.0006	−0.0005
S.D.	0.0102	0.0096	0.0099	0.0111	0.0215	0.0097	0.0088	0.0094	0.0090	0.0095	0.0101
Skew	−0.6104	0.3227	0.4614	0.4402	−2.9649	−0.2311	−0.2578	1.5542	−0.2051	0.3018	−0.7708
Kurt	13.6513	0.4530	2.6740	2.1241	17.6493	1.0632	1.2399	12.9207	1.4471	2.5777	3.3603
					3rd Quintile 1986–2000						
Mean	−0.0002	−0.0006	0.0003	−0.0009	0.0000	−0.0009	0.0018	0.0000	−0.0019	−0.0005	0.0012
S.D.	0.0119	0.0123	0.0094	0.0097	0.0139	0.0109	0.0131	0.0099	0.0171	0.0098	0.0097
Skew	−4.9062	−4.8995	−0.5308	0.1314	0.6761	0.0417	1.2659	1.7797	−10.5398	0.0034	−0.1708
Kurt	97.7701	57.8439	5.0682	−0.2152	2.9789	0.1818	6.0798	13.5173	145.989	1.1852	0.3290

Table 3 (Continued)

Moment	*Combined*	*HS*	*HIS*	*BTOP*	*BBOT*	*TTOP*	*TBOT*	*RTOP*	*RBOT*	*DTOP*	*DBOT*
					4th Quintile 1986–2001						
Mean	−0.0001	0.0005	−0.0007	0.0000	−0.0001	0.0015	−0.0003	0.0000	−0.0003	−0.0003	−0.0007
S.D.	0.0105	0.0106	0.0100	0.0099	0.0107	0.0104	0.0131	0.0097	0.0109	0.0095	0.0110
Skew	0.6161	0.1514	1.1381	0.5367	−0.3738	0.4687	−0.3978	1.6212	2.2267	0.1860	−0.4493
Kurt	6.6003	1.9315	12.5003	1.7334	2.4941	1.2966	2.4249	12.3998	17.4021	1.3370	2.2774
					Largest Quintile 1986–2001						
Mean	−0.0001	0.0007	−0.0006	0.0004	−0.0004	0.0017	−0.0004	−0.0011	0.0005	−0.0009	0.0002
S.D.	0.0111	0.0112	0.0095	0.0123	0.0107	0.0107	0.0110	0.0133	0.0089	0.0102	0.0145
Skew	−0.6895	−1.9051	−0.4540	−0.0532	−0.8595	0.7389	0.3836	−3.5000	0.9347	0.1915	2.2855
Kurt	16.9085	23.0811	1.8836	1.7270	2.6054	1.7881	3.3985	36.0094	2.4415	0.9902	13.5505
					All Stock 1986–1991						
Mean	−0.0004	−0.0004	−0.0004	0.0001	−0.0021	0.0003	0.0004	−0.0004	−0.0001	0.0016	−0.0003
S.D.	0.0125	0.0133	0.0116	0.0123	0.0178	0.0117	0.0121	0.0130	0.0120	0.0102	0.0139
Skew	−0.6714	−3.0834	0.4757	0.2842	−2.7812	0.1718	1.0738	−0.9783	0.3448	0.1902	1.6250
Kurt	17.2852	32.824	4.8908	0.1415	19.6397	−0.0663	7.7023	23.0419	9.1036	0.6343	11.576
					All Stock 1991–1996						
Mean	0.0000	0.0005	0.0003	−0.0005	0.0016	−0.0005	0.0001	−0.0004	−0.0004	0.0005	−0.0007
S.D.	0.0079	0.0070	0.0072	0.0068	0.0092	0.0072	0.0090	0.0070	0.0093	0.0086	0.0081
Skew	−1.4895	0.3146	−0.6659	−0.5532	1.7718	0.0910	0.7028	−0.0507	−6.6180	0.1585	−0.7916
Kurt	27.9843	1.7649	7.1710	1.3246	8.1191	1.8217	4.8294	1.4551	92.5483	5.3624	4.1005
					All Stock 1996–2001						
Mean	−0.0003	−0.0008	−0.0002	−0.0002	−0.0025	0.0002	0.0002	−0.0006	−0.0014	0.0002	0.0017
S.D.	0.0133	0.0121	0.0110	0.0134	0.0136	0.0124	0.0160	0.0118	0.0203	0.0121	0.0118
Skew	−2.9364	0.2807	0.0131	0.1200	−0.3312	0.4050	−0.3640	1.2093	−9.1154	0.2396	−0.0577
Kurt	55.7391	1.0793	2.2211	1.0933	0.7757	0.7317	2.8989	7.2885	107.032	0.6775	0.9200

Notes:
The conditioning technical patterns are: Head and Shoulders (HS), Inverse Head and Shoulders (IHS), Broadening Top (BTOP), Broadening Bottom (BBOT), Triangle Top (TTOP), Triangle Bottom (TBOT), Rectangle Top (RTOP), Rectangle Bottom (RBOT), Double Top (DTOP) and Double Bottom (DBOT).

The difference between the conditional and unconditional distributions is assessed using two nonparametric statistics: the chi-square goodness of fit test; and the Kolmogorov-Smirnov test. The chi-square test is used to test whether there is a difference between an observed and expected frequency of observations falling into a set of categories. In this application, the categories are chosen to be the deciles of the unconditional returns distribution. Thus, if the conditional distribution is no different to the unconditional distribution, then the expected proportion of conditional returns in each decile would be 10 percent. The chi-square test assesses the deviations from 10 percent across the whole distribution. Specifically, the test statistic, Q, is given by:

$$Q = \sum_{j=1}^{10} \frac{(n_j - 0.10n)^2}{0.10n} \overset{a}{\sim} \chi_9^2 \tag{9}$$

where n_j is the number of observations that fall into decile j and n is the total number of observations. The asymptotic distribution of n_j is given by:

$$\sqrt{n}(n_j - 0.10) \overset{a}{\sim} N(0, 0.10(1 - 0.10)) \tag{10}$$

which is used to provide asymptotic z-statistics. The Kolmogorov-Smirnov tests the difference between two cumulative frequency distributions. Specifically, the maximum difference between the two cumulative distributions, scaled by a function of the sample sizes, is compared to critical values obtained from the (limiting) distribution of the statistic.[17]

Table 4 shows the results of the goodness-of-fit tests. The frequency of conditional returns falling into the unconditional return deciles are significantly different from 10 percent for eight out of the ten patterns. The patterns that do not seem to influence the returns distribution are the broadening patterns, which were the least frequently observed patterns. Table 5 shows the results of the Kolmogorov-Smirnov tests for the equality of the distributions. For the full sample of companies across 1986 to 2001, five of the 10 patterns caused significant differences between unconditional and conditional distributions, namely HS, HIS, RBOT, RTOP and DBOT. For the

Table 4

Goodness of Fit Diagnostics for Conditional One-day Normalised Returns, Conditional on 10 Technical Pattern for a Sample of FTSE 100 and FTSE 250 Companies from 1986 to 2001

	Decile										
Pattern	*1*	*2*	*3*	*4*	*5*	*6*	*7*	*8*	*9*	*10*	*Q (p-Value)*
HS	9.23	11.11	12.46	10.77	10.30	8.82	7.41	10.10	10.44	9.36	25.67
	(−0.56)	(0.80)	(1.77)	(0.56)	(0.22)	(−0.85)	(−1.87)	(0.07)	(0.32)	(−0.46)	(0.00)
HIS	8.20	11.76	11.56	12.90	9.81	7.33	6.38	11.36	10.69	10.01	59.32
	(−0.85)	(0.85)	(0.75)	(0.39)	(−0.08)	(−1.26)	(−1.71)	(0.66)	(0.34)	(0.02)	(0.00)
BTOP	11.89	10.00	9.73	10.27	10.81	7.84	10.54	8.11	9.19	11.62	6.00
	(0.19)	(0.00)	(−0.03)	(0.03)	(0.08)	(−0.22)	(0.05)	(−0.19)	(−0.08)	(0.16)	(0.74)
BBOT	11.31	8.03	11.68	6.93	12.04	9.49	8.76	11.68	8.39	11.68	8.77
	(0.70)	(−1.04)	(0.89)	(−1.63)	(1.08)	(−0.27)	(−0.66)	(0.89)	(−0.85)	(0.89)	(0.46)
TTOP	10.82	9.41	14.59	8.94	7.53	8.47	6.82	10.82	11.29	11.29	19.45
	(0.37)	(−0.26)	(2.03)	(−0.47)	(−1.10)	(−0.68)	(−1.41)	(0.37)	(0.57)	(0.57)	(0.02)
TBOT	9.34	7.43	11.25	12.10	12.95	8.07	6.58	13.16	7.86	11.25	25.11
	(−0.27)	(−1.06)	(0.51)	(0.86)	(1.21)	(−0.79)	(−1.40)	(1.30)	(−0.88)	(0.51)	(0.00)

RTOP	8.01	9.11	11.84	12.86	8.26	12.95	7.07	12.10	8.43	9.37	51.50
	(−0.90)	(−0.40)	(0.83)	(1.30)	(−0.79)	(1.33)	(−1.33)	(0.95)	(−0.71)	(−0.29)	(0.00)
RBOT	7.75	9.15	10.54	14.81	11.06	10.80	5.66	11.06	10.54	8.62	60.95
	(−0.93)	(−0.35)	(0.22)	(1.98)	(0.44)	(0.33)	(−1.79)	(0.44)	(0.22)	(−0.57)	(0.00)
DTOP	10.77	11.14	10.41	11.86	7.63	8.60	7.99	9.32	9.81	12.47	19.64
	(0.48)	(0.70)	(0.25)	(1.15)	(−1.46)	(−0.86)	(−1.24)	(−0.42)	(−0.12)	(1.52)	(0.02)
DBOT	12.50	8.03	10.92	11.32	13.29	8.03	4.34	11.58	9.34	10.66	47.74
	(0.95)	(−0.75)	(0.35)	(0.50)	(1.25)	(−0.75)	(−2.14)	(0.60)	(−0.25)	(0.25)	(0.00)

Notes:
For each pattern the percentage of conditional returns that fall within each of the 10 unconditional return deciles is tabulated. If conditioning on the pattern provides no information, the expected percentage falling in each decile is 10%. Asymptotic z-statistics for the null hypothesis are reported in the parentheses, and χ^2 goodness of fitness test statistic Q is reported in the last column with the p-value in parentheses below the statistic. The 10 technical patterns are: Head and Shoulders (HS), Inverse Head and Shoulders (IHS), Broadening Top (BTOP), Broadening Bottom (BBOT), Triangle Top (TTOP), Triangle Bottom (TBOT), Rectangle Top (RTOP), Rectangle Bottom (RBOT), Double Top (DTOP) and Double Bottom (DBOT).

Table 5

Kolmogorov-Smirnov Test of the Equality of Conditional and Unconditional One-day Return Distribution for FTSE 100 and FTSE 250 Companies from 1986 to 2001 in 3 Sub periods and by Quintiles

Statistic	*HS*	*IHS*	*BTOP*	*BBOT*	*TTOP*	*TBOT*	*RTOP*	*RBOT*	*DTOP*	*DBOT*
				All Stocks 1986–2001						
k-s	2.02	1.729	0.802	0.713	1.093	0.868	1.41	1.935	1.212	1.81
p-value	0.001	0.005	0.541	0.69	0.183	0.438	0.037	0.001	0.106	0.003
				Smallest Quintile 1986–2001						
k-s	2.57	2.167	1.345	0.818	0.802	1.318	2.167	2.372	2.135	1.372
p-value	0.000	0.000	0.054	0.516	0.542	0.062	0.000	0.000	0.000	0.046
				Second Quintile 1986–2001						
k-s	1.356	1.745	1.362	0.741	1.112	1.231	1.657	1.637	1.626	1.553
p-value	0.051	0.005	0.049	0.643	0.169	0.097	0.008	0.009	0.01	0.016
				Third Quintile 1986–2001						
k-s	1.508	0.908	0.455	0.648	0.896	1.244	1.149	1.297	1.115	0.924
p-value	0.021	0.382	0.986	0.794	0.399	0.091	0.143	0.069	0.166	0.36
				Fourth Quintile 1986–2001						
k-s	1.258	1.875	0.829	0.535	0.809	0.805	1.516	1.854	0.753	1.667
p-value	0.084	0.002	0.497	0.937	0.53	0.536	0.02	0.002	0.622	0.008
				Largest Quintile 1986–2001						
k-s	1.34	1.241	1.107	0.58	0.532	1.557	1.12	1.111	0.719	0.689
p-value	0.055	0.092	0.173	0.889	0.94	0.016	0.163	0.17	0.679	0.73

					All Stocks 1986–91					
k-s	1.568	1.457	1.449	0.72	1.204	1.191	2.162	2.24	1.62	2.456
p-value	0.015	0.029	0.03	0.678	0.11	0.117	0.000	0.000	0.011	0.000
					All Stocks 1991–96					
k-s	1.549	1.959	1.469	0.822	1.517	0.935	1.064	1.817	1.09	1.307
p-value	0.016	0.001	0.027	0.509	0.02	0.346	0.207	0.003	0.186	0.066
					All Stocks 1996–2001					
k-s	1.447	1.129	0.847	1.092	0.668	1.085	0.871	1.521	1.169	0.843
p-value	0.03	0.157	0.471	0.184	0.764	0.19	0.434	0.02	0.13	0.476

Notes:
Conditional returns are defined as the daily return three days following the completion of an occurrence of a technical pattern. The 10 technical patterns are: Head and Shoulders (HS), Inverse Head and Shoulders (IHS), Broadening Top (BTOP), Broadening Bottom (BBOT), Triangle Top (TTOP), Triangle Bottom (TBOT), Rectangle Top (RTOP), Rectangle Bottom (RBOT), Double Top (DTOP) and Double Bottom (DBOT). All returns are normalised by the subtraction of the means and division by their standard deviations. □-values are with respect to the asymptotic distribution of the Kolmogorov-Smirnov test statistic.

head-and-shoulders (HS) and RBOT patterns, this result is stable across sub-periods and quintiles, while the others show some variation. The result for the head-and-shoulders pattern appears stronger the smaller the companies being examined and the more recent the sample.

5. DISCUSSION AND CONCLUSION

We have identified technical trading patterns in UK stock data that has been smoothed using kernel regressions. We find that the distributions of returns conditioned on these technical patterns can be significantly different from the unconditional returns distributions. Our results validate the findings of Lo et al. (2000) who examined US stock market data and found similar, but slightly stronger, results in support of the predictive ability of technical patterns.

Our results in combination, however, provide less support for the value of technical trading patterns. As pointed out by Jegadeesh (2000), the information in Tables 1 to 3 is sufficient to construct *t*-tests of the significance of the difference of the conditional mean returns away from either the unconditional mean returns or the market returns. It is clear, without calculation, that the mean returns are not significantly different. The same is also true of the study by Lo et al. (2000).

If the means of conditional and unconditional returns are not significantly different from each other, yet we find (as did Lo et al., 2000) that the distributions are significantly different from each other, then these differences must be the result of higher order moment differences. While the information in time varying standard deviations may be of importance in testing forecasts, it is difficult to interpret these higher order differences in terms of market efficiency, which is primarily mean return based. Indeed, so far as mean returns are concerned, our study provides no reason to question market efficiency.

NOTES

1 See, for example, recent studies by Chang et al. (1998) and Steeley (2001) regarding weekend effects in the US and UK stock markets, respectively, and Sehun (1993) and Chelley-Steeley (1996) regarding risk adjusted turn-of-the-year effects in the US and UK stock markets, respectively.

2 For example, Fama and French (1992) for the US market and Strong and Xu (1997) for the UK market.
3 For example, DeBondt and Thaler (1985), Fama and French (1986), Poterba and Summers (1988), Lo and MacKinlay (1990) and Mills (1991).
4 Although the weighting function is constructed from a probability density function it plays no probabilistic role. It is merely a convenient method to define a weighting scheme, that is, as a 'bell' curve.
5 See Lo et al. (2000, Figures 1–4) for a demonstration using a sine function plus a random error.
6 See, Hardle (1990) for alternative methods.
7 Searching for partially completed patterns, as Lo et al. (2000) point out, would require more structure to be placed on the otherwise flexible smoothing estimator.
8 While Lo et al.'s algorithm compares the signs of neighbouring derivatives, ours compares the product of neighbouring derivatives against zero. The effect is the same.
9 As the sampling is done with replacement, there may be companies in common across sub-periods.
10 This might suggest that the algorithm has not been correctly calibrated. In particular, it might suggest that alternative scalings be applied to the cross validation function. But, of course, the scaling used was selected by Chartists themselves.
11 One is however aware of Samuelson's cautionary note to all forecasters that 'economists have succesfully forecast seven of the last five recessions'.
12 Specifically, the price process was assumed to satisfy:

$$P_t = P_{t-1}e^{\mu+\sigma dz}$$

where dz is a normally and independently distributed random shock, μ is the mean and σ the standard deviation of the associated log returns.
13 Each simulation exercise takes approximately 3 days of pure computational time on a Pentium III 500 before manipulation of the actual results can take place, so a large number of simulations was deemed impractical to run.
14 Similar graphs were produced for each pattern and can be obtained from the authors on request.
15 In the face of a predicted stock price fall, an investor would obtain a return in excess of the market provided the market falls by no more than the security price.
16 Thus, normalised conditional returns are given by:

$$R_{i,t}^n = \frac{R_{i,t} - E(R_{i,t})}{\sigma(R_{i,t})}.$$

The mean return, $E(R_{i,t})$, and standard deviation, $\sigma(R_{i,t})$, are calculated across all conditional returns for that security in the particular sub-period. Thus:
$E(R_{i,t}) = (1/\tau_{i,k}) \sum_{t=T_{k-1}}^{T_k} R_{t,t}I_{i,t}$, and
$\sigma(R_{i,t}) = (1/\tau_{i,k}) \sum_{t=T_{k-1}}^{T_k} (R_{i,t}I_{i,t} - E(R_{i,t}))^2$,
where $T_k = 1230 \times k, k = 1,2,3$ are the sample periods, where $I_{i,t}$ is an

indicator variable taking the value 1 if stock i has a conditional return at time t and zero otherwise, and where $\tau_{i,k}$ is the number of patterns that stock i has in the kth sub-period.
The unconditional returns, $r_{i,t}$, are normalised to:

$$r^n_{i,t} = \frac{r_{i,t} - E(r_{i,t})}{\sigma(r_{i,t})}$$

where $E(r_{i,t}) = (1/100)\sum_{t=1}^{100} r_{i,t}$ and $\sigma^2(r_{i,t}) = (1/100)\sum_{t=1}^{100}(r_{i,t} - E(r_{i,t}))^2$.

17 For further details of this statistic, see Seigel and Castellan (1988). Critial values are provided in their Table $L_{III.}$

REFERENCES

Alexander S. (1961), 'Price Movements in Speculative Markets: Trends or Random Walks', *Industrial Management Review*, Vol. 2, pp. 7–26.
——— (1964), 'Price Movements in Speculative Markets: Trends or Random Walks, No.2', in P. Cootnered (ed.), *The Random Character of Stock Market Prices* (MIT Press, Cambridge, MA).
Allen F. and R. Karjalainen (1999), 'Using Genetic Algorithms to Find Technical Trading Rules', *Journal of Financial Economics*, Vol. 51, pp. 245–71.
Brock, W., J. Lakonishok and B. LeBaron (1992), 'Simple Technical Trading Rules and the Stochastic Properties of Stock Returns,' *Journal of Finance*, Vol. XLVII, No. 5 (December).
Chang, E., M. Pinegar and R. Ravichandran (1998), 'US Day-of-the-week Effects and Asymmetric Responses to Macroeconomic News', *Journal of Banking and Finance*, Vol. 22, pp. 513–34.
Chelley-Steeley, P. (1996), 'Risk, Seasonality and the Asymmetric Behaviour of Stock Returns,' *Journal of Business Finance & Accounting*, Vol. 23, pp. 145–54.
——— and J. Steeley (1997), 'The Impact of Portfolio Diversification on Trading Rule Profits: Some Evidence for UK Share Portfolios,' *Journal of Business Finance & Accounting*, Vol. 24, pp. 759–79.
Cunningham, S. (1973), 'The Predictability of British Stock Market Prices,' *Applied Statistics*, Vol. 22, pp. 315–31.
DeBondt W. and R. Thaler (1985), 'Does the Stock Market Overreact?' *Journal of Finance*, Vol. 40, pp. 793–805.
Dryden, M. (1970), 'Filter Tests of UK Share Prices,' *Applied Economics*, Vol. 1, pp. 261–75.
Edwards, R.D. and J. Magee (1996), *Technical Analysis of Stock Trends* (5th ed., John Magee Inc, Boston, Massachusetts, USA).
Fama, E. and M. Blume (1966), 'Filter Tests and Stock Market Trading,' *Journal of Business*, Vol. 39, pp. 226–41.
——— and K. French (1986), 'Permanent and Temporary Components of Stock Prices,' *Journal of Political Economy*, Vol. 98, pp. 246–82.
——— (1992), 'The Cross-section of Expected Stock Returns,' *Journal of Finance*, Vol. 47, pp. 427–65.
Girmes, D. and D. Damant (1975), 'Charts and the Random Walk,' *Investment Analyst*, Vol. 41.

Goodacre, A., J. Bosher and A. Dove (1999), 'Testing the CRISMA Trading System: Evidence from the UK Market,' *Applied Financial Economics*, Vol. 9, pp. 455–68.
Hardle, W. (1990), *Applied Nonparametric Regression* (Cambridge University Press, Cambridge, Great Britain).
Hudson, R., M. Dempsey and K. Keasey (1996), 'A Note on the Weak Form Efficency of Capital Markets: The Application of Simple Technical Trading Rule to the UK Stock Prices – 1935 to 1994,' *Journal of Banking and Finance*, Vol. 20, pp. 1121–32.
Jegadeesh, N. (2000), 'Discussion,' *Journal of Finance*, Vol. 55.
Lo, A. and C. Mackinlay (1990), 'When are Contrarian Profits Due to Stock Market Overreaction?' *Review of Financial Studies*, Vol. 3, pp. 175–206.
——— H. Mamaysky and J. Wang (2000), 'Foundations of Technical Analysis: Computational Algorithms, Statistical Inference and Empirical Implementation,' *Journal of Finance*, Vol. 55.
Mills, T. (1991), 'Assessing the Predictability of UK Stock Market Returns Using Statistics Based on Multiperiod Returns,' *Applied Financial Economics*, Vol. 1, pp. 241–45.
Osler, C. and K. Chang (1995), 'Head and Shoulders: Not Just a Flaky Pattern,' *Staff Report No. 4* (Federal Reserve Bank of New York).
Poterba J. and L. Summers (1988), 'Mean Reversion in Stock Prices: Evidence and Implications,' *Journal of Financial Economics*, Vol. 22, pp. 27–59.
Sauer, D. and C. Chen (1996), 'Mean Reversion in the United Kingdom Stock Market and its Implications for a Profitable Trading Strategy,' *Journal of Business Finance & Accounting*, Vol. 23, pp. 1379–95.
Sehun, N. (1993), 'Can Omitted Risk Factors Explain the January Effect?' A Stochastic Dominance Approach,' *Journal of Financial and Quantitative Analysis*, Vol. 28, pp. 195–211.
Siegel, S. and N.J. Castellan (1988), *Nonparametric Statistics for Behavioral Sciences*, (2nd ed., McGraw-Hill).
Silverman, B.W. (1986), *Density Estimation for Statistics and Data Analysis* (Chapman and Hall, London).
Skouras, S. (2001), 'Financial Returns and Efficiency as Seen by an Artificial Technical Analyst,' *Journal of Economic Dynamics and Control*, Vol. 25, pp. 213–44.
Steeley, J. (2001), 'A Note on Information Seasonality and the Disappearance of the Weekend Effect in the UK Stock Market,' *Journal of Banking and Finance*, Vol. 25, pp. 1941–56.
Strong, N. and X. Xu (1997), 'Explaining the Cross Section of UK Expected Stock Returns,' *British Accounting Review*, Vol. 29, pp. 1–24.
Sullivan, R., A. Timmerman, and H. White (1999), 'Data-snooping, Technical Trading Rule Performance, and the Bootstrap,' *Journal of Finance*, Vol. 54, pp. 1647–91.
Sweeney, R. (1988), 'Some New Filter Rule Tests: Methods and Results,' *Journal of Financial and Quantitative Analysis*, Vol. 23, pp. 87–92.