The University of Reading

THE BUSINESS SCHOOL FOR FINANCIAL MARKETS

Abnormal Returns in Equity Markets:

Evidence from a Dynamic Indexing Strategy

ISMA Centre Discussion Papers in Finance 2003-02
First Version: January 2003
This Version: April 2003

Carol Alexander
and
Anca Dimitriu

ISMA Centre, University of Reading, UK

The University of Reading • ISMA Centre • Whiteknights • PO Box 242 • Reading RG6 6BA • UK
Tel: +44 (0)118 931 8239 • Fax: +44 (0)118 931 4741
Email: research@ismacentre.rdg.ac.uk • Web: www.ismacentre.rdg.ac.uk

Abstract

This paper investigates the abnormal return generated through a dynamic equity indexing strategy and the extent to which this can be considered evidence against the efficient markets hypothesis. We introduce a new measure of stock price dispersion and show that it is a leading indicator for the abnormal return, where their relationship is based on a switching process of two market regimes. The entire abnormal return is associated with only one of the regimes and this is the prevalent regime during the last few years. The predictive power of the model is demonstrated over different time horizons and in different, real world and simulated stock markets. The strategy remains profitable even after introducing transaction costs, thus proving evidence of temporary market inefficiencies.

JEL classification: C32, C51, G11, G23

Keywords: equity indexing, cointegration, Markov switching, stock prices dispersion

The authors would like to thank Chris Brooks, Apostolos Katsaris and Philip Xu of ISMA Centre for valuable comments. All errors remain our responsibility.

Introduction

The phenomenon of equity indexing has attracted considerable interest in the last ten years, from both academics and practitioners. Equity indexing is the most popular form of passive investment, aiming to replicate the risk and return characteristics of a benchmark, usually a wide stock market index. Currently, it is estimated that more than $1.4 trillion are invested in index funds in the US alone (Blake, 2002). The very reason for adopting a passive strategy rests in a belief in market efficiency, which provides the theoretical foundation of indexing. Traditional capital market theory states that the market portfolio, as defined by Fama (1970),[i] offers the highest level of return per unit of risk. Moreover, empirically, active management has been shown to underperform its passive alternative most of the time, even after transaction costs and administration fees (Jensen, 1968; Elton, Gruber, Das, Hlavka, 1993; Carhart, 1997). In this theoretical framework, where the only way that investors can beat the market over the longer term is by taking greater risks, passive investment and, in particular, indexing, are natural choices.

From an operational perspective, however, one needs to make the distinction between a pure index fund, managed to replicate the performance of the market portfolio/benchmark exactly, and strategies such as enhanced index tracking, that extend it into active management. The latter are constructing well-diversified portfolios that have a stable relationship with the market index and try to take advantage of some pockets of market inefficiency.

According to Jensen's (1978) definition of efficient markets, a trading strategy producing consistent risk-adjusted economic gains, after properly defined transaction costs and over a sufficiently long period of time, is evidence against the efficient market hypothesis (EMH). This approach to markets efficiency, as compared to the previous ones, has the advantage of testability and has subsequently generated a great deal of empirical research. Most of these studies, employing for example technical analysis and filter rules (Alexander, 1964; Fama and Blume, 1966), have shown that even if different trading strategies are successful before transaction costs, after accounting for such costs the profits vanish. Published evidence of trading profitability, after properly defined transaction costs, is rather scarce, Lakonishok and Vermaelen's (1990) paper being one of the very few to document the profitability of some trading rules designed to exploit anomalous price behaviour.

In this paper, we propose the investigation of the abnormal return[ii] generated through a dynamic indexing strategy and the extent to which this can be considered evidence against the EMH. We find a leading indicator for the abnormal return and model their relationship as a switching process between two stock market regimes. The entire abnormal return is associated with only one of the regimes, occurring mostly in the last few years. The predictive power of the model is demonstrated over different time horizons and in different, real-world and simulated stock markets. The strategy remains profitable even after introducing transaction costs, thus proving evidence of temporary market inefficiencies.

The tracking strategy producing the abnormal return analysed in this paper is the cointegration based index tracking, introduced by Alexander (1999). This purely statistical strategy, based solely on the information contained in past stock prices, has some interesting features, such as the fact that a portfolio constructed based on a cointegration relationship with a price weighted index produces, out of sample, a consistent abnormal return, even after accounting for transaction costs (Alexander and Dimitriu, 2002).

The pattern of the abnormal return exhibits a pronounced time-variability: periods of stationary excess return are alternating with periods during which the excess return is accumulated consistently. Moreover, the periods during which most of the abnormal return is produced appear to coincide with the main market crises during the sample period.

Since the only information used to construct the tracking strategy is the history of the stock prices, the cause of the abnormal return should be linked to the time-variability characteristics of the stock prices in the system. We introduce a new measure of the 'cohesion' of stock prices within the market index, which we call the index dispersion, and find that this is a leading indicator of the abnormal return.

Throughout the analysis we justify the conclusions drawn from real-world and simulated stock and index prices. Beginning with the simplest two-stock scenario, we illustrate the connection between stock prices, portfolio weights and index out-performance. In the real-world universe of the Dow Jones Industrial Average (DJIA), we document a significant linear relationship between the abnormal return and the lagged dispersion, which, however, has a considerable time-variability in parameters. To address this issue, we estimate a Markov switching model for the abnormal return and find strong evidence of a latent state variable, which determines the form of the linear relationship between the abnormal return and the stock prices dispersion. The Markov switching model indicates the presence of two regimes having very different characteristics. The first regime, associated with more volatile market conditions, is responsible for the entire abnormal return generated from index tracking. This regime occurs much more frequently during the last few years: since 1999, even though stock markets have been excessively volatile, the prevailing regime has been the one producing the abnormal return. The second regime, less volatile but with no significant over-performance of the tracking strategy, is prevailing at the beginning of the data sample, and indeed during most of the 1990's.

We relate these observations to the long-run equilibrium relationships identified by cointegration and show that the strategy disregards the temporary deviations of stock prices from the equilibrium levels and tracks the index very accurately. However, when disequilibria in stock prices are no longer temporary, but instead represent transitions towards new equilibrium prices, the cointegration-based strategy generates consistent abnormal returns.

To summarise, our contributions to existing research are in two directions: first, we provide additional empirical evidence in the EMH debate and shed some light on the anomalies identified through cointegration and on the mechanism producing the abnormal return. Secondly, our findings have wide implications for the passive investment industry. We show that, without any stock selection or explicit timing attempts, which are attributes of active management, solely through smart optimisation, the benchmark performance can be significantly enhanced, even after accounting for transaction costs. Moreover, the strategy can be applied to replicate any type of value or capitalisation weighted benchmark, not only wide market indexes.

The remainder of the paper is organised as follows: section one reviews the cointegration-based index tracking strategy, defines the index dispersion and motivates a possible relationship between dispersion and the strategy performance; section two examines the real world relationship between the abnormal return and the lagged dispersion in the DJIA and motivates the need for a Markov switching framework; section three introduces the Markov switching model of dispersion and makes statistical inferences; section four demonstrates the predictive power of the Markov switching model by examining the possibility of generating profitable trading rules; section five discusses the implications for EMH; and finally, section six summarises and draws the main conclusions.

1. Cointegration and Index Dispersion

Generally, a pure index fund, managed to replicate exactly the performance of the market portfolio/benchmark, may consist of all the stocks in the benchmark or only a subset of them. Value and price weighted indexes are easy to replicate with portfolios comprising the entire set of stocks and mirroring the benchmark weights, as long as there are no changes in the index composition or in the number of shares in each issue. Such portfolios are self-adjusting to changes in the stock prices, and do not require any rebalancing, except for special events like mergers, additions or deletions, splits or dividends.

Despite the self-replication advantage, holding all the stocks in the benchmark may not always be desired or possible, mainly because of difficulties in purchasing odd lots to exactly match the market weights, or the increased transaction costs/market impact related to trading less liquid stocks. A more involved strategy is required when the replicating portfolio contains fewer stocks than the benchmark, or the latter is equally weighted. In the first case the weights will no longer be self-adjusting and will require periodic rebalancing, and in the second case frequent rebalancing is required in order to maintain equal dollar amounts in each stock.

We are concerned with the most general form of an indexing model, based on cointegration, which allows the replication of all types of indexes, with different numbers of stocks. The rationale for constructing portfolios based on a cointegration relationship with the market index, rather than correlation, rests on the following features of cointegration: the tracking error, defined as the price difference between the index and the replica portfolio is, by construction, stationary; the stock weights, being based on a large amount of history, have an enhanced stability; finally, there is a full use of the information contained in level variables such as stock prices. Moreover, cointegration relationships between the market index and portfolios comprising all or only part of their stocks should be easy to find since market indexes, either equally weighted or capitalisation weighted, are just linear combinations of stock prices.

The basic 'cointegration' model for a tracking portfolio comprising all the stocks included in the market index at a given moment is a regression of the form:

$$log(index_t) = c_1 + \sum_{k=1}^{n} c_{k+1} log(P_{k,t}) + \varepsilon_t \tag{1}$$

where the index is reconstructed historically based on the current membership of the market index, and n is the total number of stocks included in the market index. [iii]

All variables in the model, apart from the error term, are integrated of order one.[iv] The specification of the model in natural log variables has the advantage that, when taking the first difference, the expected returns on the portfolio will equal the expected returns on the market index, provided that the tracking error is a stationary process. [v]

We note that the application of ordinary least squares (OLS) to non-stationary dependent variables such as ln(index) is only valid in the special case of a cointegration relationship. The residuals in (1) are stationary if, and only if, ln(index) and the tracking portfolio are cointegrated. If the residuals from the above regressions are non-stationary, the OLS coefficient estimates will not be consistent and no further inference will be valid. Testing for cointegration is, therefore, essential in constructing cointegration optimal tracking portfolios. The Engle-Granger (1986) methodology for cointegration testing is particularly appealing in this respect for its intuitive and straightforward implementation. Moreover, its well-known limitations (small sample problems, asymmetry in treating the variables, at most one cointegration vector) are not effective in our case. The estimation sample is typically set to at least three years of daily data, there is a strong economic background to treat the market index as the dependent variable, and identifying only one cointegration vector is sufficient for our purposes. Further to estimation, the OLS coefficients in model (1) are normalised to sum up to one, thus providing the composition of the tracking portfolio.

Alexander and Dimitriu (2002) present an exhaustive analysis of the performance characteristics of cointegration optimal tracking portfolios, together with various other statistical arbitrage strategies derived from them, for different model parameters. The model parameters include the number of stocks in the portfolios, the stock selection method, the spread between the benchmarks tracked and the calibration period. The tracking portfolio comprising all stocks in the market index is shown to produce positive abnormal return in certain market conditions, even after accounting for transaction costs. This is a rather counterintuitive result, as one would expect the most complete combinations of stocks, very strongly cointegrated with the reconstructed index, to produce out of sample excess returns having zero mean.

To investigate whether this result can be replicated, we have constructed random subsets of stocks in the FTSE100, CAC40 and SP100 universes. For each index we have set up 100 random portfolios comprising a fixed number of stocks (50 for FTSE, 25 for CAC and 80 for SP100) and determined a price-weighted index for each portfolio. Each of the 300 indexes was tracked with a cointegration-optimal portfolio comprising all the stocks included in that particular index. The out-of-sample performance of portfolios was measured based on the following rebalancing method: every 10 trading days the optimal weights of the stocks are rebalanced based on the new OLS coefficients of the cointegration regression. For each re-balancing, the cointegration regression (1) is re-estimated over a fixed-length rolling calibration period of 3 years of daily data preceding the portfolio construction moment. The number of shares held in each stock is determined by the previous portfolio value, the current stock prices and the stock weights. In between re-balancings, the portfolios are left unmanaged, i.e. the number of stocks is kept constant.

Based on the rebalancing strategy detailed above, we have determined and reported in Table 1 the average annual abnormal return over the period 1997 to 2001, together with its standard deviation. For comparison, we have also reported the excess return in the real-world DJIA universe.

[*Insert Table 1 here*]

The first observation is that for the simulated indices in all four universes there is a positive average abnormal return, when measured over the entire data sample. Regarding the time distribution, the last years in the data sample are responsible for most of this excess return. This time variability is less evident in the CAC case, and most evident in the case of FTSE simulated indexes for which the abnormal return in 2000 is almost 5%, while in the case of SP100 simulated indexes, the largest abnormal return occurs during 2001, amounting to 4.6%. Given the consistency of this pattern, it is expected that the same type of out-performance will also occur in other stock markets.

Considering the above, we conclude that there is indeed evidence of abnormal returns from the cointegration-based indexing strategy, with a consistent time-variability pattern across different markets. This provides us with the motivation to investigate further the pattern of the abnormal return, restricting the analysis, for reasons of space, to the DJIA universe.

Using daily close prices for the thirty stocks in the DJIA as of 31-Dec-01 and a sample period from 01-Jan-90 to 31-Dec-01, we have estimated the out of sample performance of a portfolio constructed based on model (1) with a rolling 3-year calibration period and 10-day recalibration/rebalancing frequency. The cumulative daily abnormal return during the entire sample period is shown in Figure 1, and amounts to 11.6%, before transaction costs. The issue of the transaction costs is an important one, especially in connection with the EMH one needs to properly account for transaction costs. If we assume an amount of 20 basis points on each trade value to cover the bid-ask spread and the brokerage commissions, which is conservative for very liquid stocks such as the ones in DJIA, then the transaction costs estimated over the entire data sample sum up to no more than 2.5%. Such an amount of transaction costs can hardly be thought of affecting the overall performance of the strategy. The target of our analysis is explaining the 'pure' abnormal return from the strategy, before transaction costs. Therefore, after establishing that the overall profitability of the strategy does not disappear after transaction costs, we will perform the analysis on the abnormal returns before transaction costs.

A very noticeable feature in Figure 1 is the time variability of the excess return, which is far from being uniformly accumulated throughout the data sample. Periods of stationary excess returns alternate with periods during which there is consistently positive excess return. Moreover, the periods during which most of the abnormal return is accumulated coincide with the main market crises during the sample period: the Asian crisis, the Russian crisis and the technology market crash.

[Insert Figure 1 here]

The issue of time-variability in the performance of funds is not a new one. Generally, hedge funds and mutual funds have been found to perform better in recession periods than in boom periods. This time-variability has been associated with informational asymmetries (Shin, 2002) and changes in the investment environment, according to the phase of the business cycle (Moskowitz, 2000; Kosowski, 2001). A separate line of research concerns trend following strategies, which have been found to generate returns similar to a lookback straddle paying the owner the difference between the highest and the lowest price of the underlying asset over the observation period (Fung and Hsieh, 1997 and 2001). Trend followers appear to perform best in extreme up or down markets, and less well during calm markets. Even without highlighting the cause of this behaviour, there is considerable interest from the investment community in the fact that these funds provide a partial hedge against general market conditions (Fung and Hsieh, 1997).

However, in our case, it is a purely statistical strategy which generates this abnormal return. The over-performance of the cointegration portfolio must be connected to the portfolio weighting system and its relationship with stock price dynamics.[vi] In order to understand better the relation between the market index and the portfolio weights based on model (1), we investigate a theoretical example in the simplest case of an index comprising only 2 stocks, where the index, I_t, is computed as the average of the two stock prices. For constructing a portfolio tracking I_t from the two stocks, according to model (1) we need the weight w such that

$$log(I_t) = wlog(P_{1,t}) + (1 - w)log(P_{2,t}) \qquad (2)$$

It follows that $w = (log(1 + a) - log(2))/log(a)$ where $a = P_1/P_2$. Therefore $w > 0.5$ if and only if $P_1 > P_2$, meaning that in (2), the stock with the higher price will also have a higher weight in the portfolio. The difference between the portfolio weights in model (2) and the market weights, $w^* = a/(1 + a)$, will increase with the spread between the stock prices. The further away a is from unity, the larger the weight on the stock with the higher price in the portfolio, and, also, the larger the dispersion between the two stock prices. Therefore, a significant difference between the stock index weights and the portfolio weights will occur when the dispersion of stock prices increases.

This clear-cut result motivates our study of *index dispersion*, i.e. the cross sectional standard deviation of the prices across their mean (which is the reconstructed index), defined as:

$$d_t = \sqrt{\sum_{k=1}^{n} ((P_{k,t} - I_t)/I_t)^2 / n} \tag{3}$$

For computing the time series of index dispersion, all stock prices are rescaled to be equal to 100 at the beginning of the period, the dispersion series therefore starting from zero. In Figure 2 the bold line represents the time series of dispersion in the DJIA. After a steady increase, the dispersion increased substantially at the beginning of the technology sector boom, due to the sharp increase in the price of technology stocks, and a relative decline in price of other sectors. The highest dispersion occurred at the beginning of 2000, but since then the dispersion has decreased, most obviously during the crash of the technology bubble. We note that index dispersion in most major equity markets (whether capitalisation or equally weighted) follows a similar pattern.

2. A Basic Time Series Analysis

Both the cumulative abnormal return and index dispersion are found to be I(1) variables.[vii] Thus a basic stationary specification of their relationship will relate the (positive or negative) abnormal return (*AR*) to the daily change in index dispersion (*DD*), including also the lagged abnormal return and some lagged changes in dispersion:

$$AR_t = \alpha + \beta_1 AR_{t-1} + \beta_2 DD_t + \beta_3 DD_{t-1} + \beta_4 DD_{t-2} + \varepsilon_t \tag{4}$$

The simple regression estimation results based on the DJIA sample from Jan-92 to Dec-01 are presented in Table 2. Statistically significant coefficients are associated with the lag of the abnormal return and the first lag of the change in dispersion. The contemporaneous and the second lag of the change in dispersion are not statistically significant. The positive coefficient of the lagged abnormal return accounts for the autocorrelation in the abnormal return. Additionally, there is a negative, significant relationship between the abnormal return and the lagged change in dispersion. Thus following an increase in dispersion, there will be a relative loss in the portfolio compared with the market. The fact that the abnormal return is determined by the *lagged* change in dispersion rather than by a simultaneous variable indicates that dispersion may be a useful leading indicator of the performance of this strategy.

[Insert Table 2]

However, on further investigation the structural stability of this relationship seems questionable. A popular test for parameter instability is the Chow F-test, which, however, requires a-priori knowledge of the break date. A test which does not require knowledge of the potential break-point is the CUSUM test (Brown, Durbin and Evans, 1975), but this is known to have low asymptotic power (Ploberger and Kramer, 1990). A rolling version of the Chow test (Andrews, 1993), with the breakpoint set at different dates in the sample is only valid under the assumption of equal error variance in all the regressions. If there is heteroscedasticity in the restricted model, then the calculated F-statistic is biased upward and indicates greater instability in the coefficient estimates than in fact exists (Toyoda, 1974). A rolling Goldfeld-Quandt test (Goldfeld and Quandt, 1965) estimated for the period for the period Oct-92 to Oct-01 clearly rejects the null of errors homoscedasticity (Figure 3). Therefore, the standard Chow test cannot be used.

Provided that the sample is sufficiently large, one may use a Wald test that remains valid in the case of heteroscedastic errors. A rolling Wald test (Andrews and Fair, 1988), estimated for the period for the period Oct-92 to Oct-01 indicates that the null hypothesis of no-structural break is most significantly rejected on 16th October 2000 (Figure 4). Consequently, two separate regressions are estimated, using data before and after this date, and these have quite different results (Table 3). The main difference is on the sign of the coefficient of the lagged dispersion. The slope coefficient of the lagged change in dispersion is, until October 2000, negative, but after October 2000, the relationship between the two variables becomes positive. Additionally, when the impact of the change in dispersion is separated in the two samples, the lagged dependent variable becomes insignificant.

[*Insert Figure 4 and Table 3 here*]

Given that the weights in the portfolio are based on a long-run price equilibrium relationship, the negative relationship between the abnormal return and the lagged change in dispersion has a strong rationale. The dispersion can be interpreted of a measure of (dis)equilibrium – when prices diverge from long-run equilibrium levels, the dispersion in the entire stock prices system increases.

To illustrate this, we use the example of a higher than average priced stock. Figure 5(a) shows a smooth line representing the long-run equilibrium price of this stock and a wavy line representing the actual price of this stock. If the price of the stock increases, its weight in the market index will also increase (because the index is price-weighted). However, its weight in the portfolio, being based on a long history of prices, is not likely to react immediately to the increase in the stock price, which could be just noise from a long-run equilibrium perspective. Therefore, the portfolio will be relatively under-weighted on this particular stock while its price is increasing, and it will realise relative losses compared to the market index, during a period when dispersion is increasing. However, when the price of this stock returns towards its long-run equilibrium level, and consequently the dispersion in the system decreases, the portfolio will make a relative profit compared with the market index, because it is still under-weighted (relative to the index) on a stock whose price is declining.

[*Insert Figure 5 here*]

In this framework, the positive relationship between the abnormal return and the change in dispersion after October 2000 is rather puzzling. The only feasible explanation is that the cointegration relationship identifies new equilibrium prices, which are even further dispersed. This case is illustrated in Figure 5(b) where now the high price stock has a long-run equilibrium price that is above the actual stock price shown by the wavy line. The portfolio will realise a relative profit when the stock price is increasing because it is over-weighted (relative to the index) on the high-price stock. Similarly when the high price stock declines and the dispersion decreases, the portfolio, which is relatively over-weighted in this stock, will make a relative loss. Thus, when cointegration identifies a new equilibrium, in which the stock prices are even further dispersed, positive abnormal returns will be associated with increasing, not decreasing, dispersion.

Returning to the parameter stability test results, why should such a significant change in the behaviour of the abnormal return occur in October 2000? To answer this question, we take a closer look at the markets during the period September-December 2000.[viii] This three-month period is the time of the second great fall in the Nasdaq composite index. Index volatility reached 47.59% and the index fell 48.25%, i.e. another 745.83 points, having already fallen 425 points from March 2000. Therefore, it is reasonable to infer that October 2000 marked the end of the technology bubble.

To conclude, there is an obvious time-variability in the parameters of the estimated regressions, and this cannot be accounted for with simple tools like structural break tests without inducing a significant degree of arbitrariness. There appear to be some grounds for a structural break in the relationship between the abnormal return and dispersion in October 2000, but this does not ensure that the break identified is unique. To address these issues, we employ a Markov switching framework.

3. A Markov Switching Model

To specify and make further inference on the time-variability pattern identified in Figure 1, we have estimated a Markov switching model for the abnormal return.

Belonging to a very general class of time series models, which encompasses both non-linear and time-varying parameter models, the regime switching models provide a systematic approach to modelling multiple breaks and regime shifts in the data generating process. Increasingly, regime shifts are considered to be governed by exogenous stochastic processes, rather than being singular, deterministic events. When a time series is subject to regime shifts, the parameters of the statistical model will be time varying, but in a regime-switching model the process will be time-invariant conditional on a state variable that indicates the regime prevailing at the time.

The importance of these models has long been accepted, and the pioneering work of Hamilton (1989) has given rise to a huge research literature (Hansen, 1992 and 1996; Kim, 1994; Diebold, Lee and Weinbach, 1994; Garcia, 1998; Psaradakis and Sola, 1998). Hamilton (1989) provided the first formal statistical representation of the idea that economic recessions and expansions influence the behaviour of economic variables. He demonstrated that real output growth might follow one of two different auto-regressions, depending on whether the economy is expanding or contracting, with the shift between the two states generated by the outcome of an unobserved Markov chain.

In finance, the applications of Markov switching techniques have been many and very diverse: from modelling state dependent returns (Perez-Quiros and Timmermann, 2000) and volatility regimes (Hamilton and Lin, 1996), to option pricing (Aingworth, Das and Motwani, 2002), to detecting financial crises (Coe, 2002), bull and bear markets (Maheu and McCurdy, 2000) and periodically collapsing bubbles (Hall, Psaradakis and Sola, 1999), or to measuring mutual fund performance (Kosowski, 2001). Despite their limited forecasting abilities (Dacco and Satchell, 1988), Markov switching models have been successfully applied to constructing trading rules in equity markets (Hwang and Satchell, 1999), equity and bond markets (Brooks and Persand, 2001) and foreign exchange markets (Dueker and Neely, 2002).

The Markov switching model specified for the abnormal return assumes the presence of a latent variable (state variable), which determines the form of linear relationship between the abnormal return and the lagged dispersion in stock prices. The advantages of using a latent variable approach instead of a pre-defined indicator have been long documented. For example, when analysing business cycles, the Markov switching model produces estimates of the state conditional probabilities, which contain more precise information about the states that are driving the process than a simple binary indicator of the states, which is prone to significant measurement errors. The estimates of the conditional probability of each state allow more flexibility in modelling the switching process. An additional motivation for using a latent variable approach in this case, is the fact that there is no obvious indicator of the states of the process generating the abnormal return.

In the Markov switching model of abnormal return, the intercept, regression slope and the variance of the error terms are all assumed to be state-dependent. If we let s_t denote the latent state variable which can take one of K = 2 possible values (i.e. 1 or 2), then the regression model can be written as:

$$\mathbf{y}_t = \mathbf{z}_t' \beta_{S,t} + \varepsilon_{S,t} \tag{5}$$

where $\mathbf{y}_t$ is the (T x 1) vector of the abnormal returns; $\mathbf{z}_t = (\mathbf{1}\ \mathbf{x}_t)$ is the (T x 2) matrix of explanatory variables, with $\mathbf{x}_t$ denoting the lagged change in the prices dispersion; $\beta_{S,t} = (\gamma_{S,t}, \mu_{S,t})$ is the vector of state dependent regression coefficients; $\varepsilon_{S,t}$ is the vector of state dependent disturbances, assumed normal with state dependent variance $\sigma_{S,t}^2$.

The transition probabilities for the two states are assumed to follow a first-order Markov chain and to be constant over time:

$$P\{S_t = j \mid S_{t-1} = i, S_{t-2} = l, \ldots\} = P\{S_t = j \mid S_{t-1} = i\} = p_{ij}$$

The matrix of transition probabilities can be written:

$$\mathbf{P} = \begin{pmatrix} p_{11} & p_{21} \\ p_{12} & p_{22} \end{pmatrix} = \begin{pmatrix} p_{11} & 1-p_{22} \\ 1-p_{11} & p_{22} \end{pmatrix} = (p_{ij})$$

If we let ξ_t represent a Markov chain, with $\xi_t = (1, 0)'$ when $S_t = 1$ and $\xi_t = (0, 1)'$ when $S_t = 2$, then the conditional expectation of ξ_{t+1} given $S_t = i$ is given by:

$$E\,(\xi_{t+1} \mid S_t = i) = \begin{pmatrix} p_{i1} \\ p_{i2} \end{pmatrix} = \mathbf{P}\xi_t$$

The conditional densities of $\mathbf{y}_t$, assumed to be Gaussian, are collected in a 2x1 vector: $\eta_t = (\eta_{1t}, \eta_{2t})$ where $\eta_{it} = f(\mathbf{y}_t \mid S_t = i, \mathbf{z}_t ; \alpha)$ is the normal density function whose parameters α are conditional on the state. That is, $\eta_{it} = [(2\pi)^{1/2}\sigma_i]^{-1}\, exp\, (-(\mathbf{y}_t - \mathbf{z}_t' \beta_i)^2 / 2\, \sigma_i^{\,2})$.

The conditional state probabilities can be obtained recursively:

$$\hat{\xi}_{t|t} = \frac{\hat{\xi}_{t|t-1} \otimes \eta_t}{1'\,(\hat{\xi}_{t|t-1} \otimes \eta_t\,)} \tag{6}$$

$$\hat{\xi}_{t+1|t} = \mathbf{P}\hat{\xi}_{t|t}$$

where $\hat{\xi}_{t|t}$ represents the vector of conditional probabilities for each state estimated at time t, based on all the information available at time t, while $\hat{\xi}_{t+1|t}$ represents the forecast of the same conditional probabilities based on the information available at time t for time t+1. The symbol $\otimes$ denotes element-by-element multiplication.

The i$^{\text{th}}$ element of the product $\hat{\xi}_{t|t-1} \otimes \eta_t$ can be interpreted as the conditional joint distribution of $\mathbf{y}_t$ and $S_t = i$. The numerator in expression (6) represents the density of the observed vector $\mathbf{y}_t$ conditional on past observations.

Given the assumptions made on the conditional density of the disturbances, the log likelihood function can be written as:

$$L\,(\alpha, \mathbf{P}) = \sum_{t=1}^{T} log\; f\,(\mathbf{y}_t \mid \mathbf{z}_t\,; \alpha; \mathbf{P}\,) = \sum_{t=1}^{T} log\; \mathbf{1}'(\hat{\xi}_{t|t-1} \otimes \eta_t\,)$$

This approach allows the estimation of two sets of coefficients for the regression and variance of the residual terms, together with a set of transition probabilities.

Considering the complexity of the log likelihood function and the relatively high number of parameters to be estimated, the selection of starting values is critical for the convergence of the likelihood estimation. To reduce the risk of data mining, we have not used any state-dependent priors as starting values. Instead, we have used the unconditional estimates of the regression coefficients and the standard error of the residual term. Additionally, we have arbitrarily set $\xi_{1|1}$ to (1 0)'. A number of restrictions needed to be imposed on the coefficient values, in order to ensure their consistency with model assumptions. The transition probabilities were restricted to be between 0 and 1, while a non-negativity constraint was imposed on the standard deviation of the residuals in both states.

In Markov switching models it is essential to ensure a sufficiently long data sample for correctly identifying the time-variability of parameters. The data sample covered 10 years of daily data from 1992 to 2002. In a correctly specified switching model, i.e. one in which the entire time-variability of the parameters is captured by the regime switching and within each regime the parameter estimates are time-invariant, the use of such a long data sample should not create any difficulties.

[*Insert Table 4 here*]

Table 4 reports the Markov switching model estimation results over the entire data sample. The only coefficients statistically non-significant at 1% are the regression intercepts, for both states. A noteworthy difference between the two regimes concerns the coefficient of the lagged change in dispersion: in the first state, the coefficient is positive, while in the second state it is negative, thus the lagged change in dispersion has a different effect on the abnormal return in the two regimes. In the first regime an increase in the index dispersion is followed by a relative gain from index tracking in the next period, while in the second regime, a decrease in dispersion is the one associated with relative gains. Additionally, the standard deviation of the residuals is higher in the first regime, but as we show below, so is the abnormal return generated during this regime. Regarding the transition probabilities, the second regime appears to be more persistent than the first one: the probability of staying in regime two at time t+1 provided that at time t the process was in regime two is 0.98, while the probability of remaining in regime one, once there, is 0.88.

If we split the sample observations between the two regimes based on the criterion of estimated probability,[ix] we can determine the abnormal return associated to each regime. Based on this procedure, the number of observations in regime two is almost three times the number of observations in regime one. Also, the cumulative abnormal return generated during regime one turns out to be even higher than the abnormal return generated by the entire process, because the second regime generates a relative loss, even if not very significant.

Figure 6 shows the cumulative abnormal return from each regime: there is a consistent abnormal return produced in regime one, while the second regime produces a slightly negative abnormal return. Apart from higher returns, and higher volatility, regime one returns have a positive skewness (0.77) and higher excess kurtosis (2.99). Regime two, as well as having negative mean returns, also exhibits a negative skewness (-0.11), which indicates a higher probability of returns below the median, but relatively low excess kurtosis (0.11).

[Insert Figure 6]

Based on the same separation procedure as above, we observe that, as opposed to the replica portfolio, the market index generated smaller returns in regime one than in regime two (the equivalent of 8.75% p.a. as opposed to 12.93% p.a.). The notable difference concerns, however, the volatility of these returns: regime one returns are associated with an annual index volatility of 19%, while the returns in regime two have only 13% annual index volatility. Therefore, the tracking over-performance occurs in periods with lower returns and higher volatility for the market.

The time distribution of the states is an important feature to investigate. From Figure 7, which plots the estimated probability of the first regime, it becomes clear that in the first half of the sample regime two is the one prevailing, while towards the end of the sample, regime one becomes predominant. Over the entire data sample, observations in regime one represent 25% of the total number of observations. However, this distribution is far from being time invariant, since in the first half of the data sample, regime one accounts for only 7% of the total number of observations, while in the last two years of the data sample, i.e. 2000 and 2001, regime one occurs 87% of time.

Our main conclusion is that the two regimes have very distinctive characteristics: regime one, which occurs less frequently, but is predominant during the last few years, is responsible for producing the entire abnormal return. This regime occurs in more volatile market conditions and the over-performance follows an increase in the index dispersion. In the second regime there is a negative, but not significant excess return, with any positive excess return occurring further to a decrease of the index dispersion.

Considering the consistent abnormal return generated in regime one, and its relationship with the lagged change in dispersion, it follows that, in this regime, when the stock prices become more dispersed, the index tracking is in a relative profit position. This can occur if the portfolio is over-weighted on stocks having higher than average prices which are further increasing, and/or under-weighted on stocks having lower than average prices which are further decreasing. If these prices, after diverging, would return towards their previous levels, then the initial relative profit of the portfolio would be reversed, and there would not be any consistent abnormal return. However, since in this regime the observed abnormal return is consistent, it follows that after diverging, the prices are not returning to their previous levels. Instead, they are actually moving towards new equilibrium levels, levels that are pre-identified by the cointegration relationship and accounted for in the composition of the portfolio. Therefore, the portfolio is in a relative profit position during such transition periods, which, in regime one, are not likely to be reversed and, consequently, the profit to be eroded.

When there is an increase in dispersion in regime two, i.e. when the prices move away from equilibrium, the portfolio is in a relative loss position compared to the market index. However, since there is no consistent loss on the portfolio during regime two, it means that these stock price movements are only temporary disequilibria, and the cointegration is treating them accordingly. In regime two, the drifting of stock prices around their long-run equilibrium results in a stationary tracking error in the cointegration process.

Testing the null hypothesis of no-switching

In order to validate the above inferences about the two-state process driving the abnormal return, one needs to test and reject the null hypothesis of no switching. Even if there is evidence that the abnormal return has different patterns in the two regimes, this does not imply that the asymmetries between the two states are also statistically significant.

Standard testing methods such as likelihood ratio tests are not applicable to Markov switching models due to the presence of nuisance parameters under the null hypothesis of linearity, or no switching. The presence of nuisance parameters gives the likelihood surface sufficient freedom so that one cannot reject the null hypothesis of no switching, despite the fact that the parameters are apparently significant.

A formal test of the Markov switching models against the linear alternative of no-switching, which is designed to produce valid inference, has been proposed by Hansen (1992, 1996).

This method implies the evaluation of the log likelihood function for a grid of different values for the regression coefficients, standard deviation and the transition probabilities.

Following Hamilton (1996), we let $\alpha = (\mu_1-\mu_2, \gamma_1-\gamma_2, \sigma_1-\sigma_2, p_{11}, p_{22})'$ denote the regime switching parameters of model and $\lambda = (\mu_1, \gamma_1, \sigma_1)'$ denote the parameters which are not state dependent. The conditional log likelihood function for the parameters will be written as $L_t(\alpha, \lambda) = \log f(y_t \mid y_{t-1}, y_{t-2}, \ldots y_1; \alpha, \lambda)$.

The null hypothesis of no switching can be written as $\alpha = \alpha_0 = (0, 0, 1, 0)'$. To represent the alternative hypothesis, we have constructed a grid of 1,125 possible values for α, with A denoting the set comprising all values of α. For any α, $\hat{\lambda}(\alpha)$ denotes the value of λ that maximises the likelihood taking α as given. Hamilton (1996) defines the time series of the difference between each constraint log-likelihood function for the grid of alternatives and the constraint log-likelihood function estimated for the null hypothesis as:

$$q_t(\alpha) = l_t[\alpha, \hat{\lambda}(\alpha)] - l_t[\alpha_0, \hat{\lambda}(\alpha_0)]$$

The likelihood ratio statistic is:

$$L\hat{R} = \max_{\alpha \in A} \frac{T\bar{q}(\alpha)}{\sqrt{\sum_{t=1}^{T} [q_t(\alpha) - \bar{q}(\alpha)]^2}}$$

If the null hypothesis is true, then, for large samples, the probability that the above statistic exceeds a critical value z is less than the probability that the following statistic exceeds the same value z:

$$\max_{\alpha \in A} \frac{\sum_{k=0}^{M} \sum_{t=1}^{T} [q_t(\alpha) - \bar{q}(\alpha)] u_{t+k}}{\sqrt{(1+M)\sum_{t=1}^{T} [q_t(\alpha) - \bar{q}(\alpha)]^2}}$$

Following Hamilton (1996), we have generated Hansen's statistic for M values of 0-4 and found that the null hypothesis is strongly rejected with a p-value of 0.0000. The estimated Hansen statistic is of 5.58, while the upper bound of the simulated distribution is 2.82.

An alternative approach to the Hansen statistic uses a classical log likelihood ratio test for estimating (a) the asymmetries in the conditional mean, assuming the existence of two states in the conditional volatility, and (b) the asymmetries in the conditional volatility, assuming the existence of two states in the conditional mean. Such a test follows the standard chi-squared distribution. Table 5 reports the log likelihood estimates for the restricted and unrestricted models, the log likelihood ratios and the associated probability values.

[*Insert Table 5 here*]

We have tested the following hypotheses: (1) the intercept and slope coefficients are not significantly different between the two states, and (2) the standard deviations of the residuals of the two states are not significantly different. As shown by the results in Table 5, both tests turned out to be statistically significant, and the null hypotheses were rejected. Therefore, we conclude that there is clear evidence of the fact that the asymmetries between the two regimes identified by the model are not only economically, but also statistically significant.

4. Trading Rules based Tests of the Model's Predictive Power

In this section, in order to test the out of sample predictive power of the Markov switching model of index dispersion, we propose two market neutral strategies which exploit the regime dependent relationship between the index tracking out-performance and the stock prices dispersion. Their construction is based on the fact that the *lag* of the change in dispersion is used to explain the abnormal return, and, therefore, we have a leading indicator of portfolio performance. Also, forecasts of the latent state conditional probability can be produced for a number of steps ahead by using the unconditional transition probabilities and the current estimate of the conditional probabilities of the latent states.

The portfolio generating the abnormal relative to the index, *P*, is defined as the difference between the replica portfolio holdings and the market holdings in each stock. Both trading rules assume active trading, with daily rebalancing according to a trading signal. The first trading rule ensures that *P* is held only if there is a buy/hold signal from the Markov switching model. In the second strategy, *P* is held if there is a buy/hold signal, and is shorted otherwise. The 'buy/hold' signal occurs either after an increase in the dispersion, if the forecast of the conditional probability of the latent state indicates that the process is currently in regime one, or after a decrease in dispersion, if the forecast of the conditional probability indicates that the process is currently in regime two. As the abnormal return is not correlated with the market returns, both strategies will inherit market neutral characteristics. Moreover, they are self-financed, as the sum of all stock weights in *P* is, by construction, zero.

For pure out-of-sample tests, we have extended our initial database up to Nov-02, and the trading rules were implemented for the period Dec-01 to Nov-02. In order to obtain the signal for a given date, we have used only the information available at the moment of the signal estimation (the sign of the lagged change in dispersion and the one-period ahead forecast of the conditional probability of the regimes).

The returns of the trading strategies are plotted in Figure 8. Over the 11-months testing period, the first trading rule produced a cumulative return of 3.9%, with an average annual volatility of 2.2%. This translates into an average annual information ratio of 1.89, again, for a self-financing strategy. The second trading rule produced over the same time interval a cumulative return of 9.3%, with a slightly higher average annual volatility, i.e. 3.2% p.a. The average annual information ratio for this strategy is of 3.15.

However, these results need to be interpreted with caution. First, 11 months is a rather short sample, and secondly, the very high profitability of the trading rules during the last part of the data sample can be the result of the predominance and persistence of regime two during this period. Therefore, to estimate the impact of these limitations, we have also performed in-sample tests, for the period Dec-91 to Dec-01. The results are very similar: over the 10-year testing period, the first trading rule produced a cumulative return of 51.5%, with an average annual volatility of 1.9%. This translates into an average annual information ratio of 2.65, again, for a self-financing strategy. The second trading rule produced over the same time interval a cumulative return of 91.9%, with a slightly higher average annual volatility, i.e. 2.5% p.a. The average annual information ratio for this strategy is of 3.56.

However, the trading rules, as they are designed, require daily rebalancing, which can result in significant transaction costs. In our analysis we have not accounted for potential transaction costs, as we only aimed to test the efficiency of the model forecasts with real trading rules rather than with statistical tools. The problem of potentially high transaction costs in trading rules based on Markov switching forecasts is not new and has been dealt with either by reducing the frequency of trades, or by imposing some filtering of the signals, when trades take place only if the signal exceeds a given threshold (Dueker and Neely, 2001).

5. Implications for Market Efficiency

Turning now to the implications of our findings for market efficiency, we have provided evidence of consistent abnormal return, even after transaction costs, generated through the cointegration-based tracking model. Considering that the strategy is constructed solely on the information contained in past prices, the abnormal return alone qualifies as evidence against EMH in the weak form[x], and this happens in some of the most liquid stock universes in the world: DJIA, SP100, FTSE100 and CAC40.

Moreover, we have found a leading indicator for the abnormal return, which was shown to be predictable, even on a short time horizon, with the Markov switching model estimated in the previous sections.

However, the time-variability identified for the abnormal return, and its strict association with only one of the market regimes, indicate that the market inefficiencies exploited by the cointegration model are temporary and occur in only particular market circumstances, e.g. when the market returns are lower than average and the index volatility is higher than average. We have shown that such circumstances can indicate a transitional period in the market, where stock prices are moving towards new equilibrium levels.

Therefore, in the risk/return context, the abnormal return over the index can be seen as a reward for bearing, during transition periods, the uncertainty of new equilibrium price levels. As the cointegration model is successful in predicting the correct equilibrium levels, it earns the associated risk premium. As long as the prices converge towards the new equilibrium levels identified by the model, either further dispersed or closer to each other, the portfolio will generate abnormal return, as was the case in our sample.

Therefore, the predictive power of the cointegration model in detecting transitions towards new equilibrium prices and the success of the strategy in exploiting the information enclosed in the past stock prices, even after transaction costs, provide evidence for temporary market inefficiencies.

6. Summary and Conclusions

The aim of this paper was to investigate the abnormal return generated through a dynamic indexing strategy and to analyse its implications from a market efficiency perspective. We have introduced a new measure of stock price dispersion and found that this is a leading indicator for the abnormal return. Their relationship was modelled as a switching process between two stock market regimes having very distinctive characteristics, and we have found that almost the entire abnormal return may be associated with only one of them. This regime, prevailing during the last few years of our data sample, is characterised by higher index volatility and lower returns.

The cointegration relationship specified between the portfolio and the market index can be interpreted as a relative pricing model. For as long as stock prices are oscillating around the past equilibrium levels, the strategy generates accurate replicas of the market index. But the model also predicts transitions towards new equilibrium levels well in advance, and exploits them by producing consistent excess returns.

The predictive power of the cointegration model in detecting transitions towards new equilibrium prices and the success of the strategy in exploiting the information in the past stock prices were demonstrated over different time horizons and in different, real-world and simulated stock markets. The strategy remained profitable even after introducing transaction costs, thus proving evidence of temporary market inefficiencies.

Our findings have wide implications for the passive investment industry. We have shown that, without any stock selection or explicit timing attempts, solely through smart optimisation, the benchmark performance can be significantly enhanced. Moreover, the strategy can be implemented to replicate any type of value or capitalisation weighted benchmark, not only wide market indexes.

References

Aingworth, D., S.R. Das and R. Motwani. (2002). "Pricing Equity Options with Markov Switching State Variables." working paper, http://scumis.scu.edu/~srdas/research.htm.

Alexander, C. (1999). "Optimal Hedging Using Cointegration." Philosophical Transactions of the Royal Society, A 357, 2039-2058.

Alexander, C. and A. Dimitriu. (2002). "The Cointegration Alpha: Enhanced Index Tracking and Long-Short Market Neutral Strategies." discussion paper 2002-08, ISMA Discussion Paper Series in Finance.

Alexander, S. (1964). "Price Movements in Speculative Markets: Trends or Random Walks, No. 2", Industrial Management Review 5, 25-46.

Andrews, D. (1993). "Tests for Parameter Instability and Structural Change with Unknown Change Point." Econometrica 61, 821-856.

Andrews, D. and R. Fair. (1988). "Inference is Nonlinear Econometric Models with Structural Change." Review of Economic Studies 60, 615-639.

Blake, R. (2002). "Is Time Running Out For The S&P 500?" Institutional Investor Magazine, Americas Edition May 2002.

Brooks, C. and G. Persand. (2001). "The Trading Profitability of Forecasts of the Gilt-Equity Yield Ratio." Journal of International Forecasting 17, 11-29.

Brown, R.L., J. Durbin and J.M. Evans. (1975). "Techniques for Testing the Constancy of Regression Relationships Over Time." Journal of the Royal Statistical Society 37, 149-192.

Carhart, M. (1997). "On Persistence in Mutual Fund Performance." Journal of Finance 52, 57-82.

Coe, P. (2002). "Financial Crisis and the Great Depression: a Regime Switching Approach." Journal of Money, Credit and Banking 34, 76-93.

Dacco, R., and S. Satchell. (1999). "Why do Regime Switching Models forecast so badly?" Journal of Forecasting 18, 1-16.

Dickey, D.A., and W.A. Fuller. (1981). "Likelihood Ratio Statistics for Autoregressive Time Series with a Unit Root." Econometrica 49, 1057-1072.

Duecker, M., and C.J. Neely. (2001). "Can Markov Switching Models Predict Excess Foreign Exchange Returns?" working paper 2001-021B, Federal Reserve Bank of St. Louis.

Engle, R.F., and C.W.J. Granger. (1987). "Co-integration and Error Correction: Representation, Estimation and Testing." Econometrica 552, 251-276.

Diebold, F.X., J.H. Lee and G.C. Weinbach. (1994). "Regime Switching with Time-Varying Transition Probabilities." Oxford University Press, 283-302.

Elton, E.J., M.J. Gruber, S. Das and M. Hlavka. (1993). "Efficiency with Costly Information: A Reinterpretation of Evidence from Managed Portfolios." Review of Financial Studies 1, 1-22.

Fama, E.F., (1970). “Efficient Capital Markets: A Review of Theory and Empirical Work.” Journal of Finance 25, 383-417.

Fama, E., and M. Blume. (1966). “Filter Rules and Stock Market Trading.” Journal of Business 39, Issue 1, 226-241.

Fama, E.F., L. Fisher, M.C. Jensen and R. Roll. (1969). “The Adjustment of Stock Prices to New Information.” International Economic Review 10, 1-21.

Fung, W., and D.A. Hsieh. (1997). “Empirical Characteristics of Dynamic Trading Strategies: The Case of Hedge Funds.” Review of Financial Studies 10, 275-302.

Fung, W., and D.A. Hsieh. (2001). “The Risk in Hedge Fund Strategies: Theory and Evidence from Trend Followers.” Review of Financial Studies 14, 313-341.

Garcia, R., (1998). “Asymptotic Null Distribution of the Likelihood Ratio Test in Markov Switching Models.” International Economic Review 39, 763-788.

Goldfeld, S.M., and R.E. Quandt. (1965). “Some Tests for Homoskedasticity.” Journal of the American Statistical Association 60, 539-547.

Hall, S.G., Z. Psaradakis and M. Sola. (1999). “Detecting Periodically Collapsing Bubbles: A Markov Switching Unit Root Test.” Journal of Applied Econometrics 14, 143-154.

Hamilton, J.D., (1989). “A New Approach to the Economic Analysis of Nonstationary Time Series and the Business Cycle.” Econometrica 57, 357-384.

Hamilton, J.D., (1994). “Modeling Time Series with Changes in Regime, in Time Series Analysis.” Princeton, Princeton University Press, 677-703.

Hamilton, J.D., and G. Lin. (1996). “Stock Market Volatility and the Business Cycle.” Journal of Applied Econometrics, Volume 11, Issue 5, 573-593.

Hansen, B.E., (1992). “The Likelihood Ratio Test under Non-standard Conditions: Testing the Markov Switching Model of GNP.” Journal of Applied Econometrics 7, S61-S82.

Hansen, B.E., (1992). “Erratum: The Likelihood Ratio Test under Non-standard Conditions: Testing the Markov Switching Model of GNP.” Journal of Applied Econometrics 11, 195-199

Hendry, D., and K. Juselius. (2000). “Explaining Cointegration Analysis: Part II.” Energy Journal 21, 1-42.

Hwang, S., and S. Satchell. (2001), “Tracking Error: Ex-Ante versus Ex-Post Measures.” Journal of Asset Management Vol. 2, No. 3, 241-246.

Hwang, S., and S. Satchell. (2001), “The Disappearance of Style in US Equity Market.” Financial Econometrics Research Centre Working paper, Cass Business School.

Jensen, M.C., (1968). “The Performance of Mutual Funds in the Period 1955-1964”, Journal of Finance 23, 389-416.

Jensen, M.C., (1978). “Some anomalous evidence regarding market efficiency.” Journal of Financial Economics 6, 95-101.

Kim, C.J., (1994). "Dynamic Linear Models with Markov switching." Journal of Econometrics 60, 1-22.

Kosowski, R., (2001). "Do mutual funds perform when it matters most to investors? US mutual fund performance and risk in recessions and booms 1962-1994." working paper Financial Market Group, London School of Economics.

Lakonishok, J., and T. Vermaelen. (1990). "Anomalous Price Behavior Around Repurchase Tender Offers." Journal of Finance 45, No. 2, 455-477.

Larsen, G.A., and B. Resnick. (1998). "Empirical Insights on Indexing." Journal of Portfolio Management Vol. 25, No. 1, 51-60.

MacKinnon, J., (1991). "Critical Values for the Cointegration Tests". in R.F. Engle and C.W.J. Granger (eds.) 1991. Long-run Economic Relationships. Oxford University Press, 267-276.

Maheu, J.M., and T.H. McCurdy. (2000). "Identifying Bull and Bear Markets in Stock Returns." Journal of Business and Economic Statistics 18-1, 100-112.

Moskowitz, T.J., (2000). "Discussion: Mutual fund performance: An empirical decomposition into stock-picking talent, style, transaction costs, and expenses." Journal of Finance 55, 1655-1703.

Perez-Quiros, G., and A. Timmermann. (2000). "Firm size and cyclical variation in stock returns." Journal of Finance 50, 1229-1262.

Ploberger, W., and W. Kramer. (1990). "The Local Power of the CUSUM and CUSUM of Squares Tests." Econometric Theory 6, 335-347.

Psaradakis, Z., and M. Sola. (1998). "Finite-Sample Properties of the Maximum Likelihood Estimator in Autoregressive Models with Markov Switching." Journal of Econometrics 86, 369-386.

Shin, H.S., (2001). "Disclosures and asset prices." working paper 0371 Financial Market Group, London School of Economics.

Toyoda, T., (1974). "Use of the Chow Test under Heteroscedasticity." Econometrica 42, 601-608

Figure 1 Cumulative abnormal return

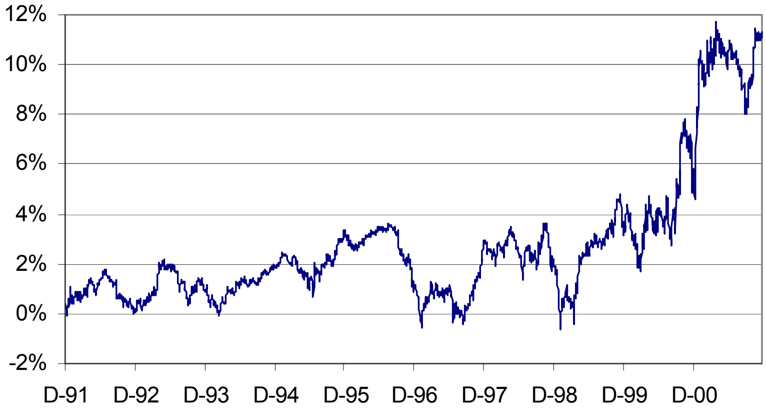

Figure 2 Price dispersion in DJIA and simulated indices

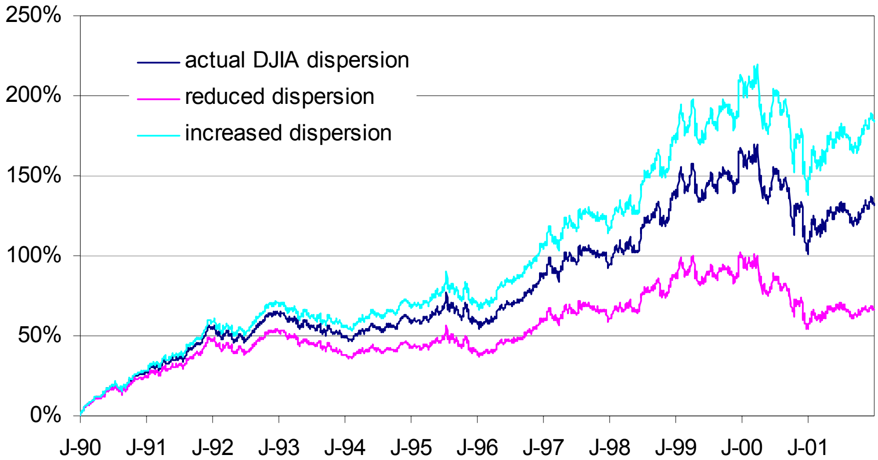

Figure 3 F-statistic Goldfeld-Quandt homoscedasticity test

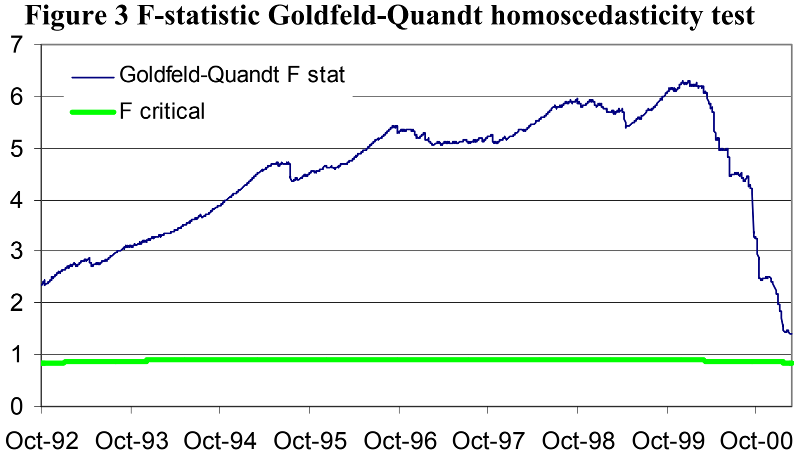

Figure 4 Wald test for parameter stability

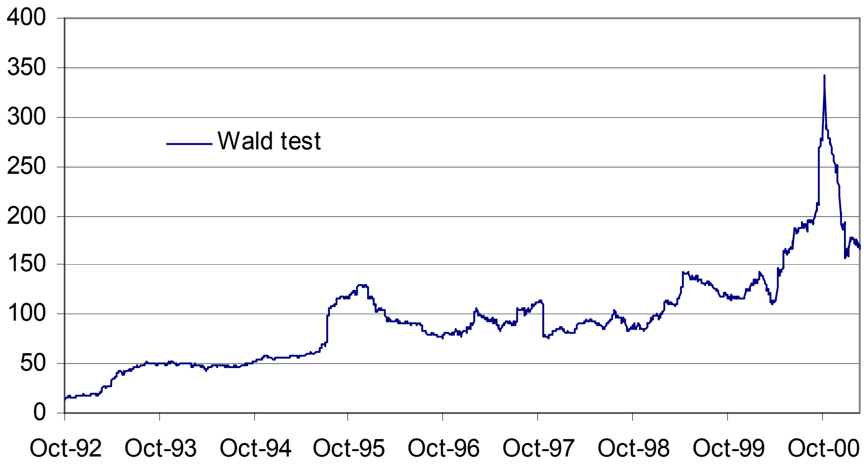

Figure 5(a) Stock prices movements in regime two

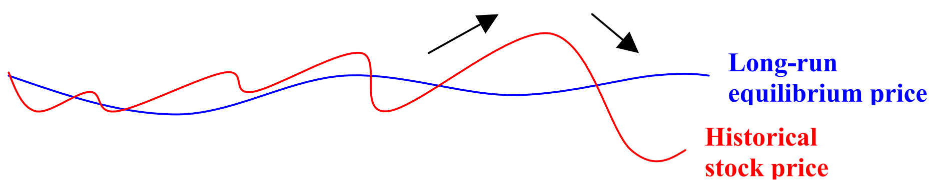

Figure 5(b) Stock prices movements in regime one

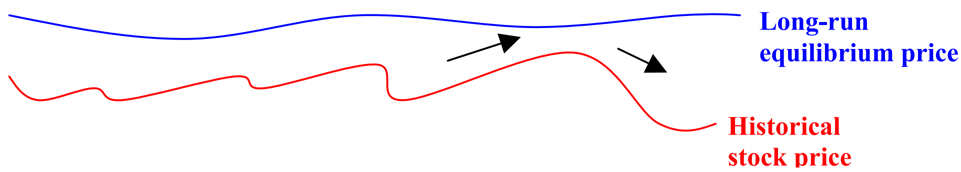

Figure 6 Regime conditional cumulative abnormal return

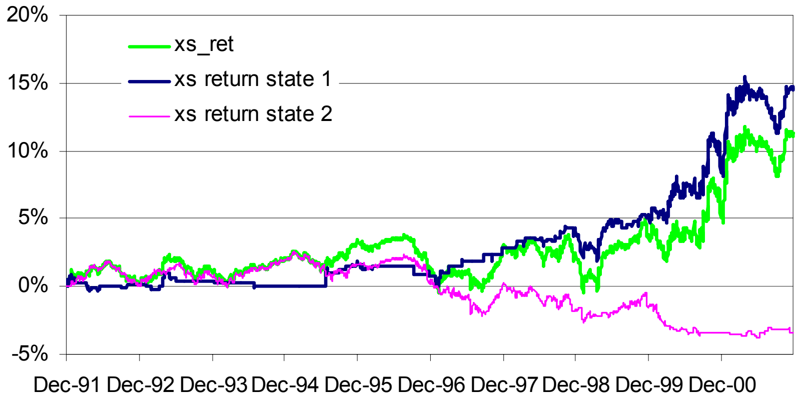

	Regime 1	Regime 2
Mean	2.32E-04	-1.80E-05
Max	0.0154	0.0036
Min	-0.0094	-0.0038
Stdev	0.0028	0.0009
Skewness	0.77	-0.11
Kurtosis	5.99	3.84
No of obs	632	1888

Figure 7 Estimated probability of regime one

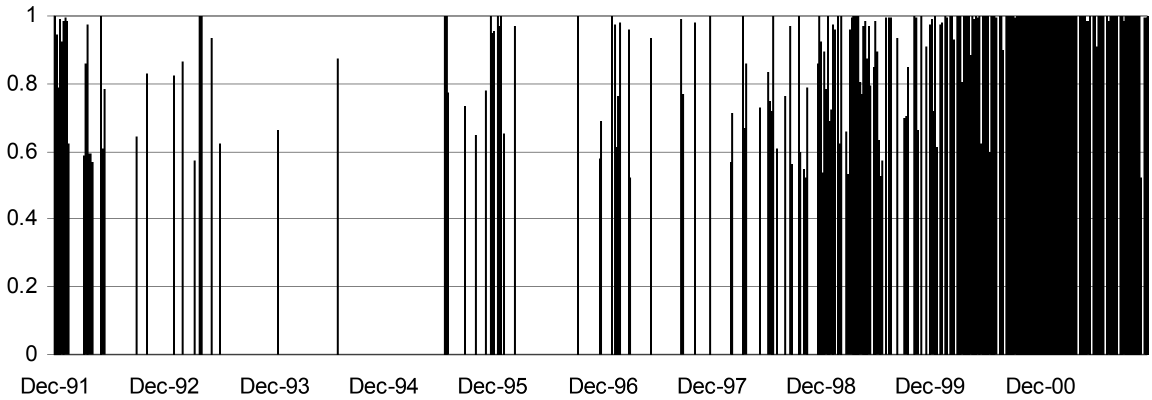

Figure 8 Cumulative returns produced by the trading rules

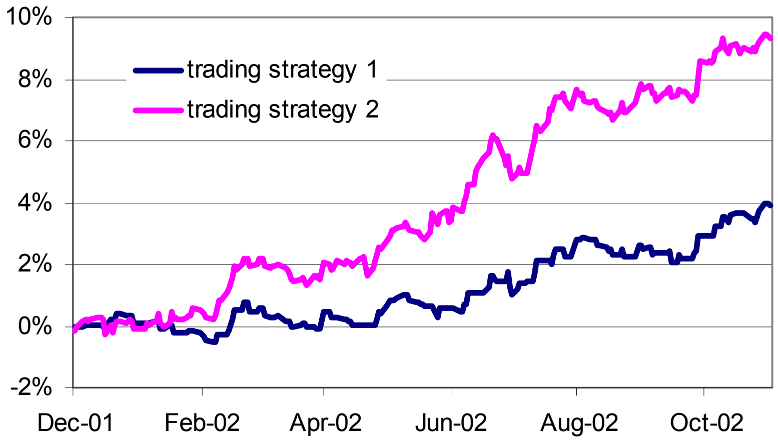

Table 1. A. Abnormal return

		1997	1998	1999	2000	2001	overall
DJIA		1.71%	-1.00%	1.46%	2.08%	5.61%	7.82%
FTSE simulated	Mean	-0.18%	-0.64%	0.92%	4.77%	0.34%	5.21%
indexes	Stdev	0.0055	0.0101	0.0136	0.0137	0.0114	0.0227
CAC simulated	Mean	1.19%	1.16%	-0.04%	1.22%	-0.13%	3.41%
indexes	Stdev	0.0045	0.0045	0.0071	0.0123	0.0100	0.0232
SP100 simulated	Mean	0.73%	-1.68%	-1.84%	1.00%	4.59%	2.79%
indexes	Stdev	0.0033	0.0071	0.0125	0.0173	0.0139	0.0280

Table 1. B. Index returns

		1997	1998	1999	2000	2001
DJIA		20.41%	14.93%	22.49%	-6.37%	-7.37%
FTSE simulated	Mean	19.86%	8.24%	8.35%	7.32%	-11.71%
CAC simulated	Mean	8.21%	27.26%	38.44%	0.45%	-21.79%
SP100 simulated	Mean	26%	19%	18%	1%	-13%

Table 2 Estimated coefficients of model (4)

	α	β_1	β_2	β_3	β_4
Coefficient	4.43E-05	0.075924	0.001347	-0.020175	0.005545
Standard error	3.25E-05	0.019892	0.002362	0.002361	0.002393
t-statistic	1.362186	3.816796	0.570434	-8.543316	2.317162
P-value	0.1733	0.0001	0.5684	0.0000	0.0206

Table 3 Stability test for model (4)

Wald Stability Test: 2227 (October 16, 2000)			
F-statistic	341.2555	Probability	0.000000

Sample Jan-92 to Oct-00			
	α	β_1	β_3
Coefficient	3.92E-0.5	0.0025	-0.0567
Standard error	2.45E-0.5	0.0179	0.0018
t-statistic	1.60	0.14	-29.92
P-value	0.1094	0.8880	0.0000

Sample Oct-00 to Dec-01			
	α	β_1	β_2
Coefficient	1.30E-0.4	0.0599	0.1068
Standard error	1.33E-0.4	0.0435	0.0070
t-statistic	0.97	1.37	15.25
P-value	0.3313	0.1695	0.000

Table 4 Estimation output for model (5)

	μ_1	μ_2	γ_1	γ_2	σ_1	σ_2	p_{11}	p_{22}
Coefficient	2.68E-04	2.63E-05	0.0163	-0.056	0.0029	0.0006	0.88	0.98
Standard error	1.25E-04	1.48E-05	2.82E-03	1.14E-03	1.02E-05	1.58E-06	0.087	0.073
Z-statistic	2.15	1.78	5.77	-49.01	-286.01	-403.08	10.03	13.35
P-value	0.031	0.074	0.000	0.000	0.000	0.000	0.000	0.000

Table 5 Log likelihood ratio tests for identical mean and volatility

	Unrestricted log likelihood	Restricted log likelihood	LR statistic	P-value
H_0: $\mu_1 - \mu_2 = 0$ $\gamma_1 - \gamma_2 = 0$	13707.64	13666.05	82.24	0.0000
H_0: $\sigma_1 - \sigma_2 = 0$	13707.64	13057.31	1300.66	0.0000

Footnotes

[ii] Portfolio in which all the assets available to investors are held in proportion to their market value.

[ii] We define the abnormal return as the difference between the strategy return and the return on a price-weighted index comprising the same stocks. In this paper, we use the terms 'abnormal return' and 'excess return' alternatively, to denote the same concept.
[iii] That is, the benchmark is reconstructed as a price weighted index of the stocks currently in the DJIA, and this represents a scale adjustment to the DJIA based on the value of the latest index divisor.

[iv] Stock prices and market indexes are usually found to be integrated of order one. The preliminary analysis of our data showed that this is also our case.

[v] Hendry and Juselius (2000) show that if level variables are cointegrated, so will be their logarithms. The level variables are cointegrated by definition, since the current weighted index is a linear combination of the stock prices.

[vi] Although the periods with abnormal return are associated with market downturns over a long time horizon, on a daily basis there is no significant negative correlation with the market returns, or with the market volatility. Additionally, the 10-day no-rebalancing period is not explaining the abnormal return.

[vii] Considering the ADF statistics for the dispersion, respectively cumulative excess return, of –0.81 and 0.18, we cannot reject the null hypothesis of a unit root in the series. However, the ADF tests on the first difference of the series clearly reject the null hypothesis of unit root (ADF statistics of –22.11 and -23.11) at any standard significance level.

[viii] We examine a three-month period rather than only October 16th because the exact date indicated by the tests as having the highest likelihood of a structural break can be an artefact of the estimation method used.

[ix] If the estimated conditional probability of regime one at time t is above 0.5, we say that the process was in regime one at time t. Alternatively, the process will be in regime two.

[x] The weak form of market efficiency assumes that prices fully reflect at all times the information comprised in past prices, as opposed to the other two forms of market efficiency, semi-strong and strong, which assume that prices also reflect all other public information, respectively, all other public and private information.