EXAMINING THE FIRST STAGES OF MARKET PERFORMANCE.
A TEST FOR EVOLVING MARKET EFFICIENCY

Anna Zalewska-Mitura
Centre for Economic Forecasting
London Business School

and

Stephen G. Hall
The School of Management
Imperial College
and
Centre for Economic Forecasting
London Business School

April 1998

The purpose of this paper is to provide an econometric tool for testing how markets evolve and become more informationally efficient. The process of ``becoming more efficient'' is important in the context of rapidly growing emerging markets, as well as developed markets which are continuously changing their regulations and structures. To test how efficiency changes, rather than to test whether a particular market is efficient or not, we propose a multi-factor model with time varying coefficients and generalised auto-regressive conditional heteroscedastic (GARCH) errors. We provide results of Monte Carlo simulations to show the usefulness of the proposed test. An illustration of its application is discussed using series available from the London Stock Exchange and one of the biggest Central European emerging markets - the Budapest Stock Exchange.

CONTENTS

1. Introduction

Since the 1970's public attention to emerging markets has grown. This is related to the growth of the number of emerging markets, as well as to the increased importance of these markets on the international financial scene. Although many of the markets are commonly named ``emerging'', differences among them are tremendous. The markets differ in capitalisation, the size of undertaken transactions, the size of involved foreign capital and organisational schemes, and most significantly, the emerging markets differ in age. Some emerging financial markets are relatively old (e.g. The Kuala Lumpur Stock Exchange and the Singapore Stock Exchange were formally established as separate markets in 1973, but their common history can be traced as far back to 1910). In contrast, many emerging markets in Africa appeared in the last decade. Another group of emerging markets appeared in Europe after the collapse of the communist regime.

When we look at emerging financial markets the question of their development appears. We agree with the classical approach, which states that: "perfect" financial market should be efficient. It seems natural to understand the "level of efficiency" as the "level of development".

In this paper we concentrate on the problem of how to measure possible changes in efficiency over time. The paper is organised as follows: Section 2 briefly discusses traditional approaches for testing for weak-form efficiency and introduces a Test for Evolving Efficiency (TEE). Section 3 tests for the power of the TEE using Monte Carlo simulations. Furthermore, to give a more direct illustration of the TEE and to see differences between developed and emerging markets some share price series and market indexes are discussed. In Section 4 the TEE is applied to the market index and a share price series from the London Stock Exchange. In Section 5, the TEE is used to examine the market index and some share price series from one of the leading Central European emerging markets - the Budapest Stock Exchange. Section 6 contains conclusions and the final discussion of the results.

2. The Concept of Market Efficiency

In this section we briefly discuss classical approaches for testing for market efficiency and define a new test allowing us to observe possible changes in efficiency.

2.1. The Classical Approach

Since the pioneering theoretical work by Bachelier (1900) on the Efficient Market Hypothesis, a large number of empirical and theoretical studies examining market behaviour and testing efficiency of developed equity markets have appeared. The concept of *efficiency* has changed over time. A great contributor to the modern definition was Roberts (1967) who distinguished among week, semistrong, and strong form of efficiency, according to an explicitly defined informational set which could be considered by earning superior returns.

Summarising the methodology of empirical studies, tests which have been carried out can be divided into two groups: tests of asset pricing models and tests of market efficiency. The first group of tests seeks to determine whether share price behaviour conforms to the Capital Asset Pricing Model. The second group of tests seeks to determine whether the pattern of changes in the

1

share price of common stocks is consistent with the notion of an efficient equity market, i.e. whether the market is informationally and allocationally efficient (Fama (1970, 1976), and Beaver (1981)). The distinction between tests of CAPM and tests of market efficiency is made only for the purpose of exposition. Tests of pricing models, in some sense, are also tests of market efficiency, and tests of market efficiency are usually carried out assuming a model of asset pricing.

Although, the tests for market efficiency can differ from a technical point of view, they have one common trait - they look at the long-run market characteristics. For instance, an examination of autocorrelation coefficients, variance ratios, or testing for long-range dependence allows us to measure the market predictability as a property steady over some predefined period. For instance, Theodore (1996) studied trading patterns in developed and developing markets for the years 1988-94. The seven year period seems rather short from the point of view of a market which has existed for some hundred years. The same period of time can be the whole life for an emerging exchange. It is logical to expect a market to behave differently on the day of its opening, and during the immediate period following, than some years later, when market participants have become more experienced and the market system better developed. If a market performs and stabilises (matures) over time we should observe how introductory inefficiency (if there was any) vanishes and market becomes more efficient.

More precisely, when we describe the classical definition of market efficiency, we say that a market is *weak-form efficient* when there is no profit opportunity based on the past movement of asset prices. It can be tested with the equation

$$r_t = \beta_0 + \sum_{i=1}^{p} \beta_i r_{t-i} + e_t, \ e_t \sim NIID(0,h) \tag{2.1}$$

where r_t is the rate of return on the asset at time t. If the market is efficient all the $\beta_i, i > 0$ coefficients should be equal to zero. In the financial literature is well documented that the error process of financial securities' returns does not often prove to have a full set of NIID properties. In particular, if the variance of the error process is changing over time in a systematic way, this will cause problems for the testing procedure and it may also affect the required rate of return. If the changing variance structure is omitted and there is a serial correlation, then again we may find spurious correlation and thus incorrectly reject market efficiency. This case can be dealt with by combining equation (2.1) with a standard GARCH-M model, i.e. by adding the volatility term to equation (2.1) in the following way:

$$r_t = \beta_0 + \sum_{i=1}^{p} \beta_i r_{t-i} + \delta h_t + e_t, \ e_t \sim N(0,h_t) \tag{2.2}$$

$$h_t = \alpha_0 + \alpha_i h_{t-1} + \alpha_2 e_{t-1}^2 \tag{2.3}$$

One might consider the question of how to modify this method to get information about possible changes in weak-form efficiency over time. For instance, one might try to use a technique of

"slowly" moving windows of a certain length and in this way divide the whole period in question into a finite number of sub-periods. Then one might estimate all the parameters required by equations (2.2)-(2.3) for each sub-period individually. In this approach, however, the length of a window is crucial and significant for the power of the estimated values. Moreover, we should note that such an approach causes the loss of some observations at the beginning and at the end of the investigated period.

We follow another approach which is, in our opinion, more self-consistent. In the next subsection we define a new test in which time-varying β_i coefficients are introduced.

2.2 The Test for Evolving Efficiency (TEE)

To test for possible changes in market efficiency we need a tool which will allow us to do the following: first, it will allow us to, at least, check for weak form efficiency; second, it will be able to detect changes in efficiency over time; and third, it will operate within stochastic series for which the error process can have the variance changing over time in a systematic way, i.e., the error process does not prove to have a full set of NIID properties. To attain all these requirements we suggest modifying equations (2.2)-(2.3) by adding a time subscript to the β_i coefficients and defining a set of p equations describing behaviour of the β_i coefficients themselves. Then, we receive

$$r_t = \beta_{0t} + \sum_{i=1}^{p} \beta_{it} r_{t-i} + \delta h_t + e_t, \quad e_t \sim N(0, h_t) \tag{2.2'}$$

$$h_t = \alpha_0 + \alpha_1 h_{t-1} + \alpha_2 e_{t-1}^2 \tag{2.3'}$$

$$\beta_{it} = \beta_{it-1} + v_{it}, \quad i = 0,...p \quad v_{it} \sim N(0, \sigma_i^2). \tag{2.4}$$

Such a model may be estimated using the standard Kalman Filter approach, with equation (2.2') being the measurement equation, and the set of state equations defined by equations (2.3') and (2.4). The parameters required to estimate time paths of β_{it}, i.e., $\delta, \alpha_0, \alpha_1, \alpha_2$ and all p values of σ_i^2 can be found by computing and maximising the likelihood function (Harvey(1996)).

3. The Monte Carlo Test

To prove the power of the above test we propose the following Monte Carlo test.

We generate a series of 1000 autocorrelated returns which follow the formula:

$$r_t = \beta_t r_{t-1} + \varepsilon_t, \quad where \, e\varepsilon_t \sim N(0, \sigma^2),$$

where the factor β_t changes over time as follows:

$$\beta_t = \frac{-0.8}{499}(t-2) + 0.8, \quad \begin{array}{l} if \;\; t = 2,...,501 \\ \\ if \;\; t = 502,...,1000. \end{array} \qquad \textbf{(3.1)}$$

To illustrate, the coefficient β_t starts from the value 0.8 and drops linearly to zero for the first 500 observations and then remains constant at zero. It means that efficiency, or rather inefficiency, changes with β_t: during the first 500 days a market is inefficient, but its inefficiency decreases linearly, and for the remaining 500 days a market remains efficient.

We perform the test for three values of the variance, σ^2:0.0005, 0.004 and 0.05. The first two values of σ^2 were taken to be close to values observed on emerging markets (the Budapest and the Warsaw Stock Exchanges, respectively). The third value of σ^2 is very large and we do not know any stock exchange characterised by such a large variance, but to analyse the influence of variability of returns on the precision of values estimated by our model, we decided to take such a large value of σ^2 into consideration. For each sequence of returns we estimate, using the Kalman Filter, all parameters required by the TEE. Hence, for each sequence $\{rLt\}_{t=1}^{1000}$ we obtain the corresponding sequence of β's. This procedure of generating the sequences $\{rLt\}_{t=1}^{1000}$ and, in consequence, the sequences $\{\beta_t\}_{t=2}^{1000}$, for each value of σ^2, i.e., $\sigma^2 = 0.0005$, =0.004 and =0.05 we repeat 2000 times[1].

For each sequence of $\{\beta_t\}$, let us call them $\{\beta_t^i\}_{i=1,...,2000}$ we calculate the mean and the standard deviation at each point of time $t=2,...,1000$, i.e.

$$\tilde{\beta}_t = \frac{1}{2000} \sum_{i=1}^{2000} \beta_t^i$$

$$\sigma_t = (\frac{1}{1999} \sum_{i=1}^{2000} (\beta_t^i - \tilde{\beta}_t)^2)^{\frac{1}{2}}$$

In Figure 3.1 we present results of the test. Since the time paths of the means and standard deviations for the three considered values of σ^2 were the same we only present one time path of the estimated β_1 coefficient[2]. In Figure 3.1.(a) the continuous line is the mean $\{\tilde{\beta}_t\}$, the dotted lines are the T-statistics, i.e. $\{\tilde{\beta} \pm 2\sigma_t\}$, and the dashed line corresponds to the assumed pattern

---

[1] The replication 10000 required by Monte Carlo was impossible due to the amount of time needed to find the maximum. to find the maximum of the Maximum Likelihood Function defined in the search process for unknown parameters from the state and measurement equations. Calculations to undertake 2000 replications for each value of σ² lasted 1 month on the UNIX machine available at London Business School.

[2] The observable differences were on the fourth and later decimal places which is not significant and rather related to numerical errors of calculations or too few performed replications (only 2000 for each σ²).

described by equation (3.1). For more direct illustration of the differences between the assumed and estimated time paths of the β coefficient Figure 3.1(b) is presented. Here the continuous line shows the difference between the continuous and dashed lines presented in part (a). The dotted lines, as previously, refer to a graphical demonstration of the t-statistics.

Figure 3.1 proves the applicability of our method for testing for evolving efficiency. The estimated and assumed time paths of β's are nearly identical. The differences are observable at the beginning of the paths estimated (we underestimate the real value), and in the middle, where the β coefficient changes its pattern (we overestimate the real value). Nevertheless, the differences are very small, less then 0.11 and 0.03, for the starting and middle departures, respectively. Moreover, they are insignificantly different from zero.

It is worth stressing that our TEE has two important properties. First, the time path of β's estimated does not depend on σ^2 - the variance of the daily returns. In other words, the variability of returns can be very large (0.05), or relatively law (0.0005), but if the returns follow the same pattern of autocorrelation, we should detect it. Second, the probability of the pattern detection does not change with changes in values of σ^2.

To compare our method with a traditional approach we repeated the test described above, but instead of searching for time varying β, we estimated the β coefficient constant over time. For the three considered values of σ^2 we obtained that β was equal to 0.26 (st. error = 0.04). It means that a hypothetical market, introduced by us, wpould appear inefficient with a significant rate over the whole period. But it is not true as the inefficiency decreased during the first 500 days and for the next 500 days the market was efficient.

For a further illustration of the TEE, in the following section, we test some empirical series recorded on a developed and developing stock exchanges.

4. The Developed Stock Exchange Test

As an example of a developed market we chose one of the oldest and biggest world exchanges: the London Stock Exchange. Since the illustration of the TEE, rather than a deep and detailed discussion of market efficiency itself, is our main target, we ignore problems such as the length of returns (daily, weekly, etc.), impact of transaction costs, and a bid-ask spread which are typically taken into account.

Testing for efficiency of a western market is also illuminating. The western markets, being a pattern for emerging markets, should show not only how an exchange should be regulated and organised, but also, in some sense, exemplary behaviour. In the light of this, we present the results of the TEE for the market index (FTSE 100) and one of the blue chips listed on the LSE (Marks & Spencer's (M&S) shares). Data series of daily prices were taken from the Extel Database provided by LBS for the period 02.01.1991 - 15.10.1997 which gave us 1711 daily data points. In the case of the share price series the closing prices were taken into consideration.

Figures 4.1-4.2 show the results of our data analysis for the FTSE 100 index and the M&S's shares. In each figure part (a) shows daily log returns and part (b) presents the result of the TEE. As previously, in Figure 3.1, as now in part (b), the continuous lines correspond to the time paths of the estimated β_{1t} coefficient, and the dotted lines to the t-statistics.

From the presented graphs we can conclude that the LSE, during last seven years, was rather stable, as the estimated β_{1t} coefficient is constant over time. Moreover, the β_{1t} coefficient estimated for the FTSE 100 index is close to zero and it is insignificantly different from this value. The β_{1t} coefficient estimated for the M&S's shares is also close to zero, but significantly different from this value. Nevertheless, a profit opportunity shown by the β_{1t} coefficient is virtually nonexistent as the coefficient itself is close to zero.

These results are confirmed by the traditional test. Namely, we estimated the value of the β_1 coefficient from the traditional approach, i.e. assuming that the autocorrelation coefficient is constant over time (equations (2.2)-(2.3)). We received exactly the same values for all the estimated parameters.

To complete this section, in Table 4.1, we present GARCH coefficients estimated from the TEE for the market index and the share in question.

shares	estimated GARCH processes
FTSE 100	$h_t = \underset{(0.00)}{0.0} + \underset{(0.017)}{0.919}\,h_{t-1} + \underset{(0.00)}{0.050}\,e_{t-1}^2$
Marks & Spencer	$h_t = \underset{(0.0)}{0.0} + \underset{(0.015)}{0.908}\,h_{t-1} + \underset{(0.022)}{0.054}\,e_{t-1}^2$

Table 4.1. The GARCH processes estimated for the FTSE 100 and the Marks & Spencer's shares. Standard errors are shown in brackets.

All these results are in agreement with the common expectation about properties of a developed stock exchange, i.e., that such an exchanges should be stable and relatively efficient. However, we would like to stress that due to the very limited series considered, the presented results cannot be treated as a deep discussion of the LSE as a whole. They should be seen as an illustration of the problem, especially, as they are in contradiction with results presented in the next section.

5. The Application to a Developing Stock Exchange

The purpose of this section is a further illustration of the TEE. This time the Budapest Stock Exchange, as an emerging market from Central Europe, will be investigated.

5.1. The Historical Background of the Budapest Stock Exchange

The idea for the creation of a stock exchange market by the Austro-Hungarian Monarchy dates from 1860, and was realised after four years of preparation, in January 1864. The market was officially opened with 21 listed companies. Because of the Revolution and the introduction of the idea of "common" property, the Stock Exchange was closed in March 1919, only to be reopened in the Autumn of the same year when the socialist Revolution was over. The Budapest Stock Exchange operated until 1948, when the Communist party took power and liquidated the multi-party systems together with the system of a private ownership.

At the end of the 1980's, because of political and economic changes in Central and Eastern Europe, once again the idea of creating a Budapest Stock Exchange, following an Anglo-Saxon pattern, was introduced. During the first trading session, which took place on 21 June 1990, 42 Members (i.e., 19 commercial banks and 23 brokers) traded shares of just one company (*Ibusz Rt.*). The next day, shares of *Konzum Kereskedel Rt*. joined the Exchange. In December 1990 there were 6 listed shares and one year later, in December 1991, 20 listed shares. In 1997, more than 50 companies offered their shares on the Budapest Stock Exchange.

Since 1990, significant changes occurred both in the trading system, in the regulation and procedures in the clearing system. In 1995, the London Stock Exchange admitted the BSE as a properly regulated stock exchange (the first one in Central Europe) and placed the BSE on its "List of Approved Exchanges".

Despite its relatively fast growth, the BSE, as many other emerging markets, suffers from thin trading and, as a consequence, from illiquidity. Thin trading on the official market is caused by little interest by potential investors as well as the market regulators themselves who allowed some transactions to be occur outside the official market. In this paper we do not want to deal with all the market phenomena and deeply analyse the market development in the context of market capitalisation, changes in regulations, etc. Instead of that, using the Test for Evolving Efficiency we want to see how market efficiency altered over time, or in other words, how the learning process evolved.

Wanting to have the highest quality data, we selected the five, best traded shares and the market index BUX as our sample for the TEE. In Table 5.1 we present general information about the chosen shares showing the day of the first listing session, the number of data points taken into consideration (all the examined samples cover the 02.01.1991-15.10.1997 period), capitalisation and a profile of a company[3].

---

[3] All the information about shares as well as the daily share prices, and the market index BUX were prepared on the basis of the Extel Data Base provided by LBS and the Bank of Hungary.

shares	1st listing session	number of observations	capitalisation (in min)	profile
Domus Kereskedelmi	23.12.92	954	882	home & office furniture retailer and wholesale
Globus Konzervipari	27.12.93	955	2,541	manufacture of preserved food, condiments, juices
Fotex	13.11.90	1706	13,098	holding company (goods and services)
Pannonplast	13.04.94	842	37,200	processing of plastic
Pick Szeged	21.12.92	1209	37,211	trader of agricultural products

Table 4.1. General information about selected companies listed on the BSE.

5.2. Results of the TEE

In Figures 5.1 - 5.6 the results of the TEE for the Hungarian market index (BUX) and five shares are presented (the notation is same as in Figure 4.1). From the very first sight we can see that the graphs differ significantly from those presented in the previous section for the LSE showing the more dynamic character of the BSE. The β_{1t} coefficient estimated for the BUX index is not constant over time as it was in the case of the FTSE 100 index. Moreover, the time path is significantly different from zero and it remains at the high level (above 20%) all the time. We do not observe any tendency towards zero either.

The results of the TEE for the Hungarian shares also are more dynamic in autocorrelation coefficients of returns than we observed in the case of the Marks & Spencer's shares. The time paths of the β_{1t} coefficients estimated for the Fotex's, Pick Szeged's and Pannonplast's shares do not show any movement towards efficiency. Even worse, the autocorrelation of the daily returns for the Pannonplast's shares seems to grow with time and the Fotex's shares manifest high fluctuations of the autocorrealation coefficient of returns without any sign of stabilisation. The Domus Kereskedelmi's and Globus Konzervipari's shares show also some changes in the autocorrelation pattern of returns. However, the estimated time paths of the β_1 coefficients remain insignificantly different from zero. It is worth noting that daily returns of these shares seem to be negatively autocorrelated.

In Table 5.1 we present the GARCH coefficients estimated from the TEE for the BUX index and the discussed shares.

shares	estimated GARCH processes
BUX	$h_t = \underset{(0.00)}{0.0} + \underset{(0.01)}{0.79}\,h_{t-1} + \underset{(0.00)}{0.096}\,e^2_{t-1}$
Domus	$h_t = \underset{(0.0)}{0.0} + \underset{(0.01)}{0.865}\,h_{t-1} + \underset{(0.05)}{0.13}\,e^2_{t-1}$
Globus	$h_t = \underset{(0.0)}{0.0} + \underset{(0.03)}{0.664}\,h_{t-1} + \underset{(0.14)}{0.206}\,e^2_{t-1}$
Fotex	$h_t\ \underset{(0.0)}{0.0} + \underset{(0.01)}{0.826}\,h_{t-1} + \underset{(0.01)}{0.197}\,e^2_{t-1}$
Pannonplast	$h_t = \underset{(0.0)}{0.0} + \underset{(0.04)}{0.220}\,h_{t-1} + \underset{(0.01)}{0.479}\,e^2_{t-1}$
Pick Szeged	$h_t = \underset{(0.0)}{0.0} + \underset{(0.02)}{0.71}\,h_{t-1} + \underset{(0.12)}{0.282}\,e^2_{t-1}$

Table 5.1. The GARCH processes estimated for the BUX index and Hungarian shares. Standard errors are shown in brackets.

Summarising, the behaviour of individual shares is quite different showing no clear evidence of the learning process. On the Budapest Stock Exchange the GARCH effects differ from stock to stock. However, we can observe that for all the stocks the influence of the component related to past shocks is stronger than the influence we observed for the London Stock Exchange.

6. Conclusions

In this paper we have extended the classical test for autocorrelation of returns by combining a multi-factor model with time varying coefficients and the GARCH-M approach to investigate evolving market efficiency. We successfully verified our test by employing Monte Carlo simulations and exemplary series from the London Stock Exchange. The application of our test to the Budapest Stock Exchange confirmed its usefulness for an investigation of the first stages of a market performance.

References

Bachelier, L. (1900), ``Theory of Speculation'', in Cootner, P.(ed.), The Random Character of Stock Market Prices, Massachusetts Institute of Technology Press, Cambridge, MA, 1964; Reprint.

Beaver, W. (1981), ``Market Efficiency'', Accounting Review, 56, January, pp. 23-37.

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

Fama, E.F. (1991), ``Efficient Capital Markets:II'', Journal of Finance, 46, pp. 1575-1618.

Harvey A.E., ``Forecasting, structural time series and the Kalman Filter'', Cambridge University Press, 1996.

Roberts (1967), ``Statistical versus Clinical Prediction of the Stock Market'', unpublished manuscript, Centre for Research in Security Prices, University of Chicago, May.

Theodore T.M. (1996), ``Trading Patterns of World Markets'', Journal of Portfolio Management, pp. 81-88.

Zalewska-Mitura A., ``Emerging Markets from Central and Eastern Europe: Problem of Thin Trading, Price Limits and Evolving Market Efficiency'', London Business School, PhD Thesis, 1998.
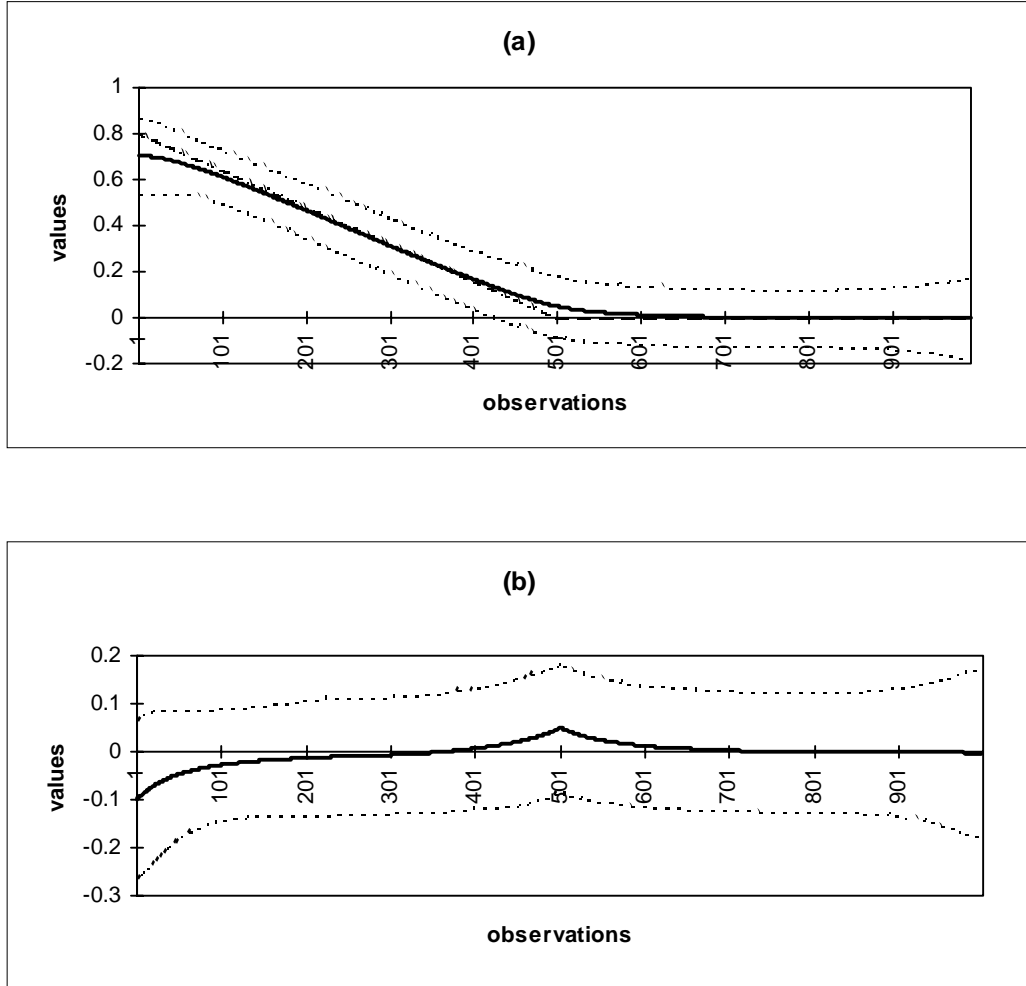
Figure 3.1. Results of the Monte Carlo simulations assessing the power of the Test for Evolving Efficiency. In part (a) the dashed line represents the assumed time path of the β coefficient, the solid back line represents the estimated time path and the dotted lines represent the T-statistics. In part (b) the black solid line corresponds to the difference between estimated and assumed values of the β coefficient. The dotted lines represent the T-statistics.
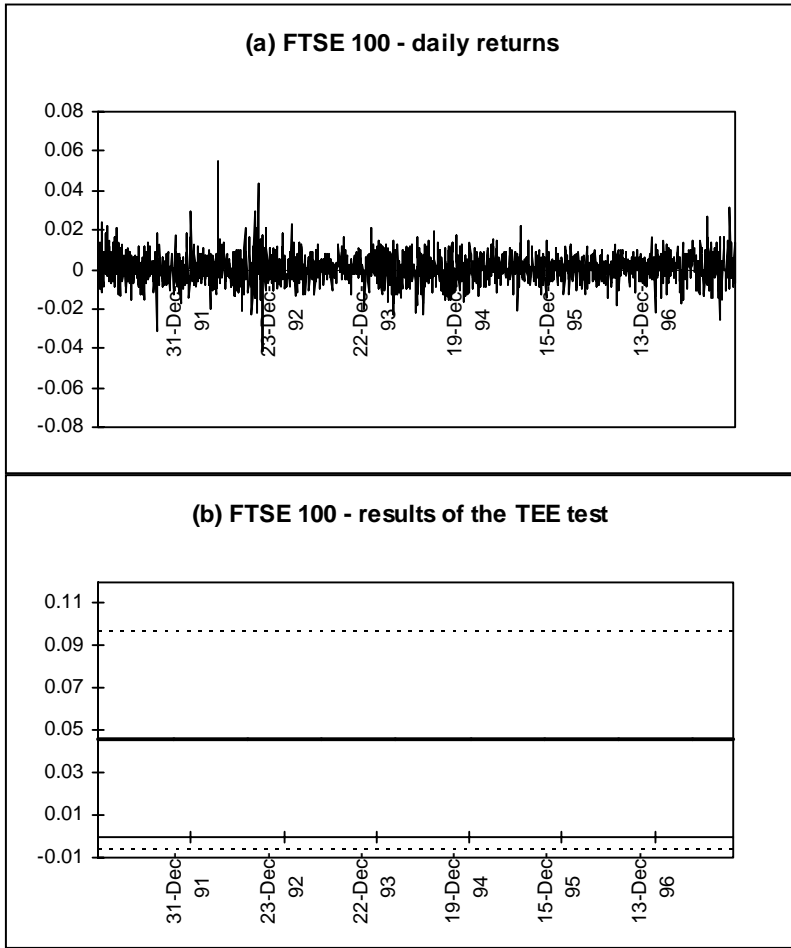
(a) FTSE 100 - daily returns

(b) FTSE 100 - results of the TEE test

Figure 4.1. Daily returns and results of the TEE for the FTSE 100 index

(a) Marks & Spencer - daily returns

(b) Marks & Spencer - results of the TEE
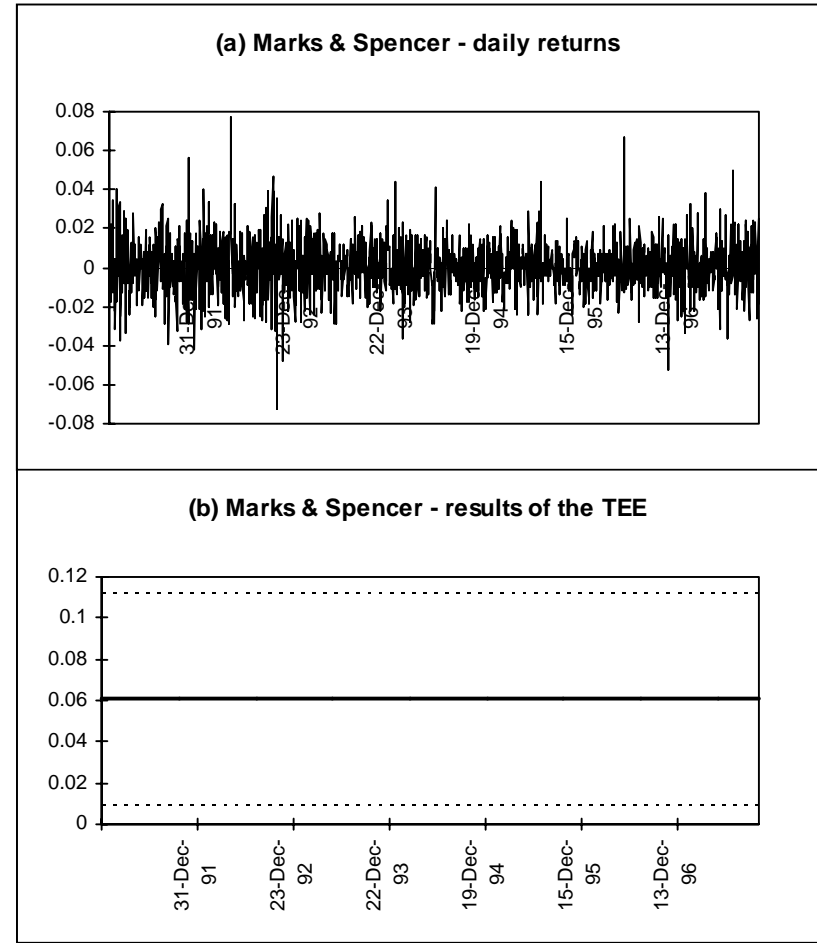
Figure 4.2. Daily returns and results of the TEE for the Marks & Spencer's shares
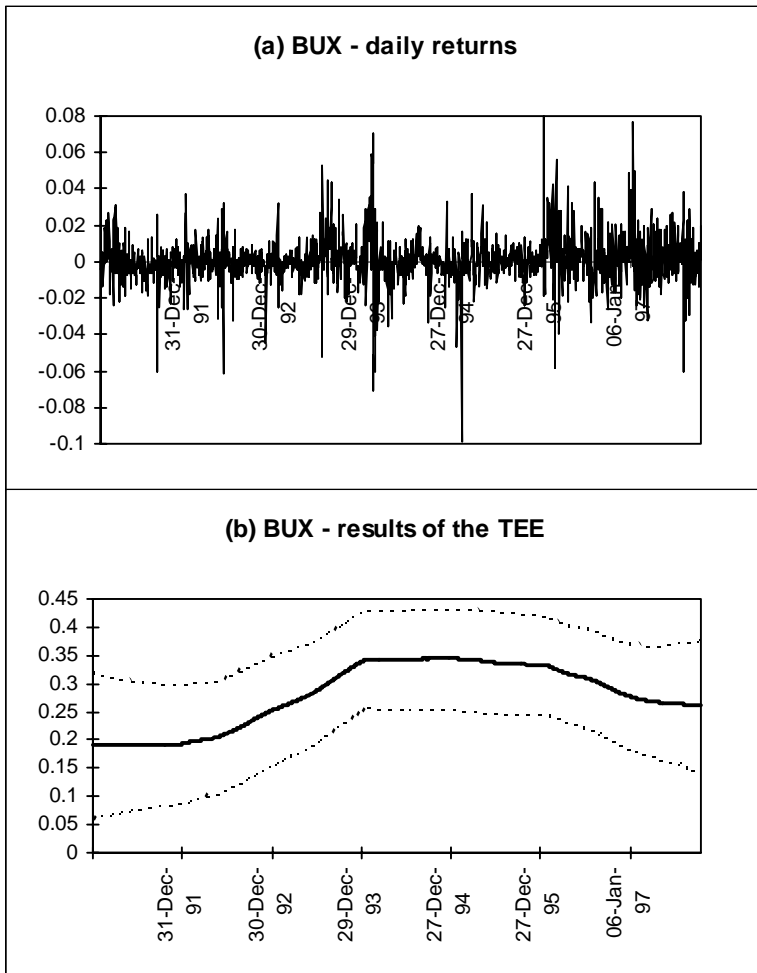
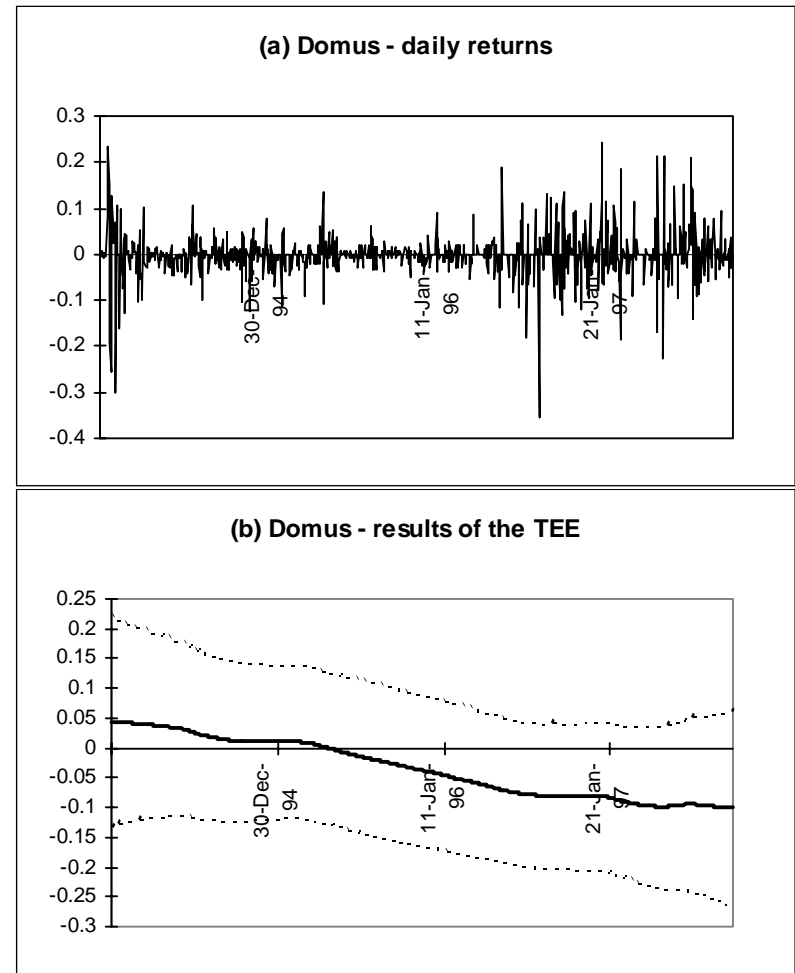Figure 5.1. Daily returns and results of the TEE for the BUX index.

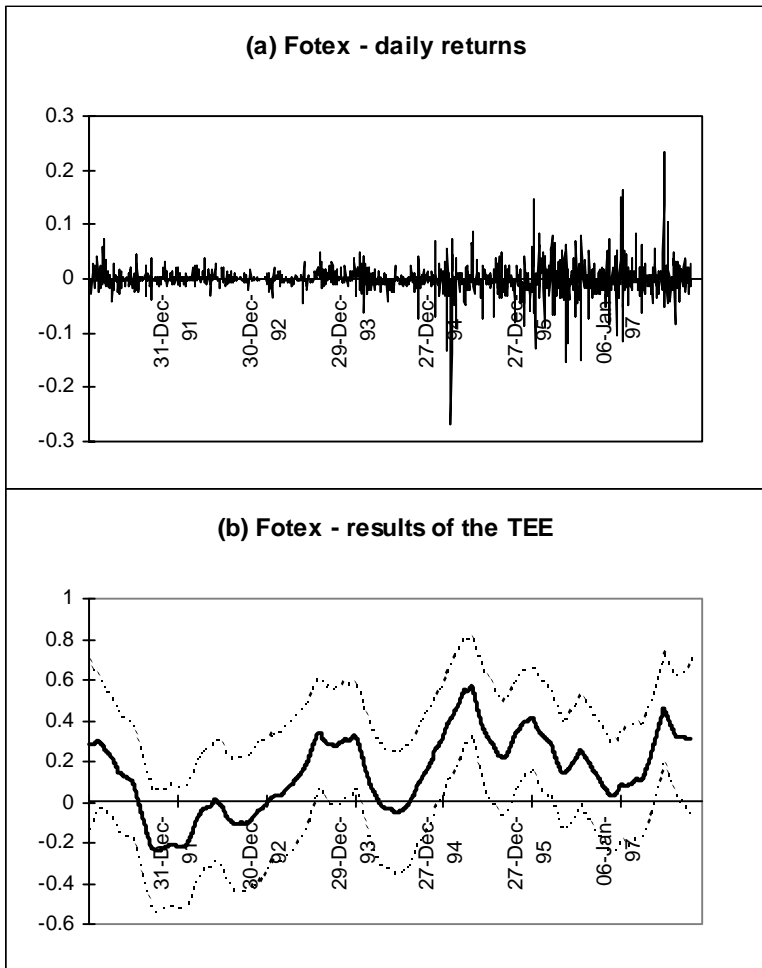Figure 5.2. Daily returns and results of the TEE for the Domus' shares.

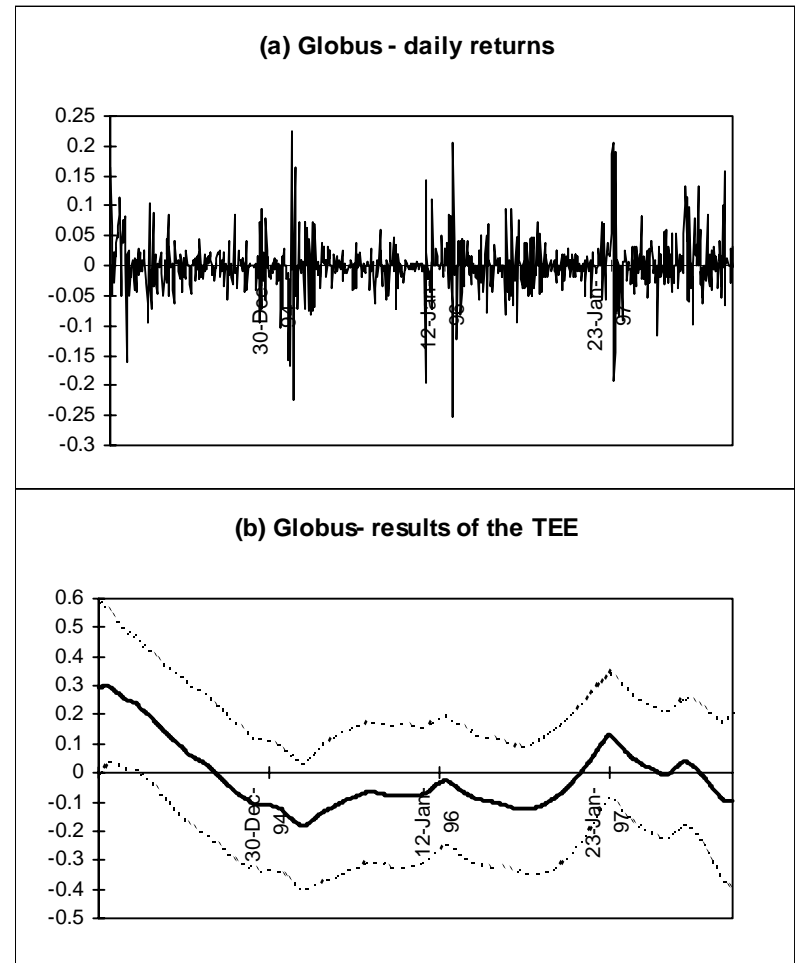Figure 5.3. Daily returns and results of the TEE for the Fotex's shares.

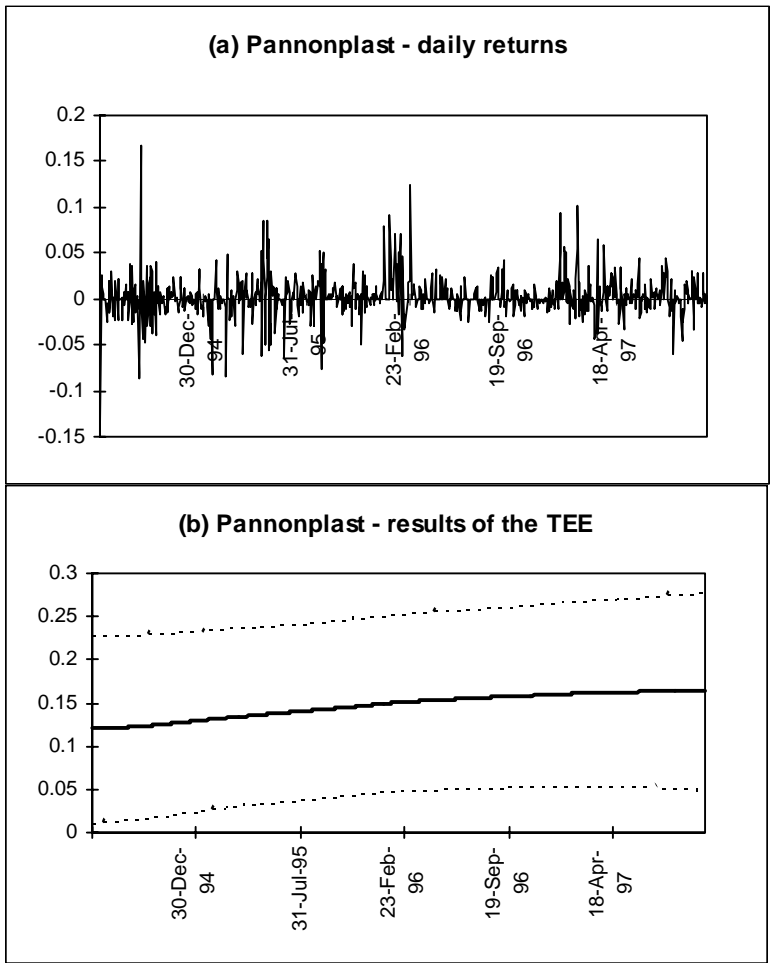Figure 5.4. Daily returns and results of the TEE test for the Globus' shares.

14

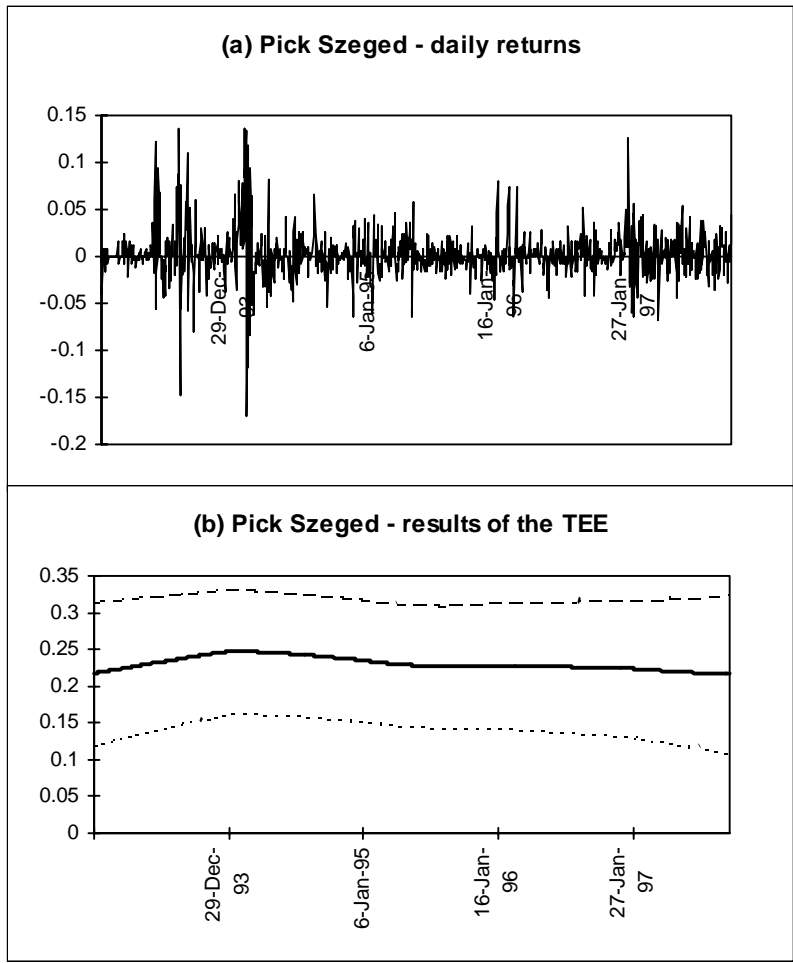Figure 5.5. Daily results of the TEE for the Pannonplast's shares.

Figure 5.6. Daily returns and results of the TEE for the Pick Szeged's shares.