BANK OF GREECE

MODELLING ECONOMIC TIME SERIES IN THE PRESENCE OF VARIANCE NON-STATIONARITY: A PRACTICAL APPROACH

Alexandros E. Milionis

Working Paper

No. 7 November 2003

MODELLING ECONOMIC TIME SERIES IN THE PRESENCE OF VARIANCE NON-STATIONARITY: A PRACTICAL APPROACH

Alexandros E. Milionis
Bank of Greece, Department of Statistics

ABSTRACT

Although non-stationarity in the level of a time series is always tested (and there is a variety of tests for this purpose), non-stationarity in the variance is sometimes neglected in applied research. In this work, the consequences of neglecting variance non-stationarity in economic time series, and the conceptual difference between variance non-stationarity and conditional variance are discussed. An ad hoc method for testing and correcting for variance non-stationarity is suggested. It is shown that the presence of variance non-stationarity leads to misspecified univariate ARIMA models and correcting for it, the number of model parameters is vastly reduced. The implications of the tests for the hypothesis of weak form market efficiency (WFME) are discussed. More specifically it is argued that the usual autocorrelation tests are inappropriate when based on the differences of asset prices. Finally, it is shown how the analysis of outliers is affected by the presence of variance non-stationarity.

Key words: Applied time series analysis, economic time series, Box-Jenkins modelling, variance non-stationarity, conditional variance, outlier analysis, efficient market hypothesis.
JEL classification:C22

The author is grateful to Steven Hall and Heather Gibson for helpful comments. The paper reflects the views of the author and not necessarily the views of the Bank of Greece.

Correspondence:
Alexandros E. Milionis
Department of Statistics,
Bank of Greece, 21 E. Venizelos Avenue,
102 50 Athens, Greece,
Tel. +30210-3203109, Fax +30210-3236035
Email:amilionis@bankofgreece.gr

1. Introduction

Although non-stationarity in the level of a time series can always be tested (and there is a variety of tests of this purpose), in applied econometric research, non-stationarity in the variance is often neglected. In this paper the consequences of neglecting variance non-stationarity are discussed. The principal problems are that univariate ARIMA models in the presence of variance non-stationarity can be seriously misspecified, and any analysis of outliers (i.e. aberrant observations) is invalid. These problems have important implications as will be argued later. The approach followed is a practical one. More specifically, in order to show the serious misspecification that can arise in ARIMA models as a result of the non-correction of variance non-stationarity, a time series of 3-month US government treasury bill rates for which a univariate ARIMA model has been suggested by Pindyck and Rubinfeld (PR henceforth, 1991), is reanalysed and a different model, in which variance non-stationarity is taken into account, is suggested. In addition, the same data set is analysed using the programme TRAMO (Gomez and Maravall, 1996) which selects the ARIMA model in an automatic way. The effect of variance non-stationarity in the identification of the various types of outliers will be examined using the above-mentioned series, as well as data on the consumer price index (CPI) for Greece. Once again, programme TRAMO will be used for this purpose. This analysis gives us the opportunity to make some additional remarks on the practical importance of variance non-satationarity on time series modelling. Finally the implications for autocorrelation tests of the hypothesis of weak-form market efficiency (WFME) are mentioned.

The rest of this paper is organised as follows: In section 2, some theoretical and practical aspects regarding non-stationarity are reviewed. The effect of variance non-stationarity on univariate ARIMA modelling is examined in section 3. Section 4 is devoted to the study of the effect of variance non-stationarity on the selection of various types of outliers. In section 5, the implications for testing of the WFME are discussed. Some concluding remarks are offered in section 6.

2. Non-stationarity

A stochastic time series is said to be strictly stationary if the joint distribution of any set of observations is unaffected by a change of time origin (Box and Jenkins, 1976). If a non-stationary series can be made stationary by just differencing it k times, the series is called homogeneously non-stationary of order k. However, as Box and Jenkins point out (1976), there is an unlimited number of ways in which a series can be non-stationary. Explosive behaviour (i.e. unit roots of the characteristic equation inside the unit circle) and evolutionary behaviour (as, for example, in the population growth of bacteria) are two cases in which the time series cannot be made stationary by just differencing. In practice, strict stationarity is difficult to test. Instead, the so-called weak or wide sense stationarity, which refers to moments up to the second order, is testable. The conditions for wide sense stationarity are that the mean and variance are constant, and that all autocovariances are a function of the time lag only (if the assumption of normality is accepted, wide sense and strict stationarity coincide).

There is a plethora of methods available for testing non-stationarity in levels (e.g. the Dickey Fuller and Augmented Dickey Fuller tests (Dickey and Fuller, 1979; 1981), the Phillips and Perron test (Phillips and Perron, 1988), the pattern of autocorrelation and partial autocorrelation functions and the pattern of the spectral density function (Box and Jenkins, 1976; Liu, 1988). However, non-stationarity may exist not only in the mean but also in the variance. Let X_t be a stochastic process. If $VAR(X_t)$ is somehow functionally related to the mean level of X_t, it is possible to select a transformation G(*) such that: $VAR(G(X_t))$= constant. The most widely used transformations for this purpose belong to the class of the power Box and Cox transformations (Box and Cox, 1964), although alternative transformations have also been suggested (e.g. Granger and Hughes, 1971). The Box and Cox transformation is given by the following expression:

$$G(X_t) = (X_t^{\lambda} - 1)/\lambda, \text{ if } \lambda \neq 0$$
$$G(X_t) = \log X_t, \text{ if } \lambda = 0$$

As Granger and Newbold (1977) note, although there are tests for variance non-stationarity, visual inspection of the data always remains a reasonable way for variance non-stationarity to be detected. It must be emphasised that variance non-stationarity is different from conditional heteroscedasticity, the latter being expressed by ARCH or GARCH type models (Engle, 1982; Bollerslev, 1986). In conditional heteroscedasticity, although the conditional variance is time varying, the unconditional variance is constant[1]. By way of an example, let the simplest case of an ARCH(1) model be considered. Then:

$$X_t = f(X_{t-1}; b) + e_t$$

where:

b is a vector of parameters,

$$e_t = v_t \sqrt{\omega + a_1 e_{t-1}^2}$$

$v_t = unit$ variance white noise

$\omega > 0$, and $0 < \alpha_1 < 1$

For the above model it is easily proved (e.g. Enders, 1995) that:

- The unconditional mean of e_t is equal to 0;
- The conditional variance of e_t is equal to $\omega + \alpha_1 e_{t-1}^2$, i.e. it is time dependent;
- The unconditional variance of e_t is equal to $\omega/(1-\alpha_1)$, i.e. it is a constant.

For such models the parameters of the conditional variance are simultaneously estimated with the parameters referring to the level of the series.

From the methodological point of view, variance non-stationarity should be corrected before removing non-stationarity in levels[2].

3. The effect of variance non-stationarity on univariate ARIMA modelling.

The general form of an $ARIMA(p,d,q)(P,D,Q)_s$ model is given by the following expression:

$$\Phi(B)\delta(B)Y_t=\Theta(B)\varepsilon_t+\mu$$

where:

B is the backward shift operator ($B^nX_t=X_{t-n}$);

s is the seasonality, if the model is seasonal;

p,d,q are the orders of the autoregressive polynomial, differencing, and moving average polynomial, respectively;

P,D,Q are the orders of the seasonal autoregressive polynomial, seasonal differencing, and seasonal moving average polynomial, respectively;

$\Phi(B) = (1-\varphi_1B-...........-\varphi_pB^p)(1-\Phi_1B^s-.....-\Phi_PB^{sP})$

$\Theta(B) = (1-\theta_1B-...........-\theta_qB^q)(1-\Theta_1B^s-.........-\Theta_qB^{sQ})$

μ is the mean level

ε_t is a white noise process

$\delta(B) = (1-B)^d(1-B^s)^D$ is the operator for regular and seasonal differencing.

As a first example of a univariate model in which variance non-stationarity has not been taken into account, the model of PR (1991, example 15.1) will be considered. The time series is that of the monthly rates of 3-month US government treasury bills from January 1950 to June 1988. The plot of this series is shown in Figure 1, while Figure 2 shows the first differences (both taken from PR). Then, using the Box-Jenkins methodology (Box and Jenkins, 1976) and experimenting with several rival univariate ARIMA models, PR select the following ARIMA(12,1,2) model as the most appropriate:

$$(1+.4211B+.4811B^2+0.0928B^3-0.2139B^4-0.0777B^5+0.2512B^6+0.1490B^7+0.1340B^8$$
$$-0.1556B^9-0.0272B^{10}-0.1171B^{11}+0.1559B^{12})(1-B)Y_t=$$
$$= 0.0109 + (1+ 0.8562B + 0.6257B^2)\varepsilon_t$$

with LBQ(36) = 28.16.

PR base their decision for the selection of the particular model mainly on the statistically insignificant value of the Ljung - Box statistic (LBQ) at the 5% level, which means that the hypothesis that the residuals of the model are white noise cannot be rejected. However, a univariate model with 15 estimated parameters (12 autoregressive coefficients, 2 moving average coefficients, and a constant) is not parsimonious, and parsimony is an essential requirement for a good model. Moreover, all parameter estimates must be statistically significant. The results of the re-estimation of the model of PR by the method of backscating (Liu, 1988) are given in Table 1. As it is evident from the results, three of the estimated coefficients and the constant are not statistically significant at the 5% level. Possibly PR were interested in the significance of all the AR coefficients jointly. However, from the correlation matrix of parameter estimates which is given in Table 2 it is apparent that several correlations are high (e.g. between θ_1 and φ_1) indicating the presence of problems of parameter redundancy often found in models where both autoregressive as well as moving average coefficients are considered.

The autocorrelation and partial autocorrelation functions (ACF and PACF, respectively) of the series are shown in Figures 3 and 4 respectively. As is evident from these plots, there are several not statistically significant autocorrelation and partial autocorrelation coefficients up to lag 12, and that explains the statistical insignificance of the coefficients discussed earlier. On the other hand, there are several significant spikes in both plots at lags higher than 12, and the intention of PR to reduce the LBQ value is another possible explanation for their selection of AR components up to lag 12. The character of both the ACF and PACF is rather confusing and does not favour the selected model.

The source of most problems in regard to the above model is the fact that PR did not take into account variance non-stationarity, which from Figure 2 is most likely to be present. PR in a subsequent chapter of their textbook use this model as part of a combined time series-regression model (chapter 18), where in a footnote they quote for the residuals of the model that *"the residuals do, however, exhibit heteroscedasticity. One could correct for this when estimating the model, but we have not chosen to do so"*. Although it is acknowledged that PR did not dealt with heteroscedasticity for pedagogic purposes, their decision has far more serious consequences than their footnote implies. The lack of homogeneity in the variance distorts the variance, as well as the autocovariances of the time series, and terms that differ from previous or subsequent ones substantially, tend to be correlated with each other, even if they are several lags apart, possibly resulting in statistically significant coefficients at some of these lags. Such problems lead to overparameterised models.

Weiss (1984) showed that ignoring conditional heteroscedasticity would also result in a similar situation, i.e. an overparameterised ARMA model.

From Figures 1 and 2 it may be seen that when the 3-month treasury bill rates rise, so does volatility. In practice the dependence of volatility on the mean level is examined using the so-called range-mean regression[3] (Gomez and Maravall, 1996). This method does not reveal the particular functional form of the relationship between variance and mean, but, due to its robustness, is considered suitable when the size of a particular time series realisation is relatively small. The logarithmic transformation is used in all cases where the regression coefficient for the mean in the range-mean regression is statistically significant.

For relatively long realisations, as the one under analysis, the following alternative methodology is suggested:

1) The series is sliced into sections of equal length and the mean of each section (Mi) is calculated.

2) For each section the existence of a time trend is examined. For those cases where a trend is present the corresponding sections are detrended.

3) The standard deviation in each section (SDi) is estimated.

4) A relationship of the form $SD_i = \alpha M_i^{\beta} u_i$ is assumed.

5) α,β in the above relationship are estimated using the regression:

$\log(SD_i) = \log\alpha + \beta\log(M_i) + \varepsilon_i$

6) From the estimated value of β the transformation that stabilizes the variance can be derived (e.g. if β=1, the variance stabilizing transformation is the logarithmic, if β= 0.5, the suitable transformation is the square root, etc, see Mills, 1991).

The advantage of this methodology over the range-mean regression is that it is suggestive regarding the identification of the particular functional form between mean and variance[4].

The estimates of the coefficients of the linear regression of $\log(SD_i)$ against $\log(M_i)$ are shown in Table 3.The estimated value of β (0.83) is closer to 1, rather than to 0.5, suggesting the log transformation. Figure 5 shows the scatterplot of the local mean levels against the local standard deviations, which provides a visual inspection of the linear relationship between local mean and local standard deviation. The correlation coefficient was found to be 0.82. On the other hand, the correlation coefficient between the local mean and the local variance was 0.77 and the corresponding plot (not shown) showed a less profound linear relationship than in the case of local mean and local standard deviation. Hence, the log-transformation was used to stabilise the variance[5].

The patterns of ACF and PACF for the first differences of the log-transformed series are shown in Figures 6 and 7 respectively. It is clear from these figures that a much smaller number of parameters is necessary, and as a matter of fact the pattern of both the ACF and PACF are a little more suggestive in regard to the underlying stochastic process. In Figure 6, there is a significant correlation of lag one of the ACF, while, in Figure 7, in the first few lags of the PACF a pattern of a damped sinus wave is recognisable. This clearly suggests a moving average process of order one (see Box and Jenkins, 1976 p.70 for a theoretical justification). Further, the two significant correlations at lags 6 and 7 in the ACF and the significant correlation at lag 6 in the PACF suggest an autoregressive process. Finally, from the significant correlations at lags 18 and 19 in the ACF and at lag 17 in the PACF, a specific process cannot be uniquely identified, but the high lags of these autocorrelations may be attributed to the existence of extreme values (outliers), which have not been smoothed out with the log-transformation. No correction of outliers was made at this stage for the sake of comparability with the model of PR. Although it is also possible that such autocorrelations at high lags are a result of remaining stationary autoregressive conditional heteroscedasticity[6], a simple reason for this large-lag significant autocorrelation may just be the dependence of the autocorrelation estimators on the same time series.

Based on these comments, models with four or three parameters were estimated and several of them were found to be adequate. Based on the minimum value of both the Akaike information criterion (Akaike, 1973) and Schwartz bayesian criterion (Schwartz, 1978) a multiplicative 3-parameter model was selected. Parameter estimates are presented in Table 3. Figure 8 shows the ACF of the residuals of the model. It can be seen that there is a significant coefficient at lag 7. However, at the 5% significance level, purely by chance, it is expected that 1 in 20 lags is significant. Alternatively this significant autocorrelation may again be attributed to the distorting effect of outliers. The values of the LBQ statistic at lags 24 and 36 are approximately 27 and 35 respectively. Both values are not significant at the 5% level. Hence, for the purposes of this work the residuals can be considered as white noise. (It is possible that higher order dependencies (ARCH effects) may be present in the residuals but such a case was not examined).

From the above analysis it is clear that by just transforming the data to obtain variance stationarity the result is really remarkable. The number of model parameters has been reduced from 15 to only 3. Needless to say the economic interpretation implied by the new model dramatically differs from that implied by the model of PR. In the 15-parameter model, changes in 3-month treasury bill rates depend on the changes of all the 12 previous months as well as the shocks of the last two months. The revised model implies that percentage changes in 3-month treasury bill rates (which are approximately equal to the logarithmic differences of the rates) depend mainly on the previous month's shock of percentage changes and, to a lesser extent, on percentage changes of 6 lags (months) apart.

In the above analysis the classical model building procedure of identification - estimation -diagnosis - metadiagnosis, proposed by Box and Jenkins (1976) was used. As an alternative to the above model building procedure, in which the role of the analyst is crucial, an additional analysis was performed using the TRAMO programme (Gomez and Maravall, 1996). This programme can suggest an ARIMA model automatically, but the orders of the regular autoregressive and moving average polynomials are restricted to up to 3 and those of the seasonal ones to 1. The programme offers the possibility of identifying automatically three types of outliers, using the methodology proposed by Box and Tiao (1976) and developed further by Tsay (1986), Chang et al. (1988) and Chen and Liu (1993) and modifying the value of the corresponding observations. More specifically, the TRAMO programme discriminates among three types of outliers according to their effect on a time series: (i) additive outliers (AO) which affect only a single observation of the series (ii) level shifts (LS) which imply a step change in the level of the series and (iii) transitory changes (TC) the effect of which is not extinguished in the next observation, as is the case with the additive outliers, but damps out gradually over a few periods.

At first, TRAMO, using the result of a range-mean regression, log-transforms the data. Without any correction for outliers TRAMO failed to identify a model with acceptable diagnostics, as it cannot incorporate higher order autoregressive and/or moving average processes. When outlier detection and correction was chosen as an option, a seasonal ARIMA(0,1,1)$(0,1,1)_{12}$ model was suggested with acceptable diagnostics. However, this was made possible only after the value of 16 observations changed as they were detected and corrected as outliers[7]. However, the nature of the data (interest rates) does not suggest a seasonal model. Again such a model was chosen by the programme as this was the only way for higher order processes to be included in the model. Higher order processes, however, are often found in ARIMA models for interest rates (e.g. Mills, 1991, example 15.1).

4. The effect of variance non-stationarity on the analysis of outliers.

The most interesting part of the analysis with TRAMO in the previous section refers not to the model itself but to the selection of outliers. The order of the observations (year, month) selected as outliers, as well as the type of outlier corresponding to each of them are shown in Table 5. From the results of Table 5 it is very remarkable that most of the outliers are centred in the first part of the series (i.e. during the 1950s), rather than during the 1980s, as one would expect looking at Figure 2. This is another important effect that the variance stabilising transformation incurs and should be taken into consideration particularly if someone tried to give any meaning to the outliers. Indeed, if the data are analysed without being log-transformed, most of the outliers are concentrated in the latest part of the series. Although, the large number of outliers detected by TRAMO may again be attributed to the fact that the programme is not capable of incorporating higher order (>3) stochastic processes in an ARIMA model, the main argument (dependence of outlier detection on the variance stabilising transformation) does not change.

In order to provide further evidence of the effect of variance non- stationarity on the selection of outliers, the series of the CPI for Greece (monthly values from January 1970 to December 1990) was also used. The CPI series itself is shown in Figure 9, and the series regularly and seasonally differenced in Figure 10. From Figure 10, it is obvious that the series is variance non-stationary. However, at first an ARIMA model was created, using TRAMO, without taking into account variance non-stationarity. TRAMO selected an $ARIMA(0,1,3)(1,1,0)_{12}$ i.e. a multiplicative seasonal model. The parameter estimates are given in Table 6. The basic diagnostics for the residuals (i.e. no linear dependencies in the residuals indicated by LBQ below the critical value and no statistically significant correlations at low-lags) denote an initially acceptable model. The outliers found, as well as their type, are listed in Table 7.

When non-stationarity in the variance is taken into account, and following the same procedure as in the previous section, it is found that the logarithmic transformation must be used to stabilize the variance. In this case (log-transformed data) TRAMO selected an ARIMA(1,1,1)(0,1,1)$_{12}$ model. The parameter estimates for this model are shown in Table 8. Again in comparison to the model for the CPI without the logarithmic transformation the number of parameters has been reduced from 4 to 3. However, the number of observations that had to be characterised as outliers in the first case was much larger than in the second case. Indeed, the outliers identified in the second case (log-transformed data) and their characteristics are listed in Table 9. A comparison of the results in Tables 7 and 9 shows that a completely different set and type of outliers is selected, after the data have been log-transformed. Without the log-transformation all (seven) outliers are located in the latest part of the series. Using the log-transformation, the only two outliers (level shifts) are located in the early part of the series.

5. Implications for the autocorrelation tests for the WFME hypothesis.

The hypothesis of efficient markets states that security prices fully reflect all available information (Fama, 1970). For the case of WFME the available information confines to the past history of prices. To make the hypothesis of WFME empirically testable it is first assumed that equilibrium conditions can be expressed in terms of expected returns. The expected returns are determined adopting a pricing model, hence, the test of WFME is in fact a joint test of WFME and the pricing model. If the adopted model is that of constant expected returns, in this risk-unadjusted framework it makes sense to perform tests for autocorrelation in the security returns[8]. Granger (1975) argues that the conclusion about the statistical significance of the autocorrelation coefficients is the same and does not depend on whether the first differences of prices, or the first differences of the logs of prices (which are the continuously compounded rates of return) are used. Elton and Gruber (1995) quote a comprehensive list of results from autocorrelation tests performed by several researchers, in which the first differences of either the prices or the logs of prices are used. From the analysis of the previous sections it is clear that such tests are valid only under the condition of variance stationarity in the series of first differences of prices. If this is not true, then the significance testing of the autocorrelation coefficients is invalid, as already shown. Variance non-stationarity in the first differences of prices is naturally expected for the following reason: if the model of constant (or near constant) expected returns is assumed, then when prices are comparatively higher (lower), so are price changes so as the returns are kept constant. Hence, the first differences of prices (price changes) are proportional to the mean level (prices) and consequently the variance is functionally related to the mean level. As the series of prices is in the vast majority of cases non-stationary of order one (in the first moment), prices are free to wander extensively and, hence, the variance is non-stationary. Consequently, autocorrelation tests cannot be applied to price changes.

6. Conclusions

In this paper three approaches to modelling variance non-stationary time series were examined. It was shown that neglecting variance non-stationarity may result to a serious distortion of a univariate time series model, and, inevitably, to a misconception of the underlying stochastic process and incorrect economic interpretation. Aside from a visual inspection of the original non-stationary series and its differences, statistically significant correlations scattered irregularly at different lags in both the ACF and PACF is another indication of variance non-stationarity and the analyst should consider the possibility of transforming the data, using a suitable transformation, before modelling the series. Although ad hoc and simple in nature, the method of splitting the series into sections of equal length and exploring qualitatively the possibility of a relationship between the means and the corresponding variances works well in revealing the existence and the character of variance non-stationarity. This ad hoc method, however, needs to be improved further, inasmuch as the separation of the series into sections of equal length is, to an extent, arbitrary and a subjective method is required.

In addition, it was shown that variance stationarity is a crucial prerequisite for outlier analysis, and the identification of outliers of any type as well as any economic significance, which might be assigned to them, is invalid if this is not the case. The variance stabilising transformation should precede the search for outliers[9].

As time series of asset prices are in most cases non-stationary and, hence, prices can wander extensively, in a framework of constant (or nearly constant) expected returns, prices are variance non-stationary, as price changes must be proportional to prices. Hence, price changes for autocorrelation tests for WFME should not be used. The log transformation, which is applied to asset prices before differencing, to a certain extent, stabilises the variance in the series of prices and then attention focuses directly on the modelling of (stationary) conditional variance of the returns in conjunction with the model for the returns themselves. However, this is not the case with interest rates and this is why in this work attention initially focused on the stabilisation of the unconditional variance. Of course after a suitable transformation has made the (unconditional) variance stationary, models for the conditional variance can be considered. This work, as mentioned previously, has been restricted to the unconditional variance only.

Finally, it is very important to note that as series which are variance non-stationary cannot have a proper autoregressive representation, if these series are not corrected for non-stationarity in the variance, any econometric methodology which is associated with autoregressive representation (VAR, ADF test, cointegration tests) cannot be legitimately applied using such series.

Endnotes

1. The special case of an integrated GARCH (IGARCH) process where volatility persists may need further investigation in regard to this aspect. It has been shown (Nelson, 1990) that, unlike the usual random walk, the process is strictly stationary, but its unconditional variance is infinite.

2. An alternative procedure has been suggested by Nelson and Granger (1979) in which the value of λ is simultaneously estimated with the values of the model parameters.

3. Robustness could possibly be better assured by performing a range-median, rather than a range-mean regression.

4. More generally, it can be assumed that changes in the variance of a time series may be attributed to the fact that the stochastic process responsible for the evolution of the series is not unique so that different processes act over different subsamples. Such differences in "regime" may be accounted for by using another variable (Hamilton 1994; Hamilton and Raj 2002), which is also random. While this more general framework is appealing, the methodology proposed here is preferred if changes in the variance are assumed to be functionally and, hence, systematically linked to changes in the mean level and this statement is confirmed empirically.

5. It must be noted at this point that before the application of the log-transformation it was checked that there were no zero values in the original series. In general for series of rates of returns the log transformation is usually applied in the series $1+Y_t$ instead of the series Y_t itself.

6. Another possibility for the appearance of significant spikes at high lags of the ACF of a time series, which, however, is not the case here, is when the series is not invertible (i.e. having a zero in its spectra). This occurs for instance in some cases when the series is filtered with a linear filter. A classic example is the seasonally adjusted series (Maravall, 1995).

7. The default parameters of the programme for outlier detection were used.

8. Caution is needed when forecasts of the originally observed series are required (see Nelson and Granger, 1979).

9. Sometimes practitioners apply such tests in a rather mechanistic way without much consideration of the assumptions under which the results of the tests are meaningful. Statistically significant correlations do not necessarily mean that the WFME

hypothesis is rejected, as what is really being tested is the joint hypothesis of efficiency along with the pricing model (constant expected returns). Moreover, even if statistically significant correlations are found and the model of constant expected returns is not rejected, these statistically significant autocorrelations may be attributed to frictions in the trading process (Cohen et. al 1980). In fact, frictions in the trading process lead to a distinction between observed returns, which reflect the frictions, and true returns, which are unobservable and would reflect a frictionless market (Cohen et. al 1980). True returns may be uncorrelated, even though observed returns are correlated.

References

Akaike, H (1976): A New Look at the Statistical Model Identification. IEE Transaction on Automatic Control, AC-19, 716-723.

Bollerslev, T. (1986): Generalised Autoregressive Conditional Heteroscedasticity, Journal of Econometrics, 31, 307-327.

Box, G., and Jenkins G. (1976): Time Series Analysis: Forecasting and Control. Holden Day, San Francisco.

Box, G., and Cox, D. (1964): An Analysis of Transformations. Journal of the Royal Statistical Society, series B, 26, 211-252.

Box, G. E. P. and Tiao, G. C. (1975). Intervention Analysis with Application to Economic and Environmental Problems, Journal of the American Statistical Association, 70, 70-79.

Chen C. and Liu L. M. (1993). Joint Estimation of Model Parameters and Outlier Effects in Time Series. Journal of the American Statistical Association 88, 284-297.

Chang I., Tiao G. C. and Chen C. (1988). Estimation of Time Series Parameters in the Presence of Outliers. Technometrics, 30, 193-204.

Cohen K., Hawawini A., Maier S., Schwartz R., and Whitcomb D. (1980).Implications of Microstructure Theory for Empirical Research on Stock Price Behaviour. The Journal of Finance, XXXV(2), 249-257.

Dickey, D., and Fuller, W. A (1979): Distribution of the Estimates for Autoregressive Time Series with Unit Root. Journal of the American Statistical Association, 74, 427-431.

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

Enders, W. (1995): Applied Econometric Time Series. Wiley, New York.

Elton E. J. and Gruber M. J. (1995). Modern Portfolio Theory and Investment Analysis, fifth edition, J. Wiley and Sons Inc, USA.

Engle, R. F. (1982): Autoregressive Conditional Heteroscedasticity, with Estimates of The Variance of United Kingdom Inflation. Econometrica, 50, 987-1007.

Fama E. (1970). Efficient Capital Markets: A Review of Theory and Empirical Work. Journal of Finance, XV(2) 383-417.

Gomez, V. and Maravall, A. (1996): Programmes SEATS and TRAMO: Instructions for the User. Working Paper No 9628, Bank of Spain

Granger C. W. J. (1975). A Survey of Empirical Studies on Capital Markets, in Elton and Gruber (eds) International Capital Markets, North Holland, Amsterdam.

Granger, C. W. J., and Hughes, A. D. (1971): A New Look at Some Old Data. Journal of the Royal Statistical Society, Series A, 134, 413-428.

Granger, C. W. J., and Newbold, P. (1977): Forecasting Economic Time Series. NewYork, Academic Press.

Hamilton, J. D., and B. Raj (2002) New Directions in Business Cycle Research and Financial Analysis. Empirical Economics, 27, 149-162.

Hamilton, J. D. (1994) Time Series Analysis, Princeton University Press, Princeton, USA.

Liu, L. M. (1988): Box-Jenkins Time Series Analysis, in BMDP Statistical Software Manual, Vol. I, University of California Press.

Maravall, A. (1995): Unobserved Components in Economic Time Series, in Handbook of Applied Econometrics, (eds) Pesaran, M. H., and Wickens, Blackwell, Oxford.

Mills, T. C. (1990). Time Series Techniques for Economists. Cambridge University Press, Cambridge

Nelson D. B. (1900). Stationarity and persistence in the GARCH(1,1) model. Econometric Theory 6, 348 334.

Nelson, H. L., and Granger, C. W. L. (1979): Experience with using the Box-Cox Transformation when Forecasting Economic Time Series. Journal of Econometrics, 10, 57-69.

Phillips, P., and Perron, P. (1988): Testing for a Unit Root in Time Series Regression. Biometrica, 75, 33-346.

Pindyck, R. S. and Rubinfeld, D. L. (1991) Econometric Models and Economic Forecasts, Third Edition, McGraw-Hill, New York.

Schwarz, G. (1978): Estimating the Dimension of a Model. Annals of Statistics, 6,461-464.

Tsay, R. S. (1986): Time Series Model Specification in the Presence of Outliers. Journal of the American Statistical Association, 81, 132-141.

Weiss, A. A. (1984): ARMA Models, with ARCH Errors. Journal of Time Series Analysis, 5, 129-143.

Table 1

Summary of the ARIMA(12,1,2) model of Pindyck and Rubinfeld

Type	Order	Estimate	t-ratio
MA	1	-0,80	-5,33
MA	2	-0,61	-4,55
AR	1	-0,37	-2,40
AR	2	-0,48	-4,65
AR	3	+0,10	+1,55
AR	4	-0,21	-3,46
AR	5	+0,08	+1,42
AR	6	-0,26	-4,50
AR	7	-0,13	-2,05
AR	8	-0,13	-2,01
AR	9	+0,15	+2,70
AR	10	+0,02	+0,33
AR	11	+0,12	+2,44
AR	12	-0,18	-3,21

Table 2

Correlation Matrix of Parameter Estimates for the ARIMA(12,1,2) model of Pindyck and Rubinfeld (estimation by backkasting method).

	θ_1	θ_2	φ_1	φ_2	φ_3	φ_4	φ_5	φ_6
θ_1	1000							
θ_2	0,695	1000						
φ_1	0,95	0,68	1000					
φ_2	0,266	0,786	0,332	1000				
φ_3	0,031	-0,369	0,127	-0,257	1000			
φ_4	0,259	0,403	0,236	0,487	0,138	1000		
φ_5	0,173	0,151	0,214	0,089	0,314	0,405	1000	
φ_6	-0,321	-0,132	-0,328	0,085	-0,09	0,322	0,281	1000
φ_7	0,457	0,19	0,481	-0,018	0,279	0,06	0,447	0,09
φ_8	0,382	0,528	0,406	0,493	-0,141	0,375	0,102	0,273
φ_9	-0,167	-0,037	-0,13	0,133	0,136	0,061	0,178	0,064
φ_{10}	-0,371	-0,405	-0,391	-0,261	0,182	0,042	-0,016	0,287
φ_{11}	-0,088	-0,003	-0,114	0,009	-0,03	0,168	0,193	0,057
φ_{12}	-0,554	-0,344	-0,575	-0,142	-0,13	-0,063	0,011	0,333

Table 3

Estimation of the model: $\log(SD_i) = \log\alpha + \beta\log(M_i) + \varepsilon_i$

Coefficient	Estimate	t-ratio
α	-1.43	-5,08
β	0.83	4.55

Table 4

Summary of the 3-parameter model

A. Correlation matrix of parameter estimates

	θ_1	φ_6	φ_{19}
θ_1	1,00		
φ_6	-0,03	1,00	
φ_{19}	-0,03	-0,02	1,00

B. Estimation of the parameters (backcasting method)

Type	Order	Estimate	t-ratio
MA	1	-0,43	-10,06
AR	6	-0,20	-4,33
AR	19	-0,11	-2,40

Table 5

Detected Outliers and their type

Observation	Date		Type
	Month	Year	
104	8	1958	LS
98	2	1958	LS
365	5	1980	LS
105	9	1958	LS
126	6	1960	LS
296	8	1980	TC
366	6	1980	TC
363	3	1980	AO
377	5	1981	AO
392	8	1892	LS
383	11	1981	LS
46	10	1953	LS
49	1	1954	LS
50	2	1954	AO
286	10	1973	AO
64	4	1955	LS

Table 6

Estimation of the ARIMA(0,1,3)(1,1,0)$_{12}$ model

Type	Order	Estimate	t-ratio
MA	1	0.56	8.64
MA	2	0.14	2.00
MA	3	0.31	4.77
AR	12	0.35	4.90

LBQ(24) in residuals= 32.1

Table 7

Detected Outliers and their type

Observation	Year		Type
	Month	Year	
245	5	1990	LS
213	9	1987	AO
225	9	1988	AO
210	6	1987	AO
193	1	1986	AO
162	6	1983	LS
246	6	1990	LS

Table 8

Estimation of the ARIMA(1,1,1)(0,1,1)$_{12}$ model

Type	Order	Estimate	t-ratio
MA	1	-0.61	-6.87
MA	12	-0.79	-14.53
AR	1	-0.87	-16.10

LBQ(24) in residuals= 12.2

Table 9

Detected Outliers and their type

Observation	Date		Type
	Month	Year	
47	11	1973	LS
109	1	1979	LS

FIGURE 1:3-MONTH US GONERNMENT TREASURY BILL RATES (MONTHLY VALUES)

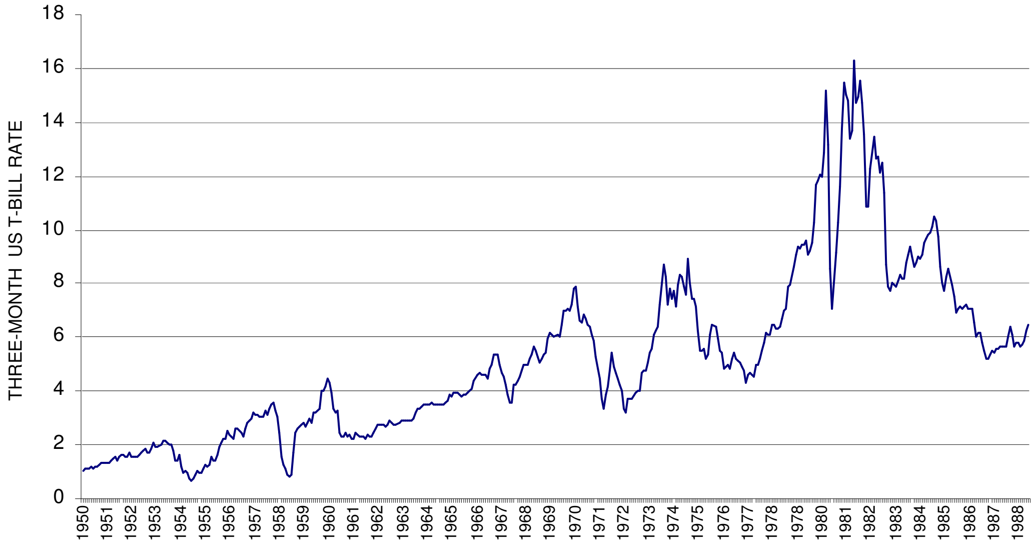

FIGURE 2: FIRST DIFFERENCES OF 3-MONTH US TRESURY BILL RATES

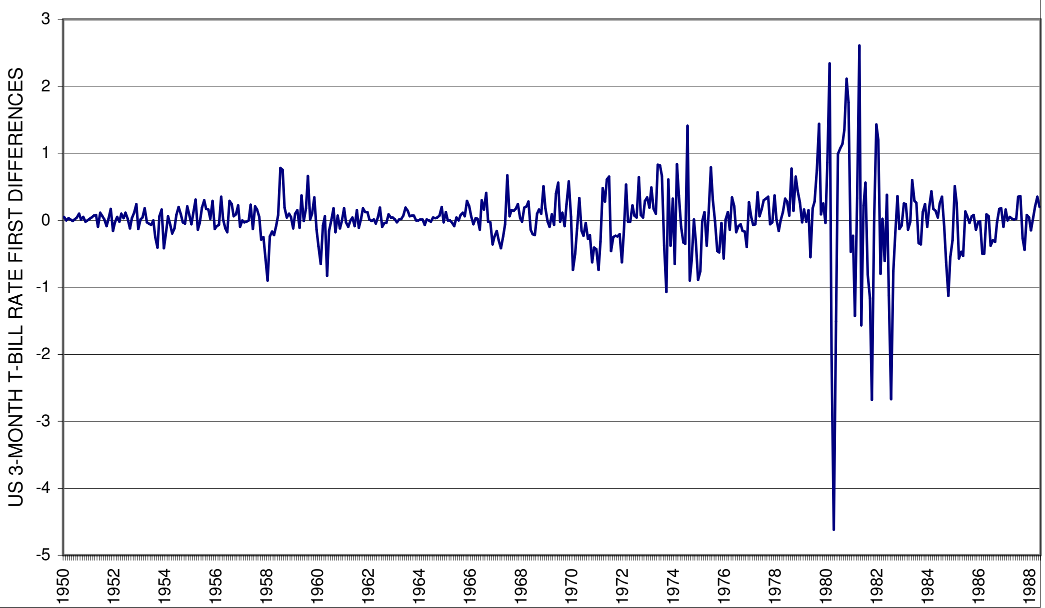

FIGURE 3: PLOT OF ACF. FIRST DIFFERENCES OF THE ORIGINAL SERIES. DASHED LINES REPRESENT 95% CONFIDENCE INTERVAL

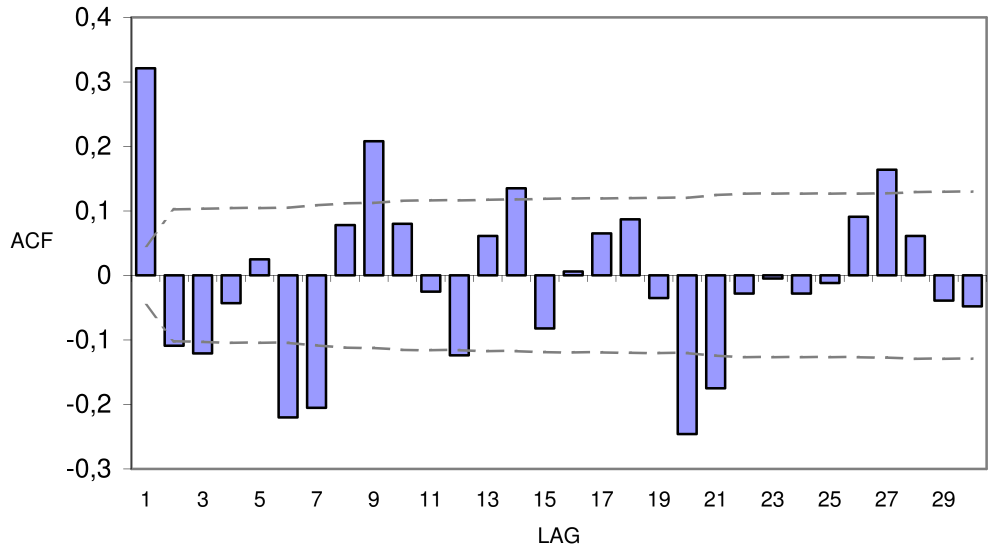

FIGURE 4: PLOT OF PACF. FIRST DIFFERENCES OF THE ORIGINAL SERIES

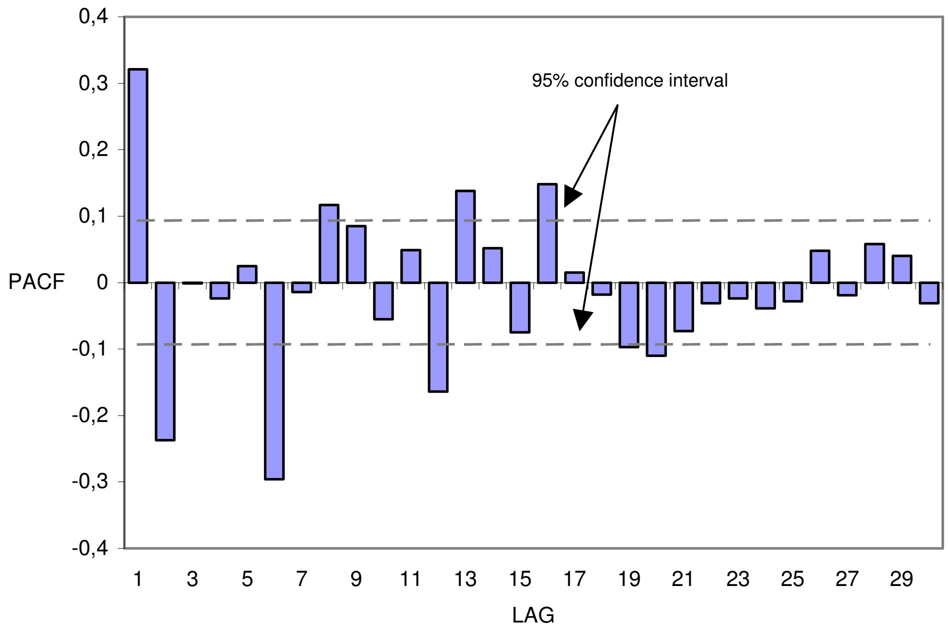

FIGURE 5: SCATTERPLOT OF THE LOCAL MEAN LEVEL AGAINST LOCAL STANDARD DEVIATION

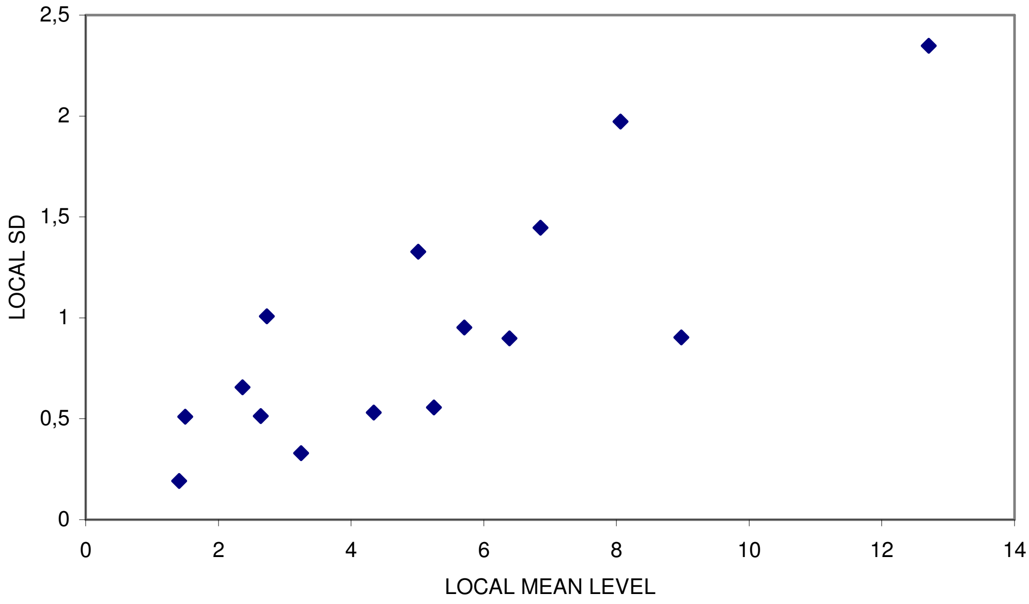

FIGURE 6:PLOT OF ACF. FIRST DIFFERENCES OF THE LOG-TRANSFORMED SERIES

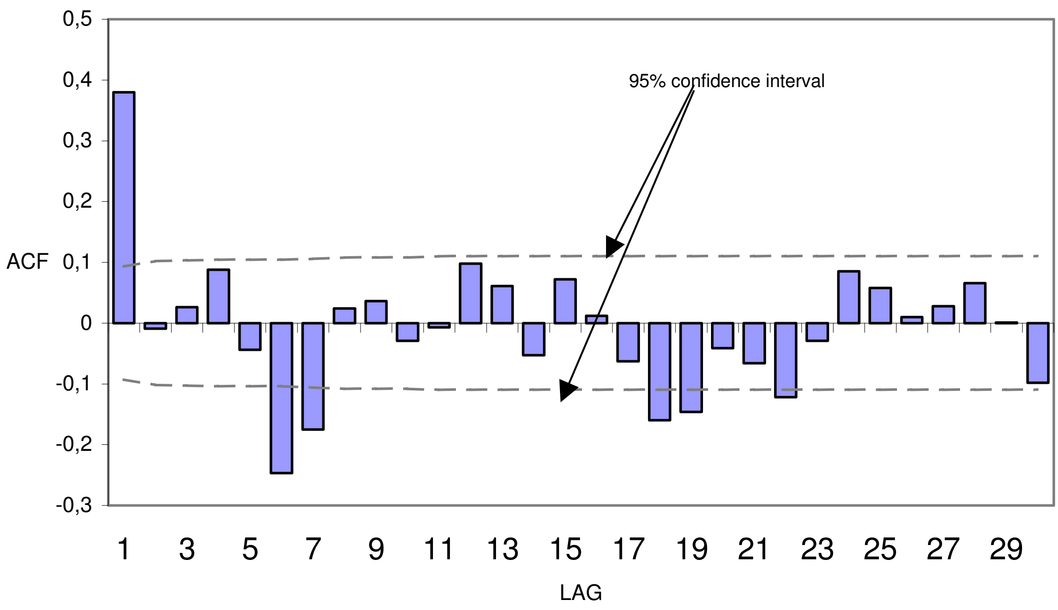

FIGURE 7: PLOT OF PACF. FIRST DIFFERENCES OF THE LOG-TRANSFORMED SERIES

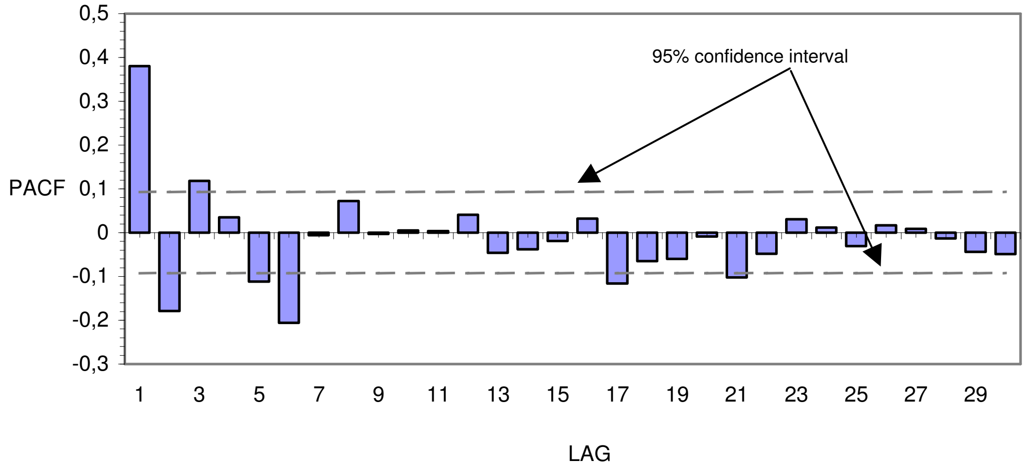

FIGURE 8: PLOT OF ACF. RESIDUALS OF THE 3-PARAMETER MODEL.

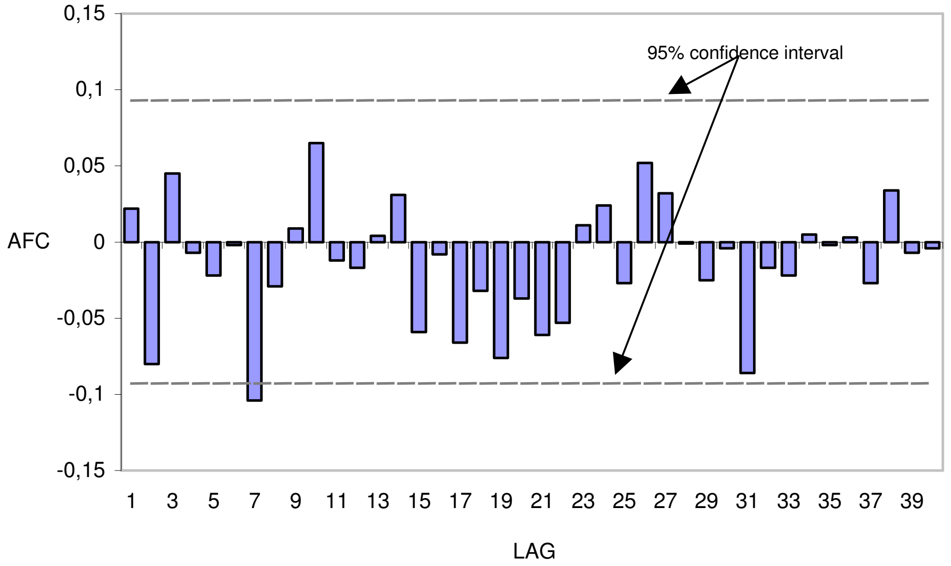

FIGURE 9: PLOT OF CPI FOR GREECE

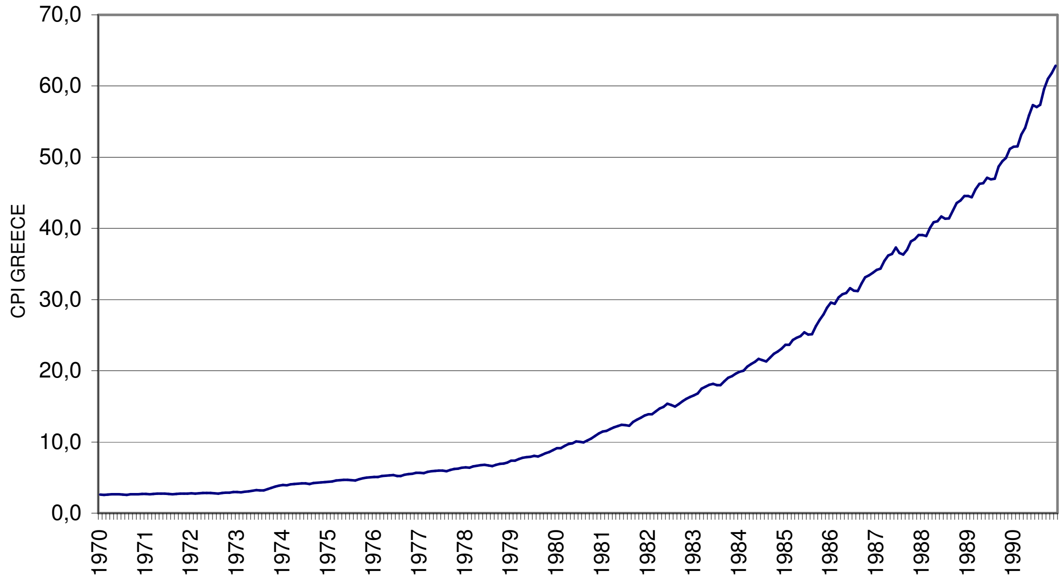

FIGURE 10: PLOT OF THE FIRST REGULAR AND SEASONAL DIFFERENCES OF CPI