Proceedings of ALGORITMY 2000
Conference on Scientific Computing, pp. 341 - 356

MODELING OF ECONOMIC TIME SERIES

JOZEF KOMORNÍK[(1)], MAGDA KOMORNÍKOVÁ[(2)]

Abstract. In this article we illustrate through several examples that business and economic time series tend to display such features as trends, seasonality, aberrant observations, conditionally varying variance, and/or non-linearity. Before we can construct sensible out-of-sample forecasts of such time series, each of these features can best be described by a time series model that is specifically designed for that purpose. In practice such models are usually combined into a single model that deals with more than one feature.

Introduction. In this article the focus is on key features of typical business and economic time series. The five key features of economic and business data can be (i) trends, (ii) seasonality, (iii) somehow influential data points, (iv) a variance that changes because of past observations (conditional heteroskedasticity), and (v) non-linearity. Typically, an economic time series displays at least one, but usually two or three of these features. To keep matters simple, however, in this article each series is analyzed for only one of these five features. In the modeling practice, the application of various possible models for each of these features is often combined.

Each of the about five data features will be illuminated using simple regression-based calculations. This should not imply that these models are also the best models we can use; they are merely helpful to show the properties of the data. Sometimes sophisticated models can better describe and forecast the observed features. A second important tool in this article is a graphical analysis. In most cases it is already quite helpful to put the data just in a graph with the values of the observations on the vertical axis and time on the horizontal axis. However, in case of, for example, many data or data with a large variance, it may sometimes be more insightful to rely on scatter or correlation plots. The latter shows the correlation between y_t and another variable x_t. Since time series analysis is our focus here, x_t is often replaced by, for example, y_{t-1}. In a practical situation, we would advise to construct each of the graphs as well as to quantify each of the simple regressions in order to obtain some overall insight in the specific data properties. Here, quantification means the estimation of the unknown parameters in the relevant regression.

In empirical time series analysis it is common practice to analyze data after the natural logarithmic transformation has been applied. Hence, if the raw data are denoted by w_t , we usually model and forecast $y_t = \log(w_t)$, where log denotes the natural logarithm. If forecasts are needed for w_t, it is usual to re-transform using

(1) Faculty of Management, Comenius University, Bratislava, Odbojárov 10, P.O.Box 1, Bratislava, (e-mail: **Jozef.Komornik@fm.uniba.sk**)
(2) Faculty of Civil Engineering STU, Radlinského 11, 813 68 Bratislava, ÚTIA AV ČR, Pod vodárenskou věží 4, 182 08 Prague, ČR, (e-mail: **magda@vox.svf.stuba.sk**)

$w_{t+h} = \exp(y_{t+h})$, preferably with some correction. One of the reasons for the logtransformation is that an exponential trend becomes linear. Although there are recent empirical and theoretical studies that question the validity of this automatism, for the moment we assume that y_t is the time series considered for modeling, and hence y_t usually denotes log (w_t). In cases where the data are already in relative format, as the unemployment rate, inflation rate or market share, the log transformation is usually not applied.

1. Trends. One of the dominant features of many economic and business time series is the trend. Such a trend can be upward or downward, it can be steep or not, and it can be exponential or approximately linear. Since a trend should definitively somehow be incorporated in a time series model, simply because it can be exploited for out-of-sample forecasting, an analysis of trend behaviour typically requires quite some research input.

To illustrate the presence of trends in economic data, consider the five graphs in figure 1, which are the annual indices of real national output (in logs) in China in five different sectors for the sample period 1952-1988. These sectors are agriculture, industry, construction, transportation, and commerce.

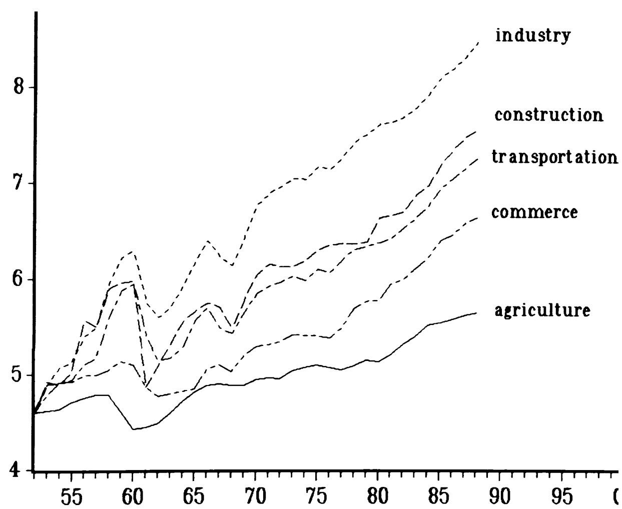

FIG. 1. *Annual indices of real national output in China, 1952 - 1988*

From this figure it can be observed that the five sectors have grown over the years at different growth rates, and also that the five sectors seem to have been affected by the likely, exogenous shocks to the Chinese economy around 1958 and 1968. These shocks roughly correspond to the two major political disturbances in China, that is, the Great Leap Forward around 1958 until 1962 and the Cultural Revolution from 1966 until 1976. It also appears from the graphs that these disturbances may not have affected each of the five sectors in a similar fashion. For example, the decline of the output in the construction sector in 1961 seems much larger than that in the industry sector in the same year. It also seems that the Great Leap Forward shock already had an impact on the output in the agriculture sector as early as 1959.

To quantify the trends in the five Chinese output series, consider the following simple regression model

$$y_t = \alpha + \delta t + u_t, \quad t = 1, 2, ..., n, \qquad (1)$$

where α and δ are unknown parameters and where u_t is an unknown residual error time series. Since the graphs in figure 1 show some cyclical behaviour, indicating that u_t may be correlated with u_{t-1}, and hence suggesting dynamic misspecification of (1), the standard errors in table 1 should be treated with care.

TABLE 1. *The significance of trends in real national output in China*

Variable:	δ in regression: $y_t = \alpha + \delta t + u_t$		μ in regression: $y_t - y_{t-1} = \mu + u_t$	
Agriculture	0.028	(0.002)	0.029	(0.011)
Industry	0.094	(0.003)	0.107	(0.028)
Construction	0.064	(0.004)	0.082	(0.042)
Transportation	0.060	(0.004)	0.074	(0.027)
Commerce	0.047	(0.003)	0.056	(0.016)

Note: The numbers in parentheses are estimated standard errors.

The left-hand panel of table 1 displays the estimates of δ in (1). It is clear that the upward trend is steepest for the industry sector (0.094) and that such growth is smallest for agriculture (0.028).

A trending pattern in economic data will be reflected by a significant average growth rate. Alternative to (1), we can therefore consider the regression

$$y_t - y_{t-1} = \mu + u_t, \quad t = 2, 3, ..., n. \qquad (2)$$

Notice that, just as in (1), we have effectively removed the trend in the regression (2), although now by differencing the data. A result of this operation is that there are now n - 1 observations that can be used to estimate μ. When it is possible to get rid of the trend by differencing the data, we say that y_t has a stochastic trend. If regression (1) is more adequate, we say that y_t has a deterministic trend. This illustrates that the concept of a trend is defined only in the context of a model.

The second panel of table 1 displays the estimates of μ in (2). It appears that μ parameters are slightly larger than the δ parameters in (1), although again the growth rate of the industry sector of about 10 percent by far exceeds the growth in the other sectors. Notice also that the estimated standard error of $\hat{\mu}$ in (2) is much larger than that of $\hat{\delta}$ in (1).

Changing trends. The trends in figure 1 are all of the familiar type, that is, many economic time series display an upward moving trend. It is however not necessary for a trend to move upwards to be called a trend. It is also possible that a trend is less smooth and may display slowly changing tendencies which once in a while change directions.

An example of such a trending pattern is given in figure 2 where the annual stock of motor cycles in The Netherlands is displayed, for 1946-1993. From 1946 to about 1960, there is an increasing trend due to the increasing popularity of motor cycles because of their successful use in World War II. From 1960 to 1973 the stock of motor cycles tends toward zero because of the increasing willingness to own a car. From 1974 onwards there is again a trend upwards, which in the last few years seems to explode. This can be attributed to the fact that car owners may want to have a motor cycle as an additional leisure vehicle. In sum, for this time series we can observe several gradual changes in the trend. Time series such as in figure 2 may have a so-called double stochastic trend, that is, the direction of trend itself is in turn a stochastic trend.

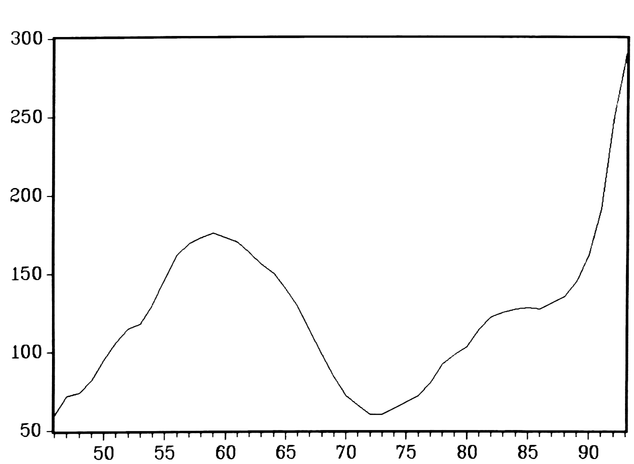

FIG. 2. *Annual stocks o motor cycles in The Netherlands, 1946 - 1993*

The motor cycles example in figure 2 highlights one possible approach to describing changing trends. An alternative method is allow the parameters in either (1) or (2) to change over time. Usually we consider such an approach in cases where it is reasonable to assume that certain exogenous shocks may have changed the direction of the trend. For example, the oil price shock in 1979.4 may have changed the direction of the trend in macroeconomic variables, see, e.g., Perron (1989). This can be illustrated by the (log of) industrial production index in the US (quarterly, seasonally adjusted), for which some regression results for (1) and (2) for different samples are presented in table 2.

TABLE 2. *The significance of trends in US industrial production*

Sample	$\hat{\delta}$ in regression: $y_t = \alpha + \delta t + u_t$		$\hat{\mu}$ in regression: $y_t - y_{t-1} = \mu + u_t$	
1960.1-1991.4	0.825	(0.016)	0.818	(0.167)
1960.1-1979.4	1.023	(0.030)	1.007	(0.216)
1980. I-1991.4	0.758	(0.041)	0.507	(0.256)

The first panel again considers regression (1), of which the parameters are estimated for the complete sample 1960.1-1991.4 and for the subsamples 1960.1 to 1979.4 and 1980.1 to 1991.4. Clearly, the estimated trend parameter $\hat{\delta}$ in the first sample (1.023) exceeds that of the second sample (0.758). A similar conclusion holds for regression (2) in the second panel of table 2 where growth in the period before the oil crisis is about 1 percent per quarter while it is only about 0.5 percent per quarter after that crisis. Again, the estimated standard error of $\hat{\mu}$ exceeds that of $\hat{\delta}$. The change in trend is visualized in figure 3.

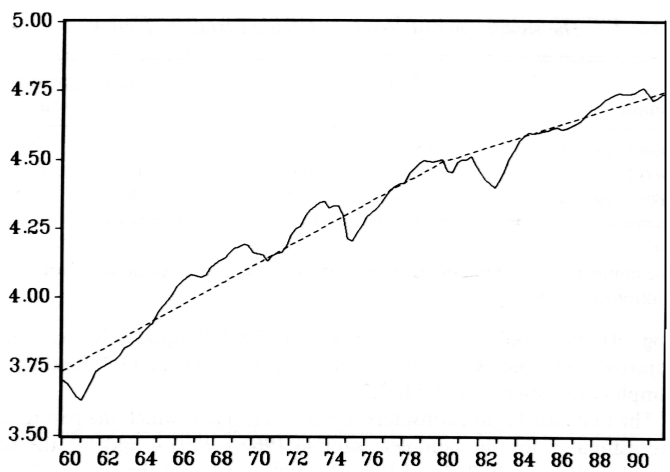

FIG. 3. *Quarterly index of industrial production in the US, 1960.1 – 1991.4 (seasonally adjusted), with changing deterministic trend after 1979.4*

A second example of changing trends is given in figure 4, which gives the monthly retail sales in The Netherlands. Until about 1979 we can observe a much steeper trend than after that period. A second feature of these data is the pronounced seasonal variation, a feature which will be discussed in the next subsection.

To conclude this introduction on trends in economic time series, it should be mentioned that there is no unique way to describe a trend. There are several different approaches, and each of these has a different impact on forecasting. When comparing figures 3 and 4 with figure 2, it is also clear that it may not be easy to make a proper choice between the versions of trend descriptions.

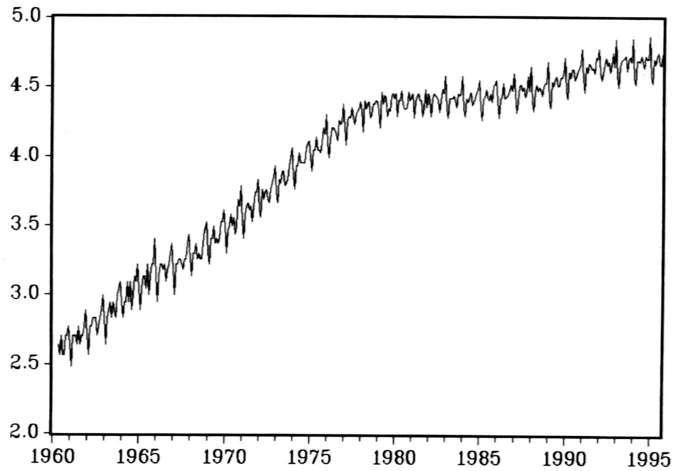

FIG. 4. *Monthly index of retail sales in The Netherlands, 1960.05 – 1995.09 (not seasonally adjusted)*

2. Seasonality. When economic time series are observed each month or quarter, it is often the case that such a series displays a seasonal pattern. Similar to the feature of a trend, where the definition of a trend depends on the model used to describe the trend, there does not exist a very precise definition of *seasonality*. Usually we refer to seasonality when observations in certain seasons display strikingly different features to those in other seasons. For example, when retail sales are always large in the fourth quarter (because of Christmas spending) as can be observed from figure 4 where the peaks occur in December, we can say that retail sales display seasonality. It may also be that seasonality is reflected in the variance of a time series. For example, for daily observed stock market returns the volatility seems often highest on Mondays, basically because investors have to digest three days of news instead of only one day.

The number of seasons is denoted by S. When a calendar year is considered the benchmark, this S equals 4 for quarterly data. In case there are N years of observations y_t, the number of observations n equals SN. Of course, other benchmarks can also be considered. For example, in empirical finance it may be relevant to consider daily observations with a week as the benchmark, that is, N refers to weeks and S = 5, or even higher frequency data like minutes within a day.

Seasonality is often noticeable rightaway from a simple graph of the time series, as in figure 4. In other cases, more experience is needed to spot the seasonal variation. For example, in figure 5, where $y_t - y_{t-1}$ of the logs of US industrial production index (seasonally unadjusted) is given, it may not be evident how relatively important seasonality is. On the other hand, it seems that the pattern after, say, 1975 is different from the pattern before that year. Notice also the dip in 1975.1, which seems an irregular observation.

In case simple graphs are not informative enough to highlight possible seasonal variation, we can rely on the following regression, which is a modification of (2)

$$y_t - y_{t-1} = \mu_1 D_{1,t} + \mu_2 D_{2,t} + ... + \mu_S D_{S,t} + u_t, \qquad t = 2, 3, ..., n, \tag{3}$$

where $D_{s,t}$ is a seasonal dummy variable with

$$D_{s,t} = \begin{cases} 1 & \text{when } t = (T - 1)S + s, \text{ with } s = 1, 2, ..., S \text{ and } T = 1, 2, ..., N \\ 0 & \text{otherwise} \end{cases} \tag{4}$$

In case the u_t process does not contain any information on seasonality, we may also consider the R^2 of (3) as giving an indication of the "amount of deterministic seasonality".

Given the tentative evidence in figure 5 of possibly changing patterns in the growth rates of US industrial production, it is often useful to fit a model like (3) for different subsamples. Here, S equals 4, and when the sample of 1960.1 to 1991.4 is divided into 1960.1-1975.4 and 1976.1 to 1991.4, the estimates of the parameters μ_1 to μ_4 are given in figure 6.

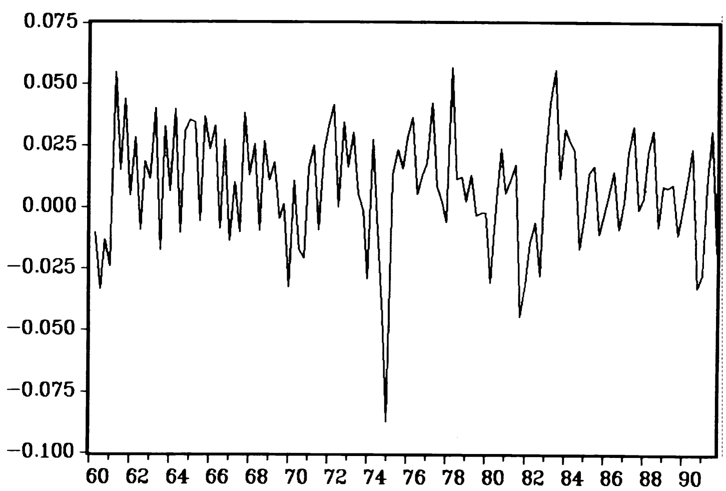

FIG. 5 *First differences of the (Iog of) the industrial production index in the US (not seasonally adjusted)*

The R^2 values for these two samples are 0.269 and 0.254, respectively. From figure 6, we can observe that the growth rates in quarters 2 and 4 are largest in the first subsample (0.027 and 0.016), while quarters 2 and 3 are important in the second sample (0.017 and 0.017). In other words, it seems that seasonality has changed in the sense that the fourth quarter observations have become less important in more recent years.

Another example of time series with pronounced seasonality is given in figure 7, where we depict (logs of the) four-weekly advertising expenditures on radio and television in The Netherlands for 1978.01-1994.13.

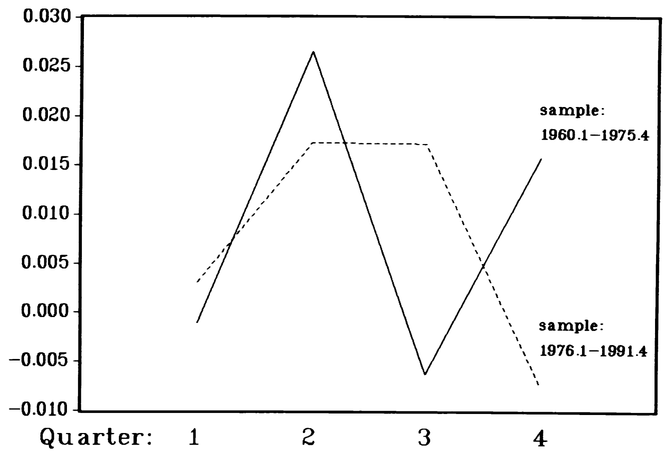

FIG..6 *Changing seasonality in US industrial production*

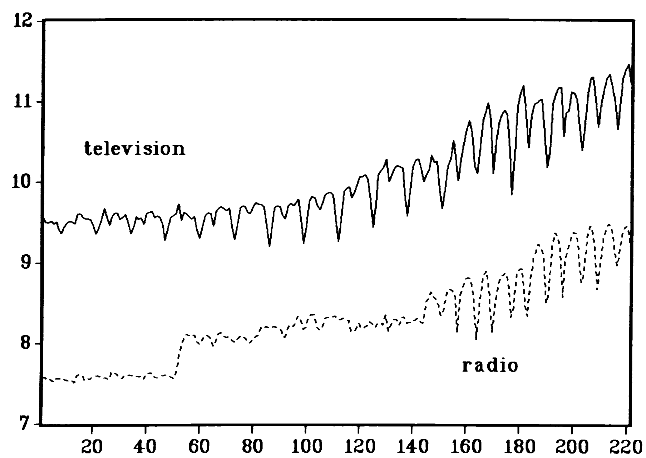

FIG. 7 *Four-weekly advertising expenditures on radio and television in The Neiherlands (1978.01-1994.13)*

For these two marketing time series it is clear that television advertising displays quite some seasonal fluctuation throughout the entire sample, where possibly there are some changes toward the end of the sample, and that radio advertising has seasonality only in about the last five years. This last period corresponds to the period where an additional commercial network (RTL-4) was introduced in The Netherlands. Furthermore, there seems to be a structural break in the radio series around observation 53. This break is related to an increase in radio broadcasting minutes in January 1982. Additionally, there is some visual evidence that the trend changes over time.

To investigate the effect of the introduction of RTL-4 on seasonality in television advertising, consider again the results of the regression in (3) for two subsamples, which are depicted in figure 8.

It is clear from figure 8 that seasonality seems to increase in the second subsample since then the estimates of μ_s, s = 1, 2, ..., 13 show more volatility.

Generally, it appears that many seasonally observed business and economic time series display seasonality in the sense that the observations in certain seasons have properties that differ from those data points in other seasons. A second feature of many economic time series is that seasonality changes over time.

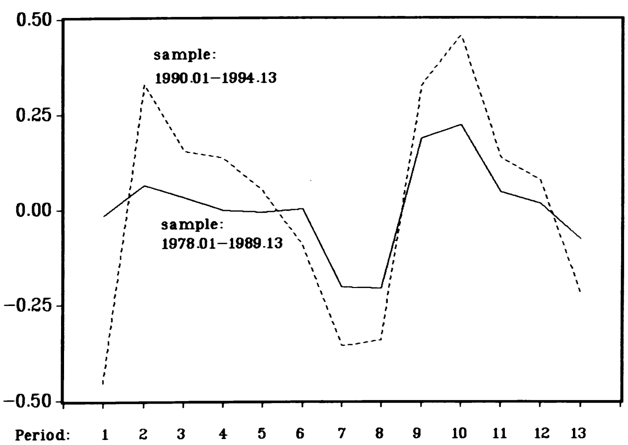

FIG. 8 *Changing seasonality in television advertising expenditures*

3. Aberrant observations. The radio advertising series in figure 7 indicates that time series may be subject to so-called regime shifts. The mean of the observations in the first part of the sample is smaller than that in the second part of the sample, and in the third part there even emerges a seasonal pattern that was previously absent. Obviously, account should be taken of such changing patterns when forecasting out of sample. For example, if we were to analyze the entire sample of the radio advertising series with a model that does not allow for seasonality, it is likely that inaccurate forecasts would be made for certain seasons. Furthermore, when the mean shift around observation 53 is neglected, we would also expect a systematic bias in forecasts. This implies that account should be taken for the possibility that there are periods or subsamples that can make time series modeling difficult.

Possibly distorting observations do not necessarily come in a sequence as in the radio advertising example. It may also be that only single observations have a major impact on time series modeling and forecasting. Such data points are called aberrant observatíons. As an illustration, consider the differenced y_t, that is, $y_t - y_{t-1}$, where $y_t = \log(w_t)$, with w_t the price level, and the inflation rate $\frac{w_t - w_{t-1}}{w_{t-1}}$ in Argentina, for the sample 1970.1-1989.4 in figure 9.

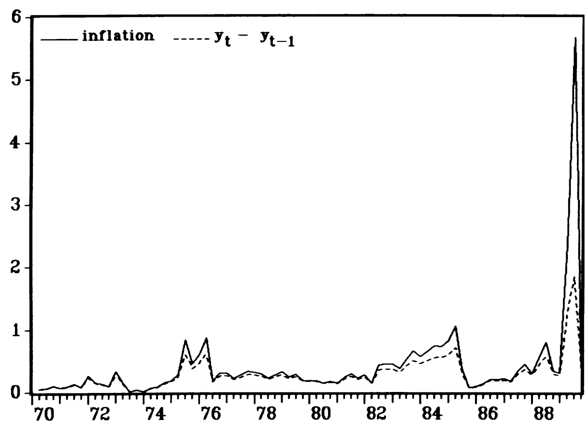

FIG.9 *First differences in log prices versus the inflation rate: the case of Argentina, 1970.1-1989.4*

The first obvious feature of these graphs is that in the case where the quarterly inflation rate is high (as is the case in 1989.3 where it is about 500 percent), the differenced y_t series is not a good approximation to the inflation rate (since the 1989.3 observation would now correspond to about 200 percent). For convenience we analyze $y_t - y_{t-1}$ in the sequel, since it smooths the variance of the time series, although of course preference could be given to modeling the inflation rate itself.

A second obvious feature of the graphs in figure 9 is that the data in 1989 seem to be quite different from those observations the years before. In fact, if there is any correlation between $y_t - y_{t-1}$, and $y_{t-1} - y_{t-2}$ such a correlation may be affected by these observations. In other words, when considering the following simple regression

$$y_t - y_{t-1} = \alpha + \rho\,(y_{t-1} - y_{t-2}) + u_t\,, \quad t = 3, 4, ..., n \tag{5}$$

we would expect that an estimate of ρ is influenced by the data points in the last year. It should be mentioned that a model as in (5) is called a first order autoregressive model AR(1) for $y_t - y_{t-1}$. It reflects the possible occurrence that this quarter's price change is likely to be correlated with the price change in the previous quarter.

When the model in (5) is quantified for n - 2 observations (which is 78 here since the estimation sample runs from 1970.3 to 1989.4), the estimated value of ρ equals $\hat{\rho}$ = 0.561 (0.094), where the estimated standard error is in parentheses. This regression result is summarized in figure 10, where the slope of the regression line equals 0.561.

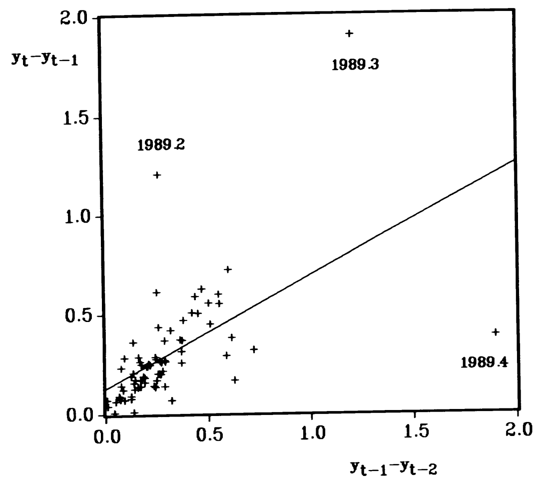

FIG.10 *Regression of first differences in log prices in Argentina at time t on differences at time t-1: al1 observations*

As is obvious from this graph, there are at least three observations that are far away from the regression line. These three observations are 1989.2, 1989.3, and 1989.4. At first sight, it is difficult to get a clear picture of how each of these observations affects the regression line and hence $\hat{\rho}$. Therefore, ρ in (5) is estimated again, while now deleting the 1989 observations. The result is displayed in figure 11.

It seems that the regression line is somewhat steeper than that in figure 10. Indeed, the ρ is now estimated as $\hat{\rho}$ = 0.704 (0.082). Even though this value may not differ significantly from the $\hat{\rho}$ for all data, we may expect the forecasts from both models to differ. The fact that in figure 11 all observations seem to correspond more to the regression line as they are in figure 10 is reflected by the smaller standard error (0.082 versus 0.094). In sum, it appears that only a few observations can affect the correlation in a time series y_t, and also the precision of an estimate of such a correlation.

4. Conditional heteroskedasticity. A fourth feature of economic time series, and in particular of financial time series, is that aberrant observations tend to emerge in clusters. Intuitively, if on a certain day news arrives on a stock market, we may react by selling or buying many stocks, while the day after the news has been digested and valued more properly, we may wish to return to the behaviour of before the arrival of the news. This pattern would be reflected by (a possibly large) increase or decrease in the returns on one day followed by an opposite change on the next day. In a sense, we can interpret these two observations as two aberrant observations in a row. On the other hand, we may assume that the second sharp increase or decrease in returns is caused by the first, and hence that the two sudden changes in returns are correlated.

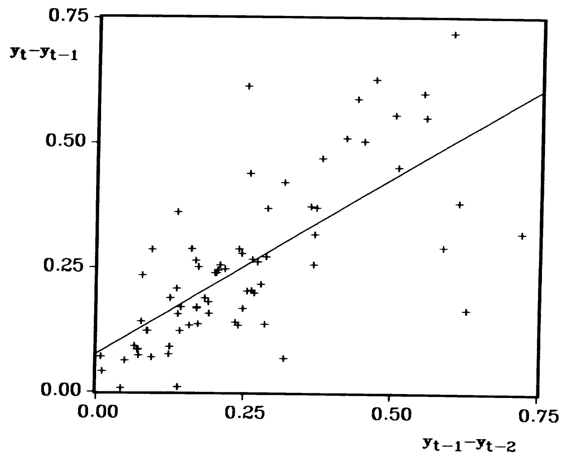

FIG.11 *Regression of first differences in log prices in Argentina at time t on differences at time t - 1: all observations without 1989.1-1989.4*

Consider for example the 770 consecutive returns on the Dow-Jones index, from week 1 in 1980 to week 39 in 1994, as they are given in figure 12.

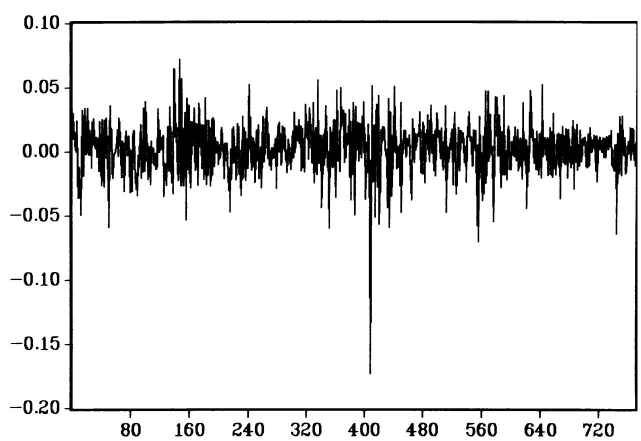

FIG. 12 *Returns of the weekly Dow-Jones index (week 1 in 1980 to week 39 in 1994)*

These returns concern the end-of the-week returns, where a week is assumed to run from Thursday to Wednesday. Observation 408 corresponds to the week commencing Monday, October 19, 1987. Around that so-called Black Monday, the returns on the Dow-Jones index decreased by about 17 percent. Immediately after that observation, we can find several data points that are large in absolute value. Additionally, in other parts of the sample, we can observe "bubbles," i.e., clusters of observations with large variances. This phenomenon is called volatility clustering or, otherwise, conditional heteroskedasticity.

Allowing for the possibility that high volatility is followed by high volatility, and that low volatility will be followed by low volatility, where volatility is defined in terms of the returns themselves, we can consider the presence of so-called conditional heteroskedasticity. This time series feature can be put forward by the simple regression

$$(y_t - y_{t-1})^2 = \alpha + \rho\, (y_{t-1} - y_{t-2})^2 + u_t, \quad t = 3, 4, ..., n, \tag{6}$$

where $(y_{t-1} - y_{t-2})^2$ represents the variance of the returns. For the sample of 768 data points (770 minus 2 starting values), we obtain $\hat{\rho}$ = 0.292 (0.035), where the estimated standard error is given in parentheses. This suggests that there is positive correlation between the variances of the returns. Additionally, when $(y_{t-1} - y_{t-2})^2$ is replaced by $|y_t - y_{t-1}|$, the absolute value of the return, the regression in (6) yields $\hat{\rho}$=0.213 (0.035). This result also points toward the presence of volatility clustering.

In case of volatility clustering, we may wish to exploit this in order to forecast future volatility. Since this variable is a measure of risk, such forecasts can be useful to evaluate investment strategies. Furthermore, it can be useful for decisions on buying or selling options or other derivatives.

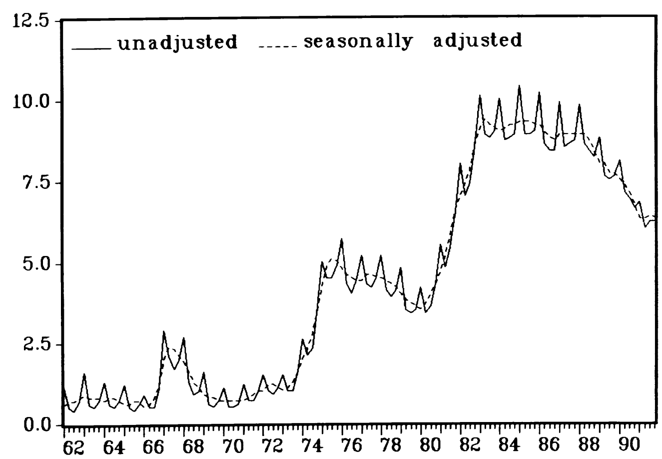

FIG. 13 *Quarterly unemployment rate in Germany, 1962.1-1991.4 (seasonally adjusted and not seasonally adjusted)*

5. Non-linearity. The fifth and final feature of several economic time series that is treated in this book is non-linearity. Although the best definition of non-linearity may perhaps be "everything other than linearity," for economic time series we usually rely on concepts such as regime switches or so-called state dependency. The latter indicates that the behaviour of a time series is different depending on its current state. Such a state may be a volatile period for a financial market, as discussed in the previous section, or it may be a recession. When the behaviour of a time series is different across such discrete states, we can also call this regime-switching behaviour. It should be mentioned that such a state (as a recession) may concern only a few observations. Furthermore, in case there are regime switches, there must also be so-called thresholds, which mark the differences between the regimes. Finally, notice that regime switching can show similarity with structural breaks. For example, the oil crisis in quarter 4 of 1979 may have changed the trend in US industrial production, see figure 3.

Non-linear behavior is often quite obvious for certain macroeconomic time series in case the growth rate ($y_t - y_{t-1}$) or the change ($w_t - w_{t-1}$) takes different average values across states. Consider for example the unemployment rate in Germany for 1962.1 to 1991.4 in figure 13.

Clearly, unemployment sometimes rises quite rapidly, usually in the recession years 1967, 1974-1975, and 1980-1982, while it decreases very slowly, usually in times of expansions. This asymmetry can be formalized by estimating the parameters in the following simple regression

$$y_t - y_{t-1} = \mu_1 I_t \text{ [expansion year]} + \mu_2 I_t \text{ [recession year]} + u_t,\ t = 2, 3, \dots n, \qquad (7)$$

where I_t[.] is an indicator variable, which allows the absolute value of the rate of change to vary across the two states, say, "decreasing y_t" and "increasing y_t" from μ_1 to μ_2, where μ_1 may be different from μ_2. For the German seasonally adjusted unemployment rate, we find that $\hat{\mu}_1$ =-0.040 and $\hat{\mu}_2$ = 0.388, indicating that when the unemployment rate increases (in recessions), it rises faster than when it goes down (in expansions). This regression result seems to confirm the visual impressions from figure 13.

When a non-linear time series model is used to describe and forecast an economic variable, much more effort is usually needed to specify the model and to estimate its parameters. A key reason is that this option leads to a wide variety of possible models. In practice therefore we start with a diagnostic measure to obtain a first and tentative impression of what type of non-linear model could be useful.

6. Common features. Many time series in economics and business usually have at least one of the above five features. For example, the graphs for the German unemployment rate in figure 13 show that this variable displays seasonality, non-linearity, and possibly also some aberrant observations. As another example, the advertising expenditure series in figure 7 seems to display structural breaks, changing seasonality, and trends. Finally, the Dow-Jones index may be heteroskedastic, but there may also be aberrant observations (such as perhaps the week around Black Monday).

In case of univariate time series modeling, that is, when we only analyze and forecast y_t given its own past pattern, it is important to take account of all of the observed features, which are illuminated by simple auxiliary regressions and specific insightful graphs. In case we want to incorporate past patterns of variables as, say, x_t and z_t to describe and forecast y_t, and possibly to use the past of y_t to describe and forecast x_t and z_t, the so-called multivariate time series models are considered.

Given all the possible features and also given the possible wish to use multivariate models, it is likely that many decisions will have to be made. One way to reduce the number of decisions is to search for the presence of common features across univariate time series. Additionally, such common features may themselves be the focus of interest. For example, a relevant question for the output in the five sectors in China, as displayed in figure 1, is whether these have common growth or whether the industry sector significantly outgrows the other four sectors. For the same five variables, we may wonder whether the transportation and commerce sectors have in fact the same trend.

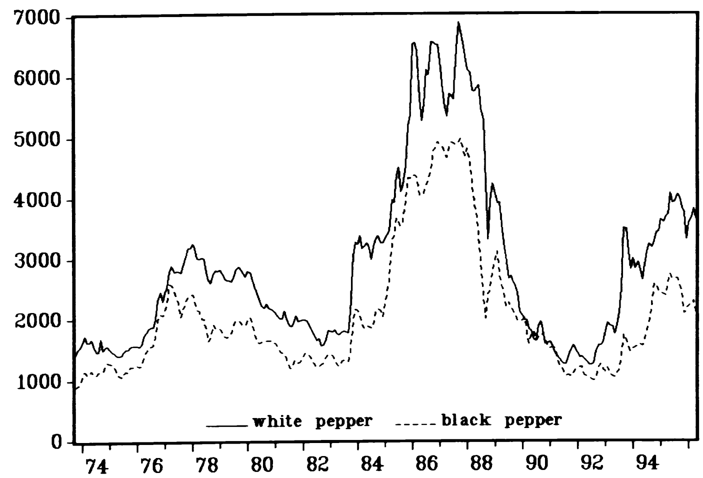

FIG. 14 *Monthly prices (no logs) of white and black pepper, 1973.10-1996.04*

Figure 14 displays the monthly prices of black and white pepper over more than two decades. Given that these two products can be substitutes, it would be reasonable to expect the prices of these commodities to move together over time. Because of international pricing agreements, temporary overproduction, or failure of crops, prices can display trending behaviour, although the direction of these trends may not be constant over time. This expectation is reflected by the graphs in figure 17 where pepper undergoes some periods with increasing prices, while also (as for example between 1977 and 1982) the prices tend downwards for longer periods. Furthermore, there are sequences of months with very large increases or decreases in prices. Hence, as well as trends pepper prices display volatility clustering. The graphs in figure 14 clearly suggest common patterns across pepper prices, such as common trends and common conditional volatility.

A final example of possibly common features is given by the series of advertising expenditures on radio and television in figure 9. Since the launch of a new commercial network in The Netherlands (in October 1989, which is about observation 154), it is noticeable that radio and television advertising seem to display common seasonality. This common feature is visualized in figure 15, where the growth rates y_t - y_{t-1} of each of these series for the sample 157 to 221 (that is, the last five years) are presented.

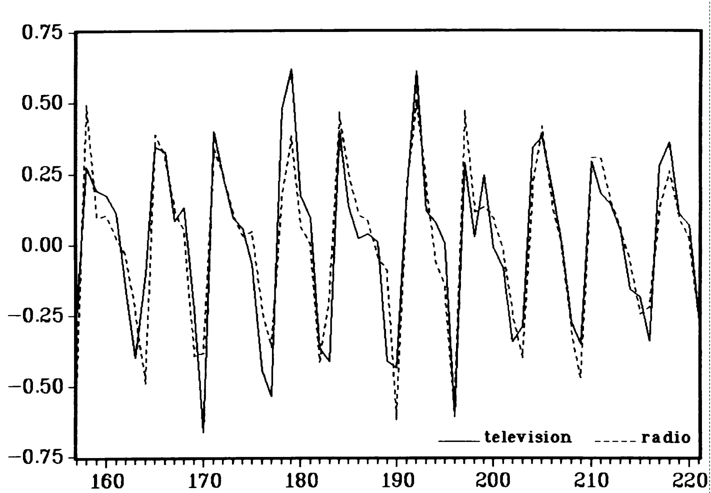

FIG.15 *Common seasonality in radio and television advertising expenditures 1990.01-1994.13*

Except for a few observations, it seems that the growth rates of both series match quite well. The simple framework that is often used to investigate common features is

$$y_t = \beta\, x_t + u_t, \tag{8}$$

where y_t and x_t have a certain feature and the residual time series u_t does not have this feature. For example, y_t and x_t may both have a trend, while u_t does not show trending behaviour.

At present, the research area concerning common features in multivariate models is in rapid development. To what extent these methods appear useful is yet to be established in empirical and simulation studies.

Acknowledgements. The work on this paper was supported by grant GA ČR 402/99/0032 and VEGA 1/7146/20, 2/6087/99.

REFERENCES

[1] Box, G.E.P., Jenkins, G. M. and Reinsel, G. C. (1994), Time Series Analysis, Forecasting and Control (third edition), Englewood Cliffs: Prentice Hall.

[2] Franses, P. H. (1998), Time series models for business and economic forecasting. Cambridge University Press.

[3] Komorník, M. Komorníková: Modelling Multivariate Financial Time Series Using Mathematica. ALGORITMY′97, Zuberec, September 1997, pp. 89 – 100