


ELEMENTS OF TIME SERIES ECONOMETRICS: AN APPLIED APPROACH

EVŽEN KOČENDA
ALEXANDR ČERNÝ

KAROLINUM



Elements of Time Series Econometrics: An Applied Approach

Evžen Kočenda, Alexandr Černý

KAROLINUM PRESS

Karolinum Press is the publishing department of Charles University in Prague

Copyright © 2015 by Charles University in Prague

Text © 2015 by Evžen Kočenda and Alexandr Černý

Photos © 2015 by Monika Kočendová

Third edition

Designed by Jan Šerých

Typeset by Studio Lacerta (www.lacerta.cz)

ISBN 978-80-246-3199-8

ISBN 978-80-246-3198-1 (pdf)

The manuscript was reviewed by Jan Kódera (University of Economics, Prague) and Miloslav Vošvrda (Institute of Information Theory and Automation, The Czech Academy of Sciences)



Univerzita Karlova v Praze
Nakladatelství Karolinum 2015

www.karolinum.cz
ebooks@karolinum.cz

CONTENTS

INTRODUCTION	-----	11
1. THE NATURE OF TIME SERIES	-----	15
1.1 DESCRIPTION OF TIME SERIES	-----	16
1.2 WHITE NOISE	-----	17
1.3 STATIONARITY	-----	17
1.4 TRANSFORMATIONS OF TIME SERIES	-----	19
1.5 TREND, SEASONAL, AND IRREGULAR PATTERNS	-----	21
1.6 ARMA MODELS OF TIME SERIES	-----	22
1.7 STYLIZED FACTS ABOUT TIME SERIES	-----	23
2. DIFFERENCE EQUATIONS	-----	27
2.1 LINEAR DIFFERENCE EQUATIONS	-----	27
2.2 LAG OPERATOR	-----	28
2.3 THE SOLUTION OF DIFFERENCE EQUATIONS	-----	29
2.3.1 PARTICULAR SOLUTION AND LAG OPERATORS	-----	30
2.3.2 SOLUTION BY ITERATION	-----	31
2.3.3 HOMOGENOUS SOLUTION	-----	33
2.3.4 PARTICULAR SOLUTION	-----	34
2.4 STABILITY CONDITIONS	-----	35
2.5 STABILITY AND STATIONARITY	-----	36

3.	UNIVARIATE TIME SERIES	41
3.1	ESTIMATION OF AN ARMA MODEL	41
3.1.1	AUTOCORRELATION FUNCTION – ACF	42
3.1.2	PARTIAL AUTOCORRELATION FUNCTION – PACF	45
3.1.3	Q-TESTS	48
3.1.4	DIAGNOSTICS OF RESIDUALS	50
3.1.5	INFORMATION CRITERIA	50
3.1.6	BOX-JENKINS METHODOLOGY	52
3.2	TREND IN TIME SERIES	54
3.2.1	DETERMINISTIC TREND	54
3.2.2	STOCHASTIC TREND	55
3.2.3	STOCHASTIC PLUS DETERMINISTIC TREND	56
3.2.4	ADDITIONAL NOTES ON TRENDS IN TIME SERIES	57
3.3	SEASONALITY IN TIME SERIES	58
3.3.1	REMOVING SEASONAL PATTERNS	59
3.3.2	ESTIMATING SEASONAL PATTERNS	60
3.3.3	DETECTING SEASONAL PATTERNS	62
3.3.4	HODRICK-PRESCOTT FILTER	62
3.4	UNIT ROOTS	64
3.4.1	DICKEY-FULLER TEST	66
3.4.2	AUGMENTED DICKEY-FULLER TEST	68
3.4.3	PHILLIPS-PERRON TEST	71
3.4.4	SHORTCOMINGS OF THE STANDARD UNIT ROOT TEST	72
3.4.5	KPSS TEST	73
3.5	UNIT ROOTS AND STRUCTURAL CHANGE	76
3.5.1	PERRON'S TEST	77
3.5.2	ZIVOT AND ANDREWS' TEST	81
3.6	DETECTING A STRUCTURAL CHANGE	84
3.6.1	SINGLE STRUCTURAL CHANGE	85
3.6.2	MULTIPLE STRUCTURAL CHANGE	91
3.7	CONDITIONAL HETEROSKEDASTICITY AND NON-LINEAR STRUCTURE	98
3.7.1	CONDITIONAL AND UNCONDITIONAL EXPECTATIONS	100
3.7.2	ARCH MODEL	101
3.7.3	GARCH MODEL	104
3.7.4	DETECTING CONDITIONAL HETEROSKEDASTICITY	107
3.7.5	THE BDS TEST	110
3.7.6	AN ALTERNATIVE TO THE BDS TEST: INTEGRATION ACROSS THE CORRELATION INTEGRAL	114
3.7.7	IDENTIFICATION AND ESTIMATION OF A GARCH MODEL	118
3.7.8	EXTENSIONS OF ARCH -TYPE MODELS	122

3.7.9	MULTIVARIATE (G)ARCH MODELS	129
3.7.10	STRUCTURAL BREAKS IN VOLATILITY	135
4.	MULTIPLE TIME SERIES	141
4.1	VAR MODELS	143
4.1.1	STRUCTURAL FORM, REDUCED FORM, AND IDENTIFICATION	144
4.1.2	STABILITY AND STATIONARITY OF VAR MODELS	145
4.1.3	ESTIMATION OF A VAR MODEL	148
4.2	GRANGER CAUSALITY	150
4.3	COINTEGRATION AND ERROR CORRECTION MODELS	154
4.3.1	DEFINITION OF COINTEGRATION	157
4.3.2	THE ENGLE-GRANGER METHODOLOGY	162
4.3.3	EXTENSIONS TO THE ENGLE-GRANGER METHODOLOGY	168
4.3.4	THE JOHANSEN METHODOLOGY	170
5.	PANEL DATA AND UNIT ROOT TESTS	177
5.1	LEVIN, LIN, AND CHU PANEL UNIT-ROOT TEST WITH A NULL OF UNIT ROOT AND LIMITED COEFFICIENTS HETEROGENEITY	178
5.2	IM, PESARAN, AND SHIN UNIT-ROOT TEST WITH A NULL OF UNIT ROOT AND HETEROGENEOUS COEFFICIENTS	181
5.3	HADRI UNIT-ROOT TESTS WITH A NULL OF STATIONARITY	184
5.4	BREUER, MCNOWN, AND WALLACE TEST FOR CONVERGENCE	186
5.5	VOGELSANG TEST FOR β -CONVERGENCE	187
	APPENDIX A: MONTE CARLO SIMULATIONS	193
	APPENDIX B: STATISTICAL TABLES	196
	REFERENCES	202
	INDEX	216

To my mother
(EK)

To Alžběta
(AČ)

INTRODUCTION

This book, in its third edition, presents the numerous tools for the econometric analysis of time series. The text is designed so that it can be used for a semester course on time series econometrics, but by no means is the text meant to be exhaustive on the topic. The major emphasis of the text is on the practical application of theoretical tools. Accordingly, we aim to present material in a way that is easy to understand and we abstract from the rigorous style of theorems and proofs.¹ In many cases we offer an intuitive explanation and understanding of the studied phenomena. Essential concepts are illustrated by clear-cut examples. Readers interested in a more formal approach are advised to consult the appropriate references cited throughout the text.²

Many sections of the book refer to influential papers where specific techniques originally appeared. Additionally, we draw the attention of readers to numerous applied works where the use of specific techniques is best illustrated because applications offer a better understanding of the presented techniques. Such applications are chiefly connected with issues of recent economic transition and European integration, and this way we also bring forth the evidence that applied econometric research offers with respect to both of these recent phenomena. The outlined style of presentation makes the book also a rich source of references.


The text is divided into five major sections. The first section, “The Nature of Time Series”, gives an introduction to time series analysis. The second section, “Difference Equations”, describes briefly the theory of difference equations with an emphasis on results that are important for time series econometrics. The third section, “Univariate Time Series”, presents the methods commonly used in univariate time series analysis, the analysis of time series of one single variable. The fourth section, “Multiple Time Series”, deals with time series models of multiple interrelated variables. The fifth section “Panel Data and Unit Root Tests”, deals with methods known as panel unit root tests that are relevant to issues of convergence. Appendices contain an introduction to simulation techniques and statistical tables.

Photographs, and illustrations based on them, that appear throughout the book are to underline the purpose of the tools described in the book. Photographs, taken by Monika Kočendová, show details and sections of Fresnel lenses used in lighthouses to collimate light into parallel rays so that the light is visible to large distances and guides ships. Tools described in this book are used to process information available in data to deliver results guiding our decisions.

1 For rigorous treatment of specific issues Greene (2008) is recommended.

2 Patterson (2000) and Enders (2009) can serve as additional references that deal specifically with time series analysis.

When working on the text we received valuable help from many people and we would like to thank them all. In particular we are grateful for the research assistance provided by Ľuboš Briatka, Juraj Stančík (first edition), and Branka Marković (second edition). We are grateful to the staff of the Karolinum Press. We also thank Professor Jan Kmenta for consenting to our use of the title of his book (Kmenta, 1986) as a part of our title. Special thanks go to Monika (EK) and Alžběta (AČ).



1.

THE NATURE OF TIME SERIES

There are two major types of data sets studied by econometrics: cross-sectional data and time series. Cross-sectional data sets are data collected at one given time across multiple entities such as countries, industries, and companies. A time series is any set of data ordered by time. As our lives pass in time, it is natural for a variable to become a time series. Any variable that registers periodically forms a time series. For example, a yearly gross domestic product (GDP) recorded over several years is a time series. Similarly price level, unemployment, exchange rates of a currency, or profits of a firm can form a time series, if recorded periodically over certain time span. The combination of cross-sectional data and time series creates what economists call a panel data set. Panel data sets can be studied by tools characteristic for panel data econometrics or by tools characteristic for multiple time series analysis.

The fact that time series data are ordered by time implies some of their special properties and also some specific approaches to their analysis. For example, the time ordering enables the estimation of models built upon one variable only – so-called univariate time series models. In such a case a variable is estimated as a function of its past values (lags) and eventually time trends as well. As the variable is regressed on its own past values, such specification is aptly called an autoregressive process, abbreviated as “AR”. Because of the time ordering of data, issues of *autocorrelation* gain prominent importance in time series econometrics.

1.1 DESCRIPTION OF TIME SERIES

A set of data ordered by time forms a time series, $\{y_t\}_{t=1}^T$. We use the term “time series” for three distinct but closely related objects: a series of random variables, a series of data that are concrete realizations of these variables, and also for the stochastic process that generates these data or random variables.

Example 1.1 The stochastic process that generates the time series can be, for example, described as a simple autoregressive process with one lag: $y_t = 0.5y_{t-1} + \varepsilon_t$ (AR(1) process), where ε_t are normal iid with mean 0 and variance σ^2 , which is some positive number. With the initial condition $y_0 = 0$, the sequence of random variables generated by this process is $\varepsilon_1, 0.5\varepsilon_1 + \varepsilon_2, 0.25\varepsilon_1 + 0.5\varepsilon_2 + \varepsilon_3$, etc. Finally, the concrete realizations of these random variables can be the numbers 0.13882, 0.034936, -1.69767, etc. When we say that we estimate a time series, it means that based on the data (concrete realizations) we estimate the underlying process that generated the time series. The specification of the process is also called the model.

The properties *frequency*, *time span*, *mean*, *variance*, and *covariance* are used to give a basic description of time series.

1. *Frequency* is related to the time elapsed between y_t and y_{t+1} . Data can be collected with yearly, quarterly, daily, or even greater frequency. In case of a greater-than-daily frequency we speak about *intra-day data*. For example stock prices may be recorded in minute or even second intervals. The term “frequency” is actually used incorrectly in the context of time series econometrics. When we say daily frequency of the data, we mean in fact that there is one data point recorded per day.
2. *Time span* is the period of time over which the data were collected. If there are no gaps in the data, the time span is equivalent to the number of observations times the frequency. Throughout the text T is reserved to indicate the sample size (the number of observations) unless stated otherwise.
3. The *mean* μ_t is defined as $\mu_t = E(y_t)$. The mean is defined for each element of the time series, so that with T observations there are T means defined.
4. The *variance* is defined as $var(y_t) = E[(y_t - \mu_t)^2]$. Similarly as with the mean, the variance is defined for each element of the time series.
5. The *covariance* is defined as $cov(y_t, y_{t-s}) = E[(y_t - \mu_t)(y_{t-s} - \mu_{t-s})]$. The covariance is defined for each time t and for each time difference s , so that in the general case there are $T^2 - T$ covariances defined; however, because of symmetry only half of them are different.

1.2 WHITE NOISE

White noise is a term frequently used in time series econometrics. As the name suggests, white noise is a time series that does not contain any information that would help in estimation (except its variance and higher moments). Residuals from a correctly specified or “true” model that captures fully the data generating process are white noise. In the text, the white noise error process will be usually denoted as ε_t . For example a *series of identically and independently distributed random variables with 0 mean is white noise*.

When we estimate a time series using a correct model as described in sections 1.6 and 3.1 then the remaining inestimable part of the time series (the errors, or residuals) must be white noise. Procedures used to test if a time series is white noise are described in section 3.1.4.

1.3 STATIONARITY

Stationarity is a crucial property of time series. If a time series is stationary, then any shock that occurs in time t has a diminishing effect over time and finally disappears in time $t + s$ as $s \rightarrow \infty$. This feature is called *mean reversion*. With a non-stationary time series this is not the case and the effect of a shock either remains present in the same magnitude in all future dates or can be considered as a source behind the “explosion” of the series over time. If the former is the case, then the time series was generated by the so-called *unit root* process. Unit root processes form a special subset of non-stationary processes. Being on the edge between stationary and non-stationary processes, unit root processes play a particularly important role in time series analysis. For more details on stationarity, non-stationarity, and unit root processes see sections 2.4, 2.5, 3.4, and 3.5.

The most frequently used stationarity concept in econometrics is the concept of *covariance stationarity*. Throughout the text, we will for simplicity usually use only the term *stationarity* instead of *covariance stationarity*. We say that a time series $\{y_t\}_{t=1}^T$ is *covariance stationary* if and only if the following formal conditions are satisfied:

1. $\mu_t = \mu_{t-s} = \mu < \infty$ for all t, s .
2. $\text{var}(y_t) = \text{var}(y_{t-s}) = \sigma^2 < \infty$ for all t, s .
3. $\text{cov}(y_t, y_{t-s}) = \text{cov}(y_{t-j}, y_{t-j-s}) = \gamma_s < \infty$ for all t, j , and s .

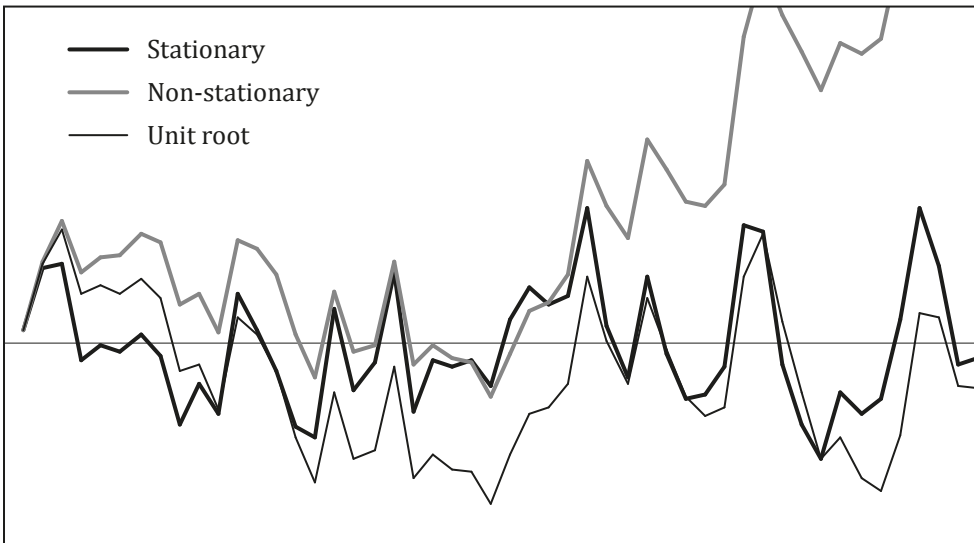
Translated into plain language the above means that a time series is covariance stationary, if its mean and variance are constant and finite over time and if the

covariance depends only on the time distance s between the two elements of the time series but not on the time t itself.

Note that any white noise time series is obviously stationary. However, a stationary time series is not automatically white noise. For white noise we need additional conditions that the mean and all covariances are 0; e.g. $\mu = 0$ and $\gamma_s = 0$ for all s .

Most economic time series are not stationary and specific transformations are needed in order to achieve stationarity. Some useful transformations are described in the next section.

Figure 1.1: Stationary, non-stationary, and unit root time series: a comparison.



Example 1.2 Figure 1.1 shows examples of stationary, non-stationary, and unit root time series. All three series were generated by a simple autoregressive process with one lag, an $AR(1)$ process defined as $y_t = a_1 y_{t-1} + \varepsilon_t$, where ε_t are normal iid with zero mean and variance $\sigma^2 = 9$. With such an $AR(1)$ process, the necessary and sufficient condition for stationarity is $|a_1| < 1$. If $|a_1| \geq 1$, then the time series is non-stationary, with $|a_1| > 1$ it explodes and with $|a_1| = 1$ it contains a unit root. The formal necessary and sufficient conditions for time series stationarity will be described in sections 2.4 and 2.5. The three time series in the figure were generated by the following processes:

stationary: $y_t = 0.6y_{t-1} + \varepsilon_t$, $a_1 = 0.6 < 1$,

non-stationary: $y_t = 1.1y_{t-1} + \varepsilon_t$, $a_1 = 1.1 > 1$,

unit root: $y_t = y_{t-1} + \varepsilon_t$, $a_1 = 1$.

We can distinguish clear visual differences between the three time series. The stationary time series tends to return often to its initial value. The non-stationary time series explodes after a while. Finally, the time series containing a unit root can resemble a stationary time series, but it does not return to its initial value as often. These differences can be more or less pronounced on a visual plot. Nevertheless, a visual plot cannot replace the formal stationarity tests described in sections 3.4, 3.5, and 5.

1.4 TRANSFORMATIONS OF TIME SERIES

In most cases some transformations of time series of economic data are necessary before we can proceed with estimation. Usually we apply transformations in order to achieve stationarity. However, sometimes it is natural to apply transformations because the transformed variable corresponds to what we are actually interested in. A typical example is a macroeconomic variable in levels versus the variable's growth rate. An example of this is prices versus inflation. If we are interested in analyzing inflation, then we want to transform prices (in levels) into inflation first. Achieving stationarity through the transformation is an extra benefit.

Depending on the type of transformation that we must apply in order to make a time series stationary, we can make a basic distinction of the time series into *difference stationary*, *trend stationary*, and *broken trend stationary* series. *Difference stationary* time series become stationary after differencing, *trend stationary* series after detrending, and *broken trend stationary* series after detrending with a structural change incorporated (more on broken trend stationary series will be introduced in section 3.5).

If a series must be differenced n times to become stationary, then it is *integrated of the order n* , which we denote as $I(n)$. Thus, a series that is stationary without any differencing can be also denoted as $I(0)$. The definition assumes that n is an integer. Its extension to fractional values of n is covered by the concept of fractional integration; the topic is beyond the scope of the book but Hosking (1981) and Mills and Markellos (2008) can serve as useful references.

Prior to *differencing* and *detrending*, the most common transformation is to *take a natural logarithm* of the data in order to deal with a sort of non-linearity or to reduce an exponential trend into a linear one. For example, if we are interested in growth rates, it is natural to apply *logarithmic differencing*, which means that we first take natural logarithms of the data and then difference them.

1. *Taking a natural logarithm* is applied when the data perform exponential growth, which is a common case in economics. For example, if a GDP of

a country grows each year roughly by 3% when compared to preceding year, then the time series of yearly GDP contains an exponential trend. In such a case, by taking a natural logarithm we receive data that grow linearly.

2. *Differencing* is the most common approach applied in order to achieve stationarity. To difference a time series, we apply the transformation $\Delta y_t = y_t - y_{t-1}$, where Δy_t is the so-called first difference. To obtain second differences denoted as $\Delta^2 y_t$ we apply the identical transformation on first differences $\Delta^2 y_t = \Delta y_t - \Delta y_{t-1}$. In this way we can create differences of even higher orders. Although any time series becomes stationary after a sufficient order of differencing, differencing of a higher than second order is almost never used in econometrics. The reason is that by each differencing we lose one observation and, more important, by each differencing we lose a part of the information contained in the data. In addition, higher order differences have no clear economic interpretation. Second differences are already linear growth rates of the linear growth rates obtained by first differencing.
3. *Detrending* is a procedure that removes linear or even higher order trends from the data. To detrend a time series, we run a regression of the series on a constant, time t , and eventually its higher powers as well. Residuals from such a regression represent the detrended time series. The degree of the time polynomial included in the regression can be formally tested by an F -test prior to detrending. Trending time series are never stationary, because their mean is not constant. Therefore, detrending also helps to make such time series stationary. More details about trends in time series will be given in the next section and in section 3.2.

Example 1.3 We can illustrate the above approaches in the following way. Usually economic data grow exponentially. This means that for a variable X we have the growth equation $X_t = (1 + g_t) X_{t-1}$ in the discrete case, or $X_t = X_{t-1} e^{g_t}$ in the continuous case, where g_t is a growth rate in between two successive periods or the rate of return depending on the nature of X . By logarithmic differencing we obtain $\ln X_t - \ln X_{t-1} = \ln(1 + g_t) \approx g_t$ in the discrete case or $\ln X_t - \ln X_{t-1} = g_t$ in the continuous case. Specifically, let us consider a time series of price levels $\{P_t\}_{t=1}^T$. By logarithmic differencing we receive the series of inflation rates $\pi_t = \ln P_t - \ln P_{t-1}$.

The above mentioned differences were always just differences between two successive periods. If the data exhibit a seasonal pattern, it is more fruitful to apply