1 Introduction and overview

This book concerns the construction of time series models for describing the dynamic properties of economic variables and for out-of-sample forecasting. The economic variables can originate from various subject areas in economics and business, including macro-economics, finance, and marketing. Specific examples of time series of interest are inflation rates, unemployment rates, stock market returns, and market shares. Out-of-sample forecasts for such variables are often needed to set policy targets. For example, the forecast for next year's inflation rate can lead to a change in the monetary policy of a central bank. A forecast of a company's market share in the next few months may lead to changes in the allocation of advertising budgets. The models in this book can be called econometric time series models because we use econometric methods for analysis.

Time series variables can display a wide variety of patterns. Typically, macroeconomic aggregates such as industrial production, consumption, and wages show an upward trending pattern. Industrial production, tourism expenditures, and retail sales, among many others, display a pronounced seasonal pattern, that is, tourism spending is usually largest during the summer and retail sales tend to peak around Christmas. Another feature is that certain observations on economic data look aberrant in the sense that they occur rarely and deviate strongly from the typical behavior of the variable. For example, if new car registrations are almost zero in a certain month because of a computer breakdown, this does not reflect the true sales of new cars. Similarly, stock markets can crash with daily returns as large as minus 20 percent. Another characteristic property of financial asset prices is that periods of large price movements alternate with relatively calm periods, suggesting that the volatility of these variables changes over time. Finally, many economic time series display asymmetric or non-linear behavior. Unemployment, for example, increases rapidly during recessions but declines only slowly during expansions.

It seems obvious that there is not a single time series model that, first, can describe all of the above features simultaneously and, second, is also reasonably accurate in out-of-sample forecasting. In fact, several models are available to describe each of these features, and all these models can be used to generate forecasts. It is the key purpose of this book to survey the most relevant of these models. We discuss how they can be implemented in practice and how their possible merits for forecasting can be evaluated. It is our opinion that just like there is no uniformly best descriptive model there is no such forecasting model. Therefore, we restrict ourselves to presenting guidelines for the selection between available models for forecasting. As will become apparent from the empirical examples in Chapter 2, the specific features of economic time series often lead to an *a priori* selection of several possibly useful models. For example, certain time series models explicitly deal with seasonality, and such models are not useful for data without seasonal variation.

An important requirement for the model construction methods discussed in this book is that the practitioner has some time available to construct his or her forecasts. Indeed, if it is necessary to generate forecasts for several hundreds or even thousands of variables every day, it may be better to rely on one of the many automatic extrapolation schemes that are available, such as smoothing algorithms and exponentially weighted moving averages. We do not wish to claim that such methods are inferior or less useful, as in fact these methods often do very well, see Makridakis *et al.* (1982), Makridakis and Hibon (2000) and Koning *et al.* (2005), among others. We do claim that the decisions involved in constructing a descriptive time series model for a time series with specific features that is also useful for out-of-sample forecasting is difficult to formalize in automatic routines.

Model building

Time series variables in economics and business are observed at different frequencies. For example, estimates of gross domestic product (GDP) and other measures of economic output are available per quarter, inflation typically is measured at a monthly frequency, product sales may be given per week, while stock prices can be recorded daily. The empirical time series used in this book as running examples reflect these different possible frequencies. It turns out that to some extent the sampling frequency determines the importance of the features discussed above. For example, if quarterly observations on consumption are used, seasonal variation is a very important characteristic to be accounted for in a time series forecasting model, while this would not be the case if annual consumption were considered. Similarly, stock returns display clear signs of time-varying volatility at daily and weekly frequencies, but much less so at the monthly frequency.

The modeling strategy described in this book exploits the key property of economic time series that the sequence of the observations is determined by calendar time. For example, the observation on unemployment in 2009 always precedes observations in 2010 and later. This seems too obvious to mention, but we believe it is not. Indeed, the value of unemployment in 2010 is likely to be influenced by that in 2009. Hence when analyzing time series data we should leave their "natural ordering" intact.

Throughout this book the value of a time series y at time t is denoted by y_t, where t takes integer values from 1 to T. Note that we assume that the time series observations are regularly spaced, that is, the time period between consecutive observations y_{t-1} and y_t is the same for all t. Models that allow for irregularly spaced data are of course available, but these are not dealt with here, see Parzen (1984) for useful reading. The key property of a time series is that observation y_t always comes after y_{t-1}. Therefore it makes sense to take y_{t-1} into account when analyzing y_t. This is in contrast to cross-sectional data, where the ordering of the data points does not matter.

Given that y_{t-1} is always measured prior to y_t, it is likely that part of the value of y_{t-1} is reflected in the value y_t. For example, it is unlikely that if this month's inflation rate is 10 percent, it will be –5 percent next month. In fact, it is more likely that it will be, say, between 8 and 12 percent. Another way of putting this is that, for many time series variables, the observations y_{t-1} and y_t are correlated. Since these observations are measurements of the same variable, we say that y_t is correlated with itself. This is called autocorrelation. If there is such autocorrelation between y_t and y_{t-1}, we can exploit this correlation for forecasting. For example, if it holds that y_t on average equals $0.8y_{t-1}$ for all $t = 1, \ldots, T$, and y_t is again the inflation rate with a value of 2 percent in the Tth month, we may forecast next month's rate y_{T+1} as 1.6 percent.

The above implies that a time series variable can be characterized by its autocorrelations. Given a sample of observations y_t for $t = 1, \ldots, n$, we can estimate these correlations simply by computing their sample counterparts. The key feature of time series analysis is that such empirical autocorrelations can be exploited to obtain a first impression of which model is possibly useful to describe and forecast the time series at hand. This follows from the fact that all time series models theoretically imply certain autocorrelation properties of the time series in case these models were the true data generating processes. For example, the so-called autoregressive model of order 1 for a non-trending time series y_t implies that the correlations between y_t and y_{t-k} decline exponentially towards zero as k increases, see Chapter 3 for details. When the estimated empirical autocorrelations suggest that this pattern holds, we may be inclined to consider such a first order autoregressive model in a first round of analysis. In brief, certain features of observed time series data suggest the possible adequacy of corresponding time series models. This resembles what medical doctors do. They know that the flu can come with fever, so if a patient is diagnosed to be feverish, it might be due to the flu.

In this book we focus on five key features of economic and business time series variables. These features are trends, seasonality, aberrant observations, conditional heteroskedasticity and non-linearity. Each of these features points towards the possible usefulness of certain classes of time series models. In order to keep matters simple, we restrict attention to regression-based models like

$$y_t = \beta' x_t + \varepsilon_t, \qquad t = 1, 2, \ldots, T, \tag{1.1}$$

where y_t is the observed time series of interest, x_t is a $k \times 1$ vector of other observed time series, ε_t is an unobserved error, β is a $k \times 1$ vector of unknown parameters, and T is the sample size. Commonly, the x_t variables contain the past of y_t. Needless to say that (1.1) is an overly simplified version of the models considered later on, but here it serves as a useful illustration. When there are T observations available, and the current time point is $t = T$, we often wish to forecast h periods ahead or "out-of-sample", that is, we want to estimate y_{T+h}. Usually, the y_t variable, which is the variable of focal interest, is known. However, the x_t variables usually have to be selected from among a potentially large number of candidates. It will become clear in later chapters that autocorrelations can be helpful to decide on the most appropriate components x_t. In this book, these x_t variables (or functions thereof) are usually assumed to be observable. There are also classes of time series models where x_t is unobserved or latent, and should be estimated as well. In these so-called unobserved components models such x_t variables can be labeled as "trend" or "seasonal fluctuations", see Harvey (1989) and Durbin and Koopman (2001) for excellent treatments of these models. Harvey (2006) provides a survey on forecasting with unobserved components models. Furthermore, also in order to limit the exposition, we confine ourselves to situations where the functional form of the relationship between y_t and x_t is known and can be characterized by a few parameters. In (1.1) this relationship is linear. In Chapters 7 and 8, we will discuss some non-linear time series models. For a detailed treatment of non-parametric methods, which allow for more flexible functional relationships between y_t and x_t, the interested reader is referred to Härdle *et al.* (1997), Fan and Gijbels (1970), and Pagan and Ullah (1999), among others.

Statistical method

There are two different approaches in analyzing time series. The first of these is called the frequency domain approach, which makes extensive use of spectral analysis. The key assumption within this approach is that any time series can be decomposed into a certain number of cyclical components with different frequencies. In fact, the lengths (and amplitudes) of these cycles can be exploited to characterize a time series. For example, a cycle of infinite length corresponds with a trend, see Granger (1966). Similarly, for a quarterly observed time series a cycle of four quarters corresponds with a seasonal pattern. Although spectral techniques can be useful to obtain an impression of the salient features of a time series and to describe, for example, business cycle properties, these are not often explicitly considered for out-of-sample forecasting. For a thorough treatment of frequency domain techniques, see, for example, Priestley (1981).

The second approach to analyzing time series, which is much more common in economics and business is called the time domain approach. Within this time domain approach, the autocorrelation function plays a central role, see Box and Jenkins (1970). Although the area of time series analysis has expanded widely since 1970, the main ideas underlying Box and Jenkins' work remain valid. In the present book we confine the discussion to this time domain approach, also since it is easier to use for data with features such as outliers, non-linearity and seasonality. In fact, in those latter cases, the application and interpretation of spectral techniques is not at all trivial.

A final remark on the statistical method concerns Bayesian and classical statistical methods. Without taking a standpoint toward favoring either of these two approaches, in this book we will use only classical statistical methods. For treatments of Bayesian methods in statistics and econometrics, see Zellner (1970), Bernardo and Smith (1994), Poirier (1995), Geweke (2005), and Koop (2003), among others. Geweke and Whiteman (2006) provide a detailed discussion and overview of Bayesian forecasting.

The data

An obvious first step when forecasting economic and business time series is to collect the relevant data for model construction. Sometimes this is easy, but in many cases we have to make many decisions before a useful set of data is available. For example, how to define a market share when only the information of about 50 percent of the market is available? Or, how should we define the unemployment rate? Does unemployment also include persons who work less than five hours per week? After how many releases of GDP figures can we say that they are final?

In this book we use several empirical time series to illustrate various concepts and models. Examples of the series used in this book for illustration are annual GDP series for Latin American countries, daily Dow Jones data, quarterly unemployment in the US, weekly observed market and distribution shares for a fast-moving consumer good, monthly new passenger car registrations, and four-weekly advertising expenditures on television and radio. The data can be obtained from http://people.few.eur.nl/djvandijk/tsmbef/data in EViews and Excel format. These data sets can be used by the reader to verify the empirical results reported here and also to try alternative modeling strategies and to examine the properties of other out-of-sample forecasts.

Forecasting

From a methodological point of view it is important to be precise about the goal of econometric time series modeling. In this book, we emphasize that time series models should both give an adequate description of the available in-sample data, and be useful for out-of-sample forecasting. The concept of adequate description will be discussed in detail in Chapter 3. When it comes to forecasting, a crucial assumption is the data in the sample that is used for model specification and estimation are similar to the out-of-sample data. If not, there is no point in spending much time on the construction of high-brow time series models, at least if the possible dissimilarities are not taken into account properly. It is crucial that we evaluate the stability of the forecasting model. For example, if all forecasts are too high or too low, we should obviously wish to re-specify the time series model.

An empirical strategy that can be helpful to assess the stability of the model and the modeling environment is the following. Suppose we have $T + P$ observations for a variable y_t to our disposition. We can then use the first T observations to construct a model and to estimate its parameters, and we can use the last P observations to evaluate its out-of-sample forecasting performance. Hence, we obtain forecasts of y_{T+h} for $h = 1, 2, \ldots, P$ from a model that is constructed using observations $y_1, y_2, \ldots, y_T$. The accuracy of these forecasts can be assessed by comparison with the true values observed for $y_{T+1}, y_{T+2}, \ldots, y_{T+P}$. When the forecasting performance is satisfactory, we may want to generate forecasts for future, unknown observations at times $T + P + 1, T + P + 2, \ldots$, where we typically re-estimate the model parameters using all $T + P$ available observations.

There are no strict guidelines for the choice of P, the number of forecasts used to evaluate the predictive accuracy of the constructed time series model. On the one hand, it is important that T is large enough to have reasonable precision for the estimates of the unknown parameters in the model. On the other hand, P should also be large enough to allow for a meaningful comparison of the predictive accuracy of various competing models. So, it is up to the user to decide on the appropriate values of T and P.

Outline of this book

The contents of this book are as follows. Chapter 2 surveys typical features of time series variables in economics and business. We limit this discussion to five such features, that is, trends, seasonality, aberrant data, time-varying variance, and non-linearity. These features correspond with a decreasing source of variation in economic time series. The most dominant source of variation is often the trend. The next dominant source is seasonal variation, while the smallest amount of variation often tends to be due to non-linearity. There is no explicit treatment of a “business cycle” here since this sometimes corresponds with cyclical dynamics in the time series itself, sometimes with short-term deviations from a trend and sometimes as a specific type of non-linearity or the occurrence of outliers. Each of the five features suggests different model structures and model specification methods.

We cover each of these features in Chapters 4 to 8. As a prelude, Chapter 3 is concerned with a discussion of several important concepts in time series analysis. Intentionally, this discussion is far from being as technically rigorous as Anderson (1971), Fuller (1976) and Hamilton (1994). Instead the focus is on discussing those techniques which can be readily applied in practice. When necessary, we give references to studies that include proofs of formal asymptotics and other results.

Most of the discussion in Chapter 3, as well as that in Chapters 4 to 8, deals with univariate time series. In Chapter 9 some of the concepts in Chapter 3 are extended to multivariate time series. In addition, the focus in this chapter is on common aspects across economic time series, in particular common trends. Finally, most chapters make mention of recent or even current research topics. The research area of time series analysis is very active, and we can expect many new developments in the near future.

2 Key features of economic time series

In this chapter the focus is on key features of business and economic time series. It also serves to introduce several of the empirical time series that will be used throughout this book as running examples.

The five key features of economic and business data that we consider are (i) trends, (ii) seasonality, (iii) influential data points, (iv) time-varying (conditional) variance (conditional heteroskedasticity) and (v) non-linearity. Typically, an economic time series displays at least one, but usually two or more of these features. To keep matters simple, however, in this chapter each series is analyzed for only one of these five features at a time. In later chapters, the various possible models for each of these features will sometimes be combined to illustrate the practice of time series modeling.

In this chapter, each of the five data features will be illuminated using simple regression-based calculations. This should not imply that these regression models are the best we can do. They are merely helpful tools to demonstrate the properties of the data. A second important tool in this chapter is graphical analysis. In many cases it is already quite helpful just to put the data in a graph with the values of the observations (or some transformation thereof) on the vertical axis and time on the horizontal axis. However, in case of many observations or of a time series with large variance, it may sometimes be more insightful to construct 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 can be replaced by, for example, y_{t-1}. In practical situations, we would advise to construct each of the graphs and to estimate each of the simple regression models in order to obtain an overall insight in the specific data properties.

The data that are considered in practice usually come in their original format, that is, they have not been transformed. In empirical time series analysis it is quite common to analyze data after applying the natural logarithmic transformation. Hence, if the original data are denoted by w_t, we usually model and forecast $y_t = \log(w_t)$, where log denotes the natural logarithm. One of the reasons for the log-transformation is that it makes an exponential trend to become a linear trend. Also, if the time series shows increasing variation over time, the log-transformation dampens this trend. When the

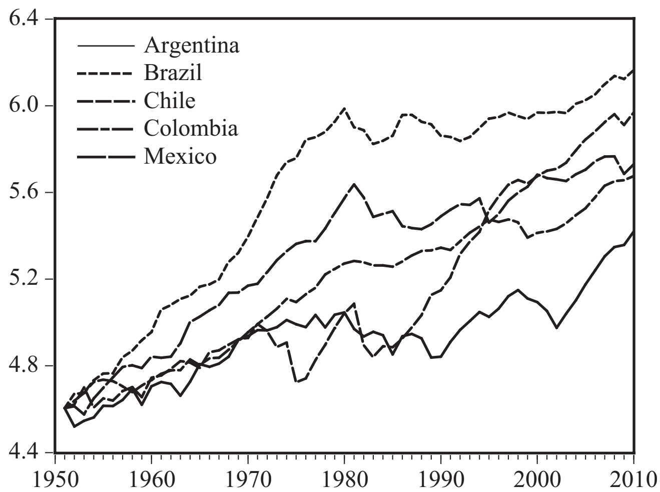

Figure 2.1: Annual indices of log real GDP per capita in Latin American countries, 1951–2010 (1951 = 100).

data are already in relative format, as in the case of the unemployment rate, interest rates or market shares, for example, the log transformation is usually not applied.

2.1 Trends

One of the key features of many economic and business time series is the trend, by which we mean, at least for the moment, that the data show a general tendency to increase or decrease over time. Such a trend can take different shapes. It can be upward or downward sloping, it can be steep or moderate, and it can be exponential or approximately linear. As will become clear below, for many time series the trend is the dominant source of variation, which makes it of crucial importance for out-of-sample forecasting. If a trend is wrongly incorporated in a time series model, forecasts will be poor, especially in the long run.

To illustrate the presence of trends in economic data, consider the five time series shown in Figure 2.1, which are annual indices of real gross domestic product (GDP) per capita (in logs) in the five largest Latin American economies for the sample period 1951–2010, that is, Argentina, Brazil, Chile, Colombia and Mexico. From the graph it can be observed that over the complete 60-year sample period all five countries have grown considerably, although at different rates, with Brazil and Argentina showing the highest and lowest average growth rates, respectively. Furthermore, although the

Table 2.1: Trends in real GDP per capita in Latin American countries, 1951–2010

Country	$\hat{\beta}$ in regression: $y_t = \alpha + \beta t + u_t$		$\hat{\mu}$ in regression: $y_t - y_{t-1} = \mu + u_t$	
Argentina	1.06	(0.07)	1.42	(0.62)
Brazil	2.48	(0.14)	2.64	(0.51)
Chile	2.28	(0.12)	2.31	(0.75)
Colombia	1.78	(0.04)	1.76	(0.32)
Mexico	1.92	(0.09)	1.99	(0.53)

Note: The numbers in parentheses are estimated standard errors. All numbers are multiplied by 100.

(indexed) levels of real GDP per capita in Chile, Colombia and Mexico are approximately the same in the year 1995, such that their average growth rates are rather close, their developments during the preceding period were quite different. GDP growth in Colombia appears to have been rather stable at an almost constant pace. Mexico experienced rapid growth up to 1980 when a long recession started, which lasted until 1988 and was followed by a gradual recovery until another recession occurred in 1995. Finally, GDP growth in Chile was about the same as in Colombia until the early 1970s when the economy was hit by a severe recession, followed by a second one in the early 1980s. Starting in 1984, a steep recovery occurred that took GDP per capita back to the levels of Mexico and Colombia within a decade.

To quantify the trends in the five GDP per capita series, consider the following simple regression model

$$y_t = \alpha + \beta t + u_t, \qquad t = 1, 2, \ldots, T, \tag{2.1}$$

where α and β are unknown parameters and where u_t is an unknown residual error time series with mean zero. Note that the GDP series in Figure 2.1 show some cyclical behavior around their trends, indicating that u_t may be correlated with u_{t-1}, and hence suggesting misspecification of (2.1). The standard errors of the parameter estimates should therefore be interpreted with care.

The left-hand panel of Table 2.1 displays the estimates of β in (2.1), multiplied by 100 for convenience. It is clear that the upward trend is steepest for Brazil ($\hat{\beta} = 2.48$) and that average growth is lowest for Argentina ($\hat{\beta} = 1.06$). As expected, the estimates of β are fairly close for Chile, Colombia, and Mexico. Note, though, that the standard errors for Mexico and Chile are two and three times as large as for Colombia, reflecting the much steadier growth pattern of the latter country.