Contents

1 Multiplicative SARIMA models

Rong Chen, *Rainer Schulz* and *Sabine Stephan*

1.1 Introduction

In the history of economics, the analysis of economic fluctuations can reclaim a prominent part. Undoubtedly, the analysis of business cycle movements plays the dominant role in this field, but there are also different perspectives to look at the ups and downs of economic time series. Economic fluctuations are usually characterized with regard to their periodic recurrence. Variations that last several years and occur in more or less regular time intervals are called business cycles, whereas seasonality (originally) indicates regularly recurring fluctuations within a year, that appear due to the season. Such seasonal patterns can be observed for many macroeconomic time series like gross domestic product, unemployment, industrial production or construction.

The term seasonality is also used in a broader sense to characterize time series that show specific patterns that regularly recur within fixed time intervals (e.g. a year, a month or a week). Take as an example the demand for Christmas trees: the monthly demand in November and especially in December will be generally very high compared to the demand during the other months of the year. This pattern will be the same for every year—irrespective of the total demand for Christmas trees. One can also detect seasonal patterns in financial time series like in the variance of stock market returns. The highest volatility is often observed on Monday, mainly because investors used the weekend to think carefully about their investments, to obtain new information and to come to a decision.

As we saw so far, seasonality has many different manifestations. Consequently, there are different approaches to model seasonality. If we focus on macroeconomic time series the class of seasonal models is confined to processes with dynamic properties at periods of a quarter or a month. However, when financial time series are studied, then our interest shifts to seasonal patterns at the daily level together with seasonal properties in higher moments. Therefore, it is no surprise, that a rich toolkit of econometric techniques has been developed to model seasonality.

In the following we are going to deal with seasonality in the mean only (for seasonality in higher moments see Ghysels and Osborn (2001)), but there are still different ways to do so. The choice of the appropriate technique depends on whether seasonality is viewed as *deterministic* or *stochastic*. The well-known deterministic approach is based on the assumption, that seasonal fluctuations are fix and shift solely the level of the time series. Therefore, deterministic seasonality can be modelled by means of seasonally varying intercepts using seasonal dummies. Stochastic seasonality however is a topic in recent time series analysis and is modelled using appropriate ARIMA models (Diebold, 1998, Chapter 5). Since these *seasonal* ARIMA models are just an extension of the usual ARIMA methodology, one often finds the acronym SARIMA for this class of models (Chatfield, 2001).

The topic of this chapter is the modelling of seasonal time series using SARIMA models. The outline of this chapter is as follows: the next Section 1.2.1 illustrates, how to develop an ARIMA model for a seasonal time series. Since these models tend to be quite large, we introduce in Section 1.2.2 a parsimonious model specification, that was developed by Box and Jenkins (1976)—the *multiplicative SARIMA model*. Section 1.3 deals with the identification of these models in detail, using the famous airline data set of Box and Jenkins for illustrative purposes. Those who already studied the Section *ARIMA model building* in Chapter 4 on *Univariate Time Series Modeling*, will recognize that we use the same tools to identify the underlying data generation process. Finally, in Section 1.4 we focus on the estimation of multiplicative SARIMA models and on the evaluation of the fitted models.

All quantlets for modelling multiplicative SARIMA models are collected in XploRe's `times` library.

1.2 Modeling seasonal time series

1.2.1 Seasonal ARIMA models

Before one can specify a model for a given data set, one must have an initial guess about the data generation process. The first step is always to plot the time series. In most cases such a plot gives first answers to questions like: "Is the time series under consideration stationary?" or "Do the time series show a seasonal pattern?"

Figure 1.1 displays the quarterly unemployment rate u_t for Germany (West) from the first quarter of 1962 to the forth quarter of 1991. The data are published by the OECD (Franses, 1998, Table DA.10). The solid line represents

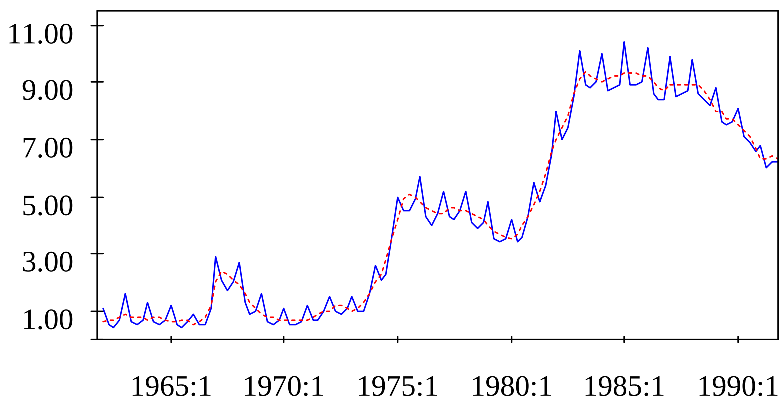

Figure 1.1: Quarterly unemployment rate for Germany (West) from 1962:1 to 1991:4. The original series u_t is given by the solid blue line and the seasonally adjusted series is given by the dashed red line.

XEGmsarima1.xpl

the original series u_t and the dashed line shows the seasonally adjusted series. It is easy to see, that this quarterly time series possesses a distinct seasonal pattern with spikes recurring always in the first quarter of the year.

After the inspection of the plot, one can use the sample autocorrelation function (ACF) and the sample partial autocorrelation function (PACF) to specify the order of the ARMA part (see `acf`, `pacf`, `acfplot` and `pacfplot`). Another convenient tool for first stage model specification is the extended autocorrelation function (EACF), because the EACF does not require that the time series under consideration is stationary and it allows a simultaneous specification of the autoregressive and moving average order. Unfortunately, the EACF can not be applied to series that show a seasonal pattern. However, we will present the EACF later in Section 1.4.5, where we use it for checking the residuals resulting from the fitted models.

Figures 1.2, 1.3 and 1.4 display the sample ACF of three different transformations of the unemployment rate u_t for Germany. Using the difference—or backshift—operator L, these kinds of transformations of the unemployment rate can be written compactly as

$$\Delta^d \Delta_s^D u_t = (1-L)^d (1-L^s)^D u_t \ ,$$

where L^s operates as $L^s u_t = u_{t-s}$ and s denotes the seasonal period. Δ^d and Δ_s^D stand for nonseasonal and seasonal differencing. The superscripts d and D indicate that, in general, the differencing may be applied d and D times.

Figure 1.2 shows the sample ACF of the original data of the unemployment rate u_t. The fact, that the time series is neither subjected to nonseasonal nor to seasonal differencing, implies that $d = D = 0$. Furthermore, we set $s = 4$, since the unemployment rate is recorded quarterly. The sample ACF of the unemployment rate declines very slowly, i.e. that this time series is clearly nonstationary. But it is difficult to isolate any seasonal pattern as all autocorrelations are dominated by the effect of the nonseasonal unit root.

Figure 1.3 displays the sample ACF of the first differences of the unemployment rate Δu_t with

$$\Delta u_t = u_t - u_{t-1} \ .$$

Since this transformation is aimed at eliminating only the nonseasonal unit root, we set $d = 1$ and $D = 0$. Again, we set $s = 4$ because of the frequency of the time series under consideration. Taking the first differences produces a very clear pattern in the sample ACF. There are very large positive autocorrelations at the seasonal frequencies (lag 4, 8, 12, etc.), flanked by negative autocorrelations at the 'satellites', which are the autocorrelations right before and after the seasonal lags. The slow decline of the seasonal autocorrelations

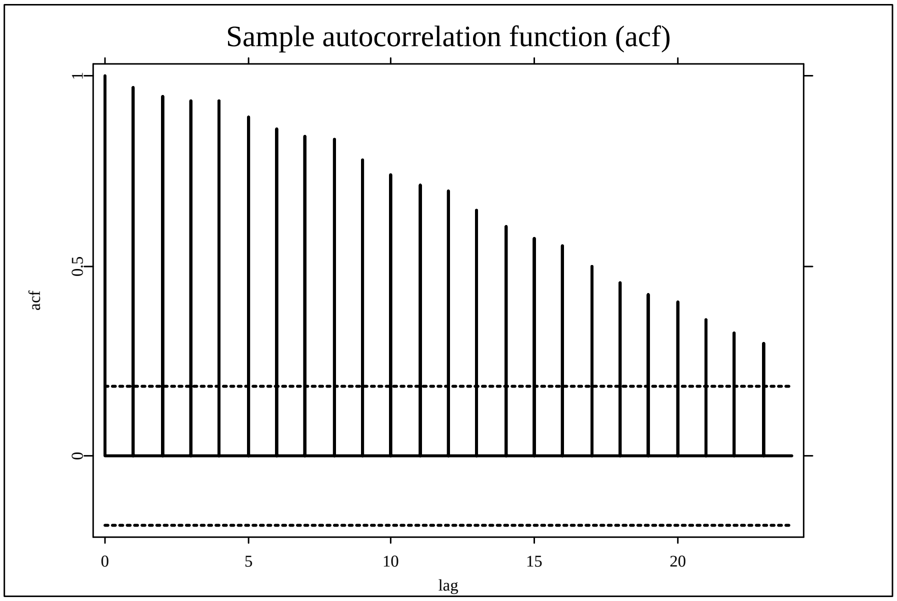

Figure 1.2: Sample ACF of the unemployment rate u_t for Germany (West) from 1962:1 to 1991:1.

XEGmsarima2.xpl

indicates seasonal instationarity. Analogous to the analysis of nonseasonal nonstationarity, this may be dealt by seasonal differencing; i.e. by applying the $\Delta_4 = (1 - L^4)$ operator in conjunction with the usual lag operator $\Delta = (1 - L)$ (Mills, 1990, Chapter 10).

Eventually, Figure 1.4 displays the sample ACF of the unemployment rate that was subjected to the final transformation

$$\begin{aligned} \Delta\Delta_4 u_t &= (1 - L)(1 - L^4)u_t \\ &= (u_t - u_{t-4}) - (u_{t-1} - u_{t-5}) \,. \end{aligned}$$

Since this transformation is used to remove both the nonseasonal and the seasonal unit root, we set $d = D = 1$. What the transformation $\Delta\Delta_4$ finally

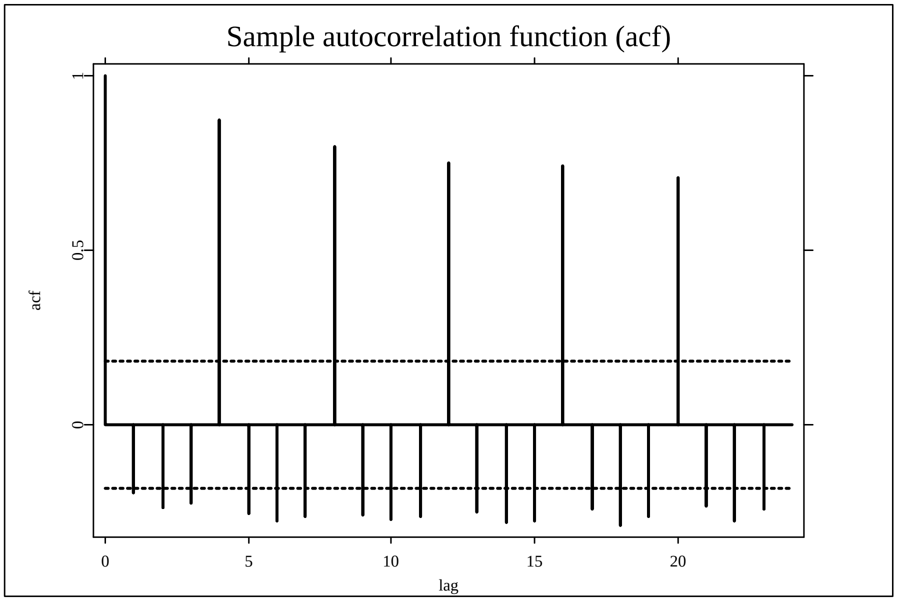

Figure 1.3: Sample ACF of the first differences of the unemployment rate Δu_t for Germany (West) from 1962:1 to 1991:1.

XEGmsarima3.xpl

does is seasonally differencing the first differences of the unemployment rate. By means of this transformation we obtain a stationary time series that can be modeled by fitting an appropriate ARMA model.

After this illustrative introduction, we can now switch to theoretical considerations. As we already saw in practice, a seasonal model for the time series $\{x_t\}_{t=1}^T$ may take the following form

$$\Delta^d \Delta_s^D x_t = \frac{\Theta(L)}{\Phi(L)} a_t \ , \tag{1.1}$$

where $\Delta^d = (1-L)^d$ and $\Delta_s^D = (1-L^s)^D$ indicate nonseasonal and seasonal differencing and s gives the season. a_t represents a white noise innovation.

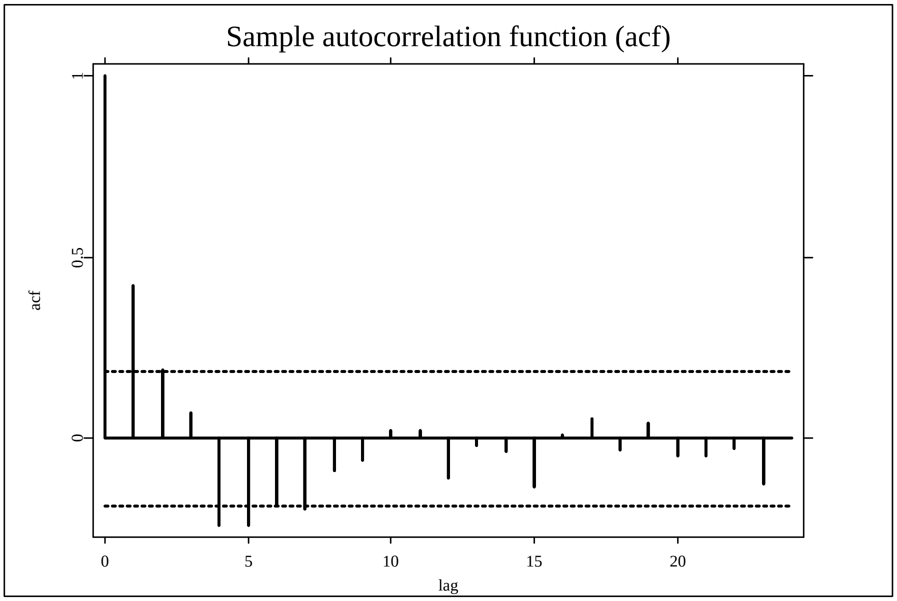

Figure 1.4: Sample ACF of the seasonally differenced first differences of the unemployment rate $\Delta\Delta_4 u_t$ for Germany (West) from 1962:1 to 1991:1.

XEGmsarima4.xpl

$\Phi(L)$ and $\Theta(L)$ are the usual AR and MA lag operator polynomials for ARMA models

$$\Phi(L) \equiv 1 - \phi_1 L - \phi_2 L^2 - \ldots - \phi_p L^p$$

and

$$\Theta(L) \equiv 1 + \theta_1 L + \theta_2 L^2 + \ldots + \theta_q L^q \ .$$

Since the $\Phi(L)$ and $\Theta(L)$ must account for seasonal autocorrelation, at least one of them must be of minimum order s. This means that the identification of models of the form (1.1) can lead to a large number of parameters that have to be estimated and to a model specification that is rather difficult to interpret.

1.2.2 Multiplicative SARIMA models

Box and Jenkins (1976) developed an argument for using a restricted version of equation (1.1), that should be adequate to fit many seasonal time series. Starting point for their approach was the fact, that in seasonal data there are two time intervals of importance. Supposing that we are still dealing with a quarterly series, we expect the following to occur (Mills, 1990, Chapter 10):

- a seasonal relationship between observations for the same quarters in successive years, and
- a relationship between observations for successive quarters in a particular year.

Referring to Figure 1.1 that displays the quarterly unemployment rate for Germany, it is obvious that the seasonal effect implies that an observation in the first quarter of a given year is related to the observations of the first quarter for previous years. We can model this feature by means of a *seasonal model*

$$\Phi_s(L)\Delta_s^D x_t = \Theta_s(L)v_t \ . \tag{1.2}$$

$\Phi_s(L)$ and $\Theta_s(L)$ stand for a seasonal AR polynomial of order p and a seasonal MA polynomial of order q respectively:

$$\Phi_s(L) = 1 - \phi_{s,1}L^s - \phi_{s,2}L^{2s} - \ldots - \phi_{s,P}L^{Ps}$$

and

$$\Theta_s(L) = 1 + \theta_{s,1}L^s + \theta_{s,2}L^{2s} + \ldots + \theta_{s,Q}L^{Qs} \ ,$$

which satisfy the standard stationarity and invertibility conditions. v_t denotes the error series. The characteristics of this process are explained below.

It is obvious that the above given seasonal model (1.2) is simply a special case of the usual ARIMA model, since the autoregressive and moving average relationship is modeled for observations of the same seasonal time interval in different years. Using equation (1.2) relationships between observations for the same quarters in successive years can be modeled.

Furthermore, we assume a relationship between the observations for successive quarters of a year, i.e. that the corresponding error series (v_t, v_{t-1}, v_{t-2}, etc.) may be autocorrelated. These autocorrelations may be represented by a *nonseasonal model*

$$\Phi(L)\Delta^d v_t = \Theta(L)a_t \ . \tag{1.3}$$

v_t is ARIMA(p, d, q) with a_t representing a process of innovations (white noise process).

Substituting (1.3) into (1.2) yields the *general multiplicative seasonal model*

$$\Phi(L)\Phi_s(L)\Delta^d\Delta_s^D x_t = \delta + \Theta(L)\Theta_s(L)a_t \ . \tag{1.4}$$

In equation (1.4) we additionally include the constant term δ in order to allow for a deterministic trend in the model (Shumway and Stoffer, 2000). In the following we use the short-hand notation SARIMA $(p, d, q) \times (s, P, D, Q)$ to characterize a multiplicative seasonal ARIMA model like (1.4).

1.2.3 The expanded model

Before starting with identification and estimation of a multiplicative SARIMA model a short example may be helpful. This example sheds some light on the connection between a multiplicative SARIMA $(p, d, q) \times (s, P, D, Q)$ and a simple ARMA (p, q) model. It reveals that the SARIMA methodology leads to parsimonious models.

Polynomials in the lag operator are algebraically similar to simple polynomials $ax + bx^2$. So it is possible to calculate the product of two lag polynomials (Hamilton, 1994, Chapter 2).

Given that fact, every multiplicative SARIMA model can be telescoped out into an ordinary ARMA(p, q) model in the variable

$$y_t \stackrel{\text{def}}{=} \Delta_s^D \Delta^d x_t \ .$$

For example, let us assume that the series $\{x_t\}_{t=1}^T$ follows a SARIMA$(0, 1, 1) \times (12, 0, 1, 1)$ process. In that case, we have

$$(1 - L^{12})(1 - L)x_t = (1 + \theta_1 L)(1 + \theta_{s,1}L^{12})a_t \ . \tag{1.5}$$

After some calculations one obtains

$$y_t = (1 + \theta_1 L + \theta_{s,1}L^{12} + \theta_1\theta_{s,1}L^{13})a_t \tag{1.6}$$

where $y_t = (1 - L^{12})(1 - L)x_t$. Thus, the multiplicative SARIMA model has an ARMA(0,13) representation where only the coefficients

$$\theta_1 \ , \quad \theta_{12} \stackrel{\text{def}}{=} \theta_{s,1} \quad \text{and} \quad \theta_{13} \stackrel{\text{def}}{=} \theta_1\theta_{s,1}$$

are not zero. All other coefficients of the MA polynomial are zero.

Thus, we are back in the well-known ARIMA(p, d, q) world. However, if we know that the original model is a SARIMA(0,1,1)×(12,0,1,1), we have to estimate only the two coefficients θ_1 and $\theta_{s,1}$. For the ARMA(0,13) we would estimate instead the three coefficients θ_1, θ_{12}, and θ_{13}. Thus it is obvious that SARIMA models allow for a parsimonious model building.

In the following, a model specification like (1.6) is called an *expanded model*. In Section 1.4 it is shown that this kind of specification is required for estimation purposes. Only an expanded multiplicative model can be estimated directly.

1.3 Identification of multiplicative SARIMA models

This section deals with the identification of a multiplicative SARIMA model. The required procedure is explained step by step, using the famous airline data of Box and Jenkins (1976, Series G) for illustrative purposes. The date give the number of airline passengers (in thousands) in international air travel from 1949:1 to 1960:12. In the following G_t denotes the original series.

The identification procedure comprises the following steps: plotting the data, possibly transforming the data, identifying the dependence order of the model, parameter estimation, and diagnostics. Generally, selecting the appropriate model for a given data set is quite difficult. But the task becomes less complicated, if the following approach is observed: one thinks first in terms of finding difference operators that produce a roughly stationary series and then in terms of finding a set of simple ARMA or multiplicative SARMA to fit the resulting residual series.

As with any data analysis, the time series has to be plotted first so that the graph can be inspected. Figure 1.5 shows the airline data of Box and Jenkins. The series G_t shows a strong seasonal pattern and a definite upward trend. Furthermore, the variability in the data grows with time. Therefore, it is necessary to transform the data in order to stabilize the variance. Here, the natural logarithm is used for transforming the data. The new time series is defined as follows

$$g_t \stackrel{\text{def}}{=} \ln G_t \ .$$

Figure 1.6 displays the logarithmically transformed data g_t. The strong sea-

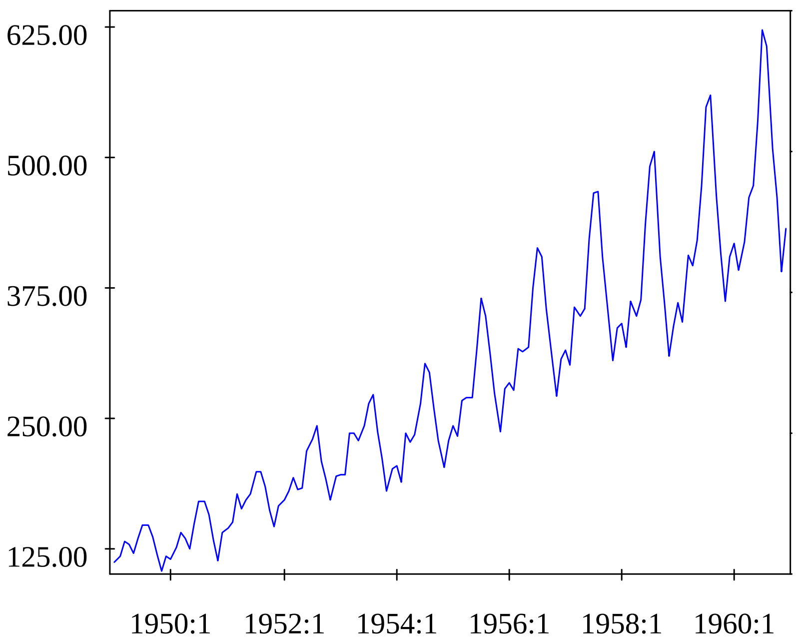

Figure 1.5: Number of airline passengers G_t (in thousands) in international air travel from 1949:1 to 1960:12.

XEGmsarima5.xpl

sonal pattern and the obvious upward trend remain unchanged, but the variability is now stabilized. Now, the first difference of time series g_t has to be taken in order to remove its nonseasonal unit root, i.e. we have $d = 1$. The new variable

$$\Delta g_t = (1 - L)g_t \tag{1.7}$$

has a nice interpretation: it gives approximately the monthly growth rate of the number of airline passengers.

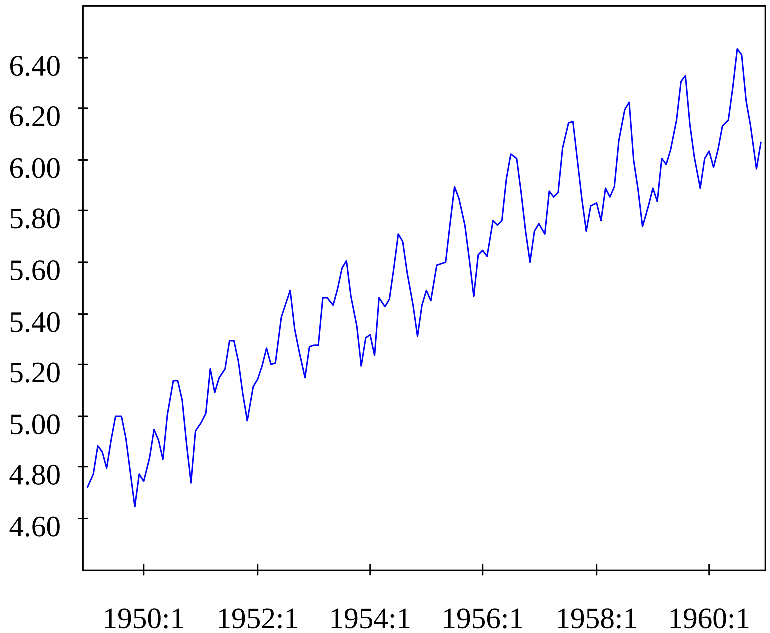

Figure 1.6: Log number of airline passengers g_t in international air travel from 1949:1 to 1960:12.

 XEGmsarima6.xpl

The next step is plotting the sample ACF of the monthly growth rate Δg_t. The sample ACF in Figure 1.7 displays a recurrent pattern: there are significant peaks at the seasonal frequencies (lag 12, 24, 36, etc.) which decay slowly. The autocorrelation coefficients of the months in between are much smaller and follow a regular pattern. The characteristic pattern of the ACF indicates that the underlying time series possesses a seasonal unit root. Typically, $D = 1$ is sufficient to obtain seasonal stationarity. Therefore, we take the seasonal

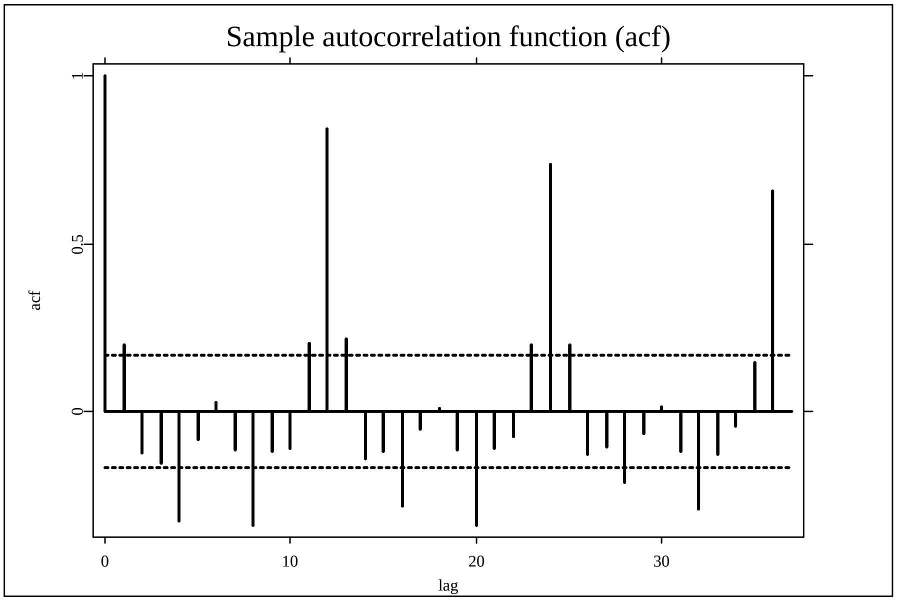

Figure 1.7: Sample ACF of the monthly growth rate of the number of airline passengers Δg_t.

XEGmsarima7.xpl

difference and obtain the following time series

$$\Delta_{12}\Delta g_t = (1 - L)(1 - L^{12})g_t$$

that neither incorporates an ordinary nor a seasonal unit root.

After that, the sample ACF and PACF of $\Delta_{12}\Delta g_t$ has to be inspected in order to explore the remaining dependencies in the stationary series. The autocorrelation functions are given in Figures 1.8 and 1.9. Compared with the characteristic pattern of the ACF of Δg_t (Figure 1.7) the pattern of the ACF and PACF of $\Delta_{12}\Delta g_t$ are far more difficult to interpret. Both ACF and PACF show significant peaks at lag 1 and 12. Furthermore, the PACF displays autocorrelation for many lags. Even these patterns are not that clear, we might feel that we face a seasonal moving average and an ordinary MA(1). Another possible specification could be an ordinary MA(12), where only the coefficients θ_1 and θ_{12} are different from zero.

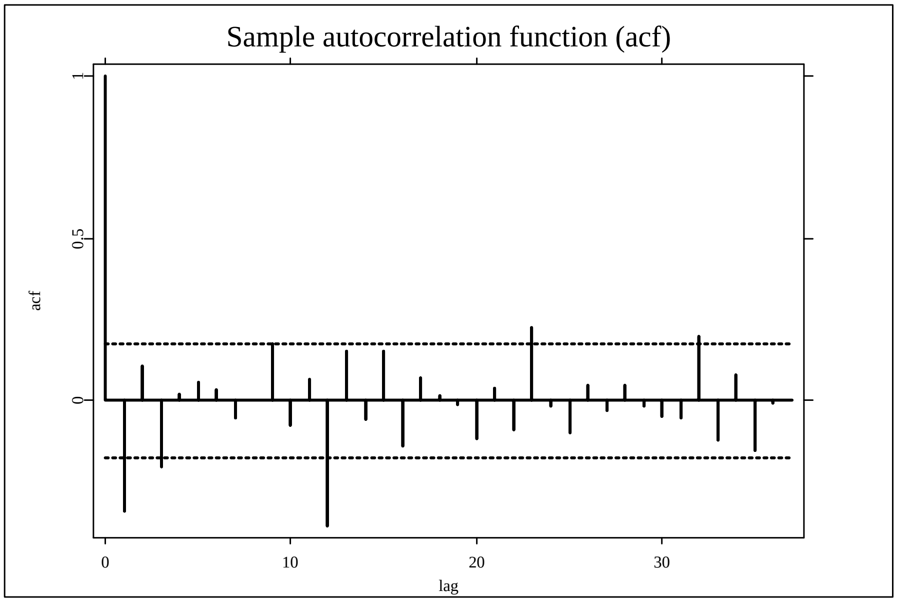

Figure 1.8: Sample ACF of the seasonally differenced growth rate of the airline data $\Delta_{12}\Delta g_t$.

XEGmsarima8.xpl

Thus, the identification procedure leads to two different multiplicative SARIMA specifications. The first one is a SARIMA(0,1,1)×(12,0,1,1). Using the lag-operator this model can be written as follows:

$$\begin{aligned}(1-L)(1-L^{12})G_t &= (1+\theta_1 L)(1+\theta_{s,1}L^{12})a_t \\ &= (1+\theta_1 L+\theta_{s,1}L^{12}+\theta_1\theta_{s,1}L^{13})a_t.\end{aligned}$$

The second specification is a SARIMA(0,1,12)×(12,0,1,0). This model has the

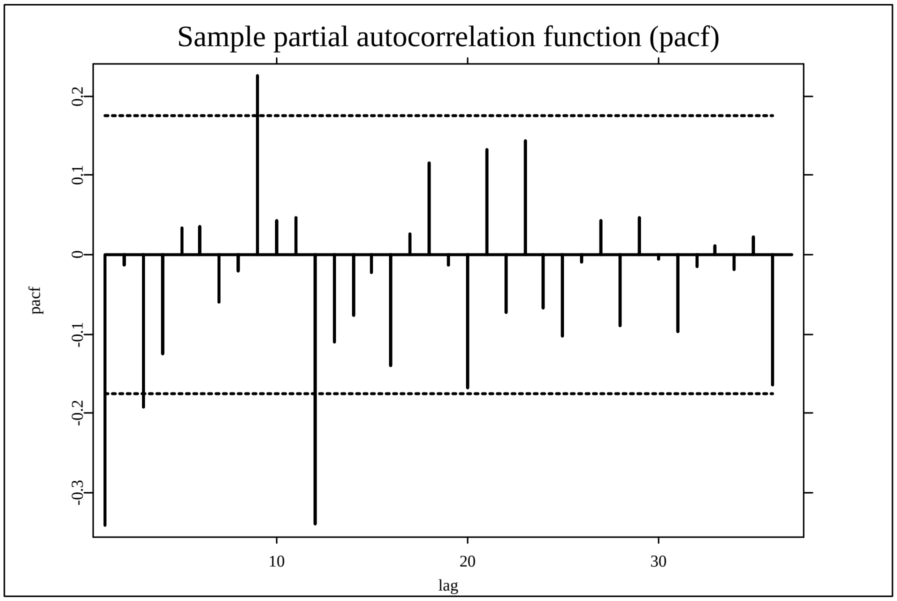

Figure 1.9: Sample PACF of the seasonally differenced growth rate of the airline data $\Delta_{12}\Delta g_t$.

XEGmsarima8.xpl

following representation:

$$(1 - L)(1 - L^{12})G_t = (1 + \theta_1 L + \theta_{12} L^{12})a_t.$$

Note, that in the last equation all MA coefficients other than θ_1 and θ_{12} are zero. With the specification of the SARIMA models the identification process is finished. We saw that modeling the seasonal ARMA after removing the nonseasonal and the seasonal unit root was quite difficult, because the sample ACF and the PACF did not display any clear pattern. Therefore, two different SARIMA models were identified that have to be tested in the further analysis.

1.4 Estimation of multiplicative SARIMA models

This section deals with the estimation of identified SARIMA models, i.e. we want to estimate the unknown parameters ψ of the multiplicative SARIMA model. If the model contains only AR terms, the unknown coefficients can be estimated using ordinary least squares. However, if the model contains MA terms too, the task becomes more complicated, because the lagged values of the innovations are unobservable. Consequently, it is not possible to derive explicit expressions to estimate the unknown coefficients and therefore one has to use maximum likelihood for estimation purposes. In the next subsection 1.4.1 the theoretical background of maximizing a likelihood function is briefly outlined.

In order to convey an idea of what follows, we will shortly outline the procedure: first, one sets up the multiplicative SARIMA model—in the following also called *original model*— with some initial values for the unknown parameters ψ. In subsection 1.4.2 it is explained how to set the original SARIMA model using the quantlet `msarimamodel`. Restrictions can be imposed on the coefficients. The simplest restriction is that some of the coefficients are zero. Then the value of the likelihood function—given the initial parameters—is evaluated.

Unfortunately, in most cases the original SARIMA model cannot be estimated directly. If one looks at the SARIMA(3,1,1)×(12,1,0,0) model in section 1.4.3—equation (1.18)—one recognizes on the left hand side the product of two expressions. Both of them contain lag-operators. Such expressions have to be telescoped out first. XploRe provides a very convenient tool to do so: `msarimaconvert`. This quantlet is explained in detail in subsection 1.4.3. The result you get from `msarimaconvert` is an ordinary ARMA(p,q) model which can be estimated.

Under the assumption that an ARMA model is stationary and invertible and that the observations are normally distributed, it can be estimated using the maximum likelihood approach. By making suitable assumptions about initial conditions, the maximum likelihood estimators can be obtained by minimizing the conditional sum of squares. In subsection 1.4.4 the quantlet `msarimacond` is presented. It calculates the conditional sum of squares function and allows for zero restrictions on the coefficients. Given this function, numerical methods have to be applied to maximize the likelihood function with respect to ψ.

To evaluate the fit of the estimated model, the quantlet `msarimacond` also delivers several criteria for diagnostic checking. The residuals of the model should be white noise. The quantlet `eacf` provides an easy way to check the behavior of the residuals.

However, the conditional sum of squares is not always very satisfactory for seasonal series. In that case the calculation of the exact likelihood function becomes necessary (Box and Jenkins, 1976, p. 211). One approach is to set up the likelihood function via Kalman filter techniques. We briefly discuss how to set up the airline model in state space form and how to use the Kalman filter to evaluate the exact likelihood function. Once again, numerical methods are necessary to maximize the exact likelihood function.

1.4.1 Maximum likelihood estimation

The approach of maximum likelihood (ML) requires the specification of a particular distribution for a sample of T observations y_t. Let

$$f_{Y_T,Y_{T-1},\ldots,Y_1}(y|\psi)$$

denote the probability density of the sample given the unknown $(n \times 1)$ parameters ψ. y is the vector of all observations $(y_T, y_{T-1}, \ldots, y_1)$. $f(y|\psi)$ can be interpreted as the probability of having observed the given sample (Hamilton, 1994, p. 117).

With the sample y at hand, the above given probability can be rewritten as a function of the unknown parameters given the sample. We use the notation $L(\psi|y)$ to denote the likelihood function evaluated at the given sample and given specified values of the unknown parameters ψ. We have to find the value of ψ that maximizes $L(\psi|y)$. This explains the name of the ML approach: given a specified distribution and a sample of observations, find the values of the unknown parameters that maximize the likelihood of the observed sample.

In most cases it is easier to work with the log likelihood function $l(\psi|y) = \ln L(\psi|y)$. Due to the fact that the logarithm is a strictly monotone increasing function, the maximization of the log likelihood function is equivalent to directly maximizing $L(\psi|y)$. Let $\tilde{\psi}$ denote the parameter vector that maximizes the likelihood for the observed sample y. The ML estimator satisfies the so-called likelihood equations, which are obtained by differentiating $l(\psi|y)$ with respect to each of the unknown parameters of the vector ψ and setting the derivatives to zero (Harvey, 1993)

$$\frac{\partial l(\tilde{\psi}|y)}{\partial \psi} = 0. \tag{1.8}$$

As a rule, the likelihood equations are non-linear. Therefore, the ML estimates must be found in the course of an iterative procedure. This is true for the exact likelihood function of every Gaussian ARMA(p,q) process (see Hamilton, 1994, Chapter 5).

As already mentioned above, there are two different likelihood functions in use: the conditional and the exact likelihood function. Both alternatives can be estimated using XploRe.

In many applications of ARMA models the **conditional likelihood function** is an alternative to the exact likelihood function. In that case, one assumes that the first p observations of a Gaussian ARMA(p,q) process are deterministic and are equal to its observed values y. The initial residuals a_t for $t \in \{p, \ldots, p - q + 1\}$ are set to its expected values 0. In that case, the log likelihood function is

$$l(\psi) = -\frac{1}{2}(T-p)\ln 2\pi - \frac{1}{2}(T-p)\ln \sigma^2 - \frac{S(\psi')}{2\sigma^2} \tag{1.9}$$

where $\psi = (\psi', \sigma^2)$ and $S(\psi')$ denotes the sum of squares

$$S(\psi') = \sum_{t=p+1}^{T} \left(a_t(\psi')\right)^2. \tag{1.10}$$

The notation $a_t(\psi')$ emphasizes that a_t is no longer a disturbance, but a residual which depends on the value taken by the variables in ψ'.

Note, that the parameter σ^2 is an additional one, that is not included in vector ψ'. It is easy to see that (1.9) is maximized with respect to ψ' if the sum of squares $S(\psi')$ is minimized. Using the condition (1.8), this leads to

$$\sum_{t=1}^{T} \frac{\partial a_t(\psi')}{\partial \psi'} \, a_t = 0. \tag{1.11}$$

Thus, the ML estimate for ψ' can be obtained by minimizing (1.10). Furthermore, we obtain from (1.9) and (1.8) that

$$\tilde{\sigma}^2 = \frac{S(\tilde{\psi}')}{T-p} \; . \tag{1.12}$$

Thus, given the parameter vector $\tilde{\psi}'$ that maximizes the sum of squares—and thus the conditional log likelihood function (1.9)—one divides the sum of squares by $T-p$ to obtain the ML estimate $\tilde{\sigma}^2$. Another estimator for the variance of the innovations corrects furthermore for the number of estimated coefficients k.

As already mentioned, one approach to calculate the **exact log likelihood** is to use the Kalman filter. We want to show this for the airline model. This model can be written in state space form (SSF) as

$$\alpha_t = \begin{bmatrix} 0 & I_{13} \\ 0 & 0 \end{bmatrix} \alpha_{t-1} + \begin{bmatrix} 1 & \theta_1 & 0 & \ldots & \theta_{s,1} & \theta_1\theta_{s,1} \end{bmatrix}^\top a_t \tag{1.13a}$$

$$y_t = \begin{bmatrix} 1 & 0 & \ldots & 0 \end{bmatrix} \alpha_t \tag{1.13b}$$

where I_{13} is an (13×13) identity matrix (Koopman, Shephard and Doornik, 1999). Here,

$$\begin{bmatrix} 0 & I_{13} \\ 0 & 0 \end{bmatrix} \tag{1.14}$$

is the so-called transition matrix.

Given Gaussian error terms a_t, the value of the log likelihood function $l(\psi)$ for the above given state space form is

$$-\frac{T}{2}\ln 2\pi - \frac{1}{2}\sum_{t=1}^{T} \ln |F_t| - \frac{1}{2}\sum_{t=1}^{T} v_t^\top F_t^{-1} v_t \; . \tag{1.15}$$

Here,

$$v_t \stackrel{\text{def}}{=} y_t - Z_t \mathrm{E}[\alpha_t | \mathcal{F}_{t-1}]$$

are the *innovations* of the Kalman filtering procedure and $\mathcal{F}_{t-1}$ is the information set up to $t-1$. Z_t is the matrix from the above given state space form that contains the identity matrix. The matrix F_t is the covariance matrix of the innovations in period t and it is a by-product of the Kalman filter. The above log likelihood is known as the *prediction error decomposition form* (Harvey, 1989).

Given some initial values for ψ, the above exact log likelihood is evaluated with the Kalman filter. Using numerical methods, the function is maximized with respect to ψ.

Once the ML estimate $\tilde{\psi}$ is calculated, one wants to have standard errors for testing purposes. If T is sufficiently large, the ML estimate $\tilde{\psi}$ is approximately normally distributed with the inverse of the information matrix divided by T as covariance matrix. The inverse of the Hessian for $l(\tilde{\psi}|y)$ is one way to estimate the covariance matrix (Hamilton, 1994, Section 5.8). One can calculate the Hessian applying numerical methods.

1.4.2 Setting the multiplicative SARIMA model

```
msarimamodelOut = msarimamodel(d,arma,season)
```
sets the coefficients of a multiplicative seasonal ARIMA model

The original model is specified by means of the quantlet `msarimamodel`. The three arguments are lists that give the difference orders (d, D), the ordinary ARMA parts $\Phi(L)$ and $\Theta(L)$, and the seasonal AR and MA polynomials $\Phi_s(L)$ and $\Theta_s(L)$. If the model has no seasonal difference, one just omits D.

The `arma` list has at most four elements: the first element is a vector that specifies the lags of the AR polynomial $\Phi(L)$ that are not zero. The second element is a vector that specifies the lags of the MA polynomial $\Theta(L)$. If the model has only one polynomial, one sets the lags of the other one to 0.

The third element of the `arma` list is a vector with

- the AR coefficients
    i) if the model has both an AR and a MA polynomial
    ii) if the model has only an AR polynomial
- the MA coefficients if the model has only a MA polynomial .

If the model has both an AR and a MA polynomial then the fourth argument is necessary. It is a list that contains the coefficients of the MA polynomial. For example,

```
arma = list((1|3),0,(0.1|-0.25))
```

specifies an ARMA(3,0) part with $\phi_2 = 0$, $\phi_1 = 0.1$ and $\phi_3 = -0.25$.

The last argument `season` is a list that contains the information concerning the seasonal AR and MA polynomials. This list has at most five elements: the first element specifies the season s. If the data show no seasonal pattern, one sets $s = 0$ as the only argument of the list `season`. The second element is the lag structure of the seasonal AR polynomial. You have to fill in the lags that are different from zero. The third element is the lag structure of the seasonal MA polynomial. The last two elements are for the coefficient vectors of the polynomials. As explained for the `arma` list, one can omit the respective vector if the model has only one polynomial. For example,

```
season = list(12,0,(1|4),(-0.3|0.1))
```

gives a model with no seasonal AR polynomial and with the seasonal MA(4) polynomial

$$1 - 0.3L^{12} + 0.1L^{48} \ . \tag{1.16}$$

To understand what `msarimamodel` does, let us assume that the multiplicative SARIMA model is given as

$$(1 - 0.1L + 0.25L^3)\Delta x_t = (1 - 0.3L^{12} + 0.1L^{48})\varepsilon_t \ . \tag{1.17}$$

Here $d = 1$ and $D = 0$, so that we can set `d=1`. The lists for the ordinary and seasonal polynomials are given above. To have a look at the output of `msarimamodel`, one just compiles the following piece of XploRe code

```
arma   = list((1|3),0,(0.1|-0.25))   ; ordinary ARMA part
season = list(12,0,(1|4),(-0.3|0.1)) ; seasonal ARMA part
msarimamodelOut = msarimamodel(1,arma,season)
msarimamodelOut ; shows the list msarimamodelOut in the output window
```

The output is

```
Contents of msarimamodelOut.d
[1,]        1
Contents of msarimamodelOut.arlag
[1,]        1
[2,]        3
Contents of msarimamodelOut.malag
[1,]        0
Contents of msarimamodelOut.s
```

```
[1,]       12
Contents of msarimamodelOut.sarlag
[1,]        0
Contents of msarimamodelOut.smalag
[1,]        1
[2,]        4
Contents of msarimamodelOut.phi
[1,]      0.1
[2,]    -0.25
Contents of msarimamodelOut.theta
[1,]        0
Contents of msarimamodelOut.Phi
[1,]        0
Contents of msarimamodelOut.Theta
[1,]     -0.3
[2,]      0.1
```

and it resembles our example in general notation (see equation (1.4)). The list `msarimamodelOut` is an easy way to check the correct specification of the model.

1.4.3 Setting the expanded model

```
{y,phiconv,thetaconv,k} = msarimaconvert(x,msarimamodelOut)
```

sets the coefficients of an expanded multiplicative seasonal ARIMA model

If you want to estimate the coefficients in (1.4), you have to telescope out the original model. Given the specification by the list `msarimaOut` (see Subsection 1.4.2) and the time series $\{x_t\}_{t=1}^T$, the quantlet `msarimaconvert` telescopes out the original model automatically.

Let us consider the following SARIMA(3,1,1)×(12,1,0,0) model with $\phi_2 = 0$:

$$(1 - \phi_{s,1}L^{12})(1 - \phi_1 L - \phi_3 L^3)\Delta x_t = (1 + \theta_1 L)a_t. \tag{1.18}$$

Telescoping out the polynomials on the left-hand side leads to an ordinary

ARMA model:

$$(1 - \phi_1^e L - \phi_3^e L^3 - \phi_{12}^e L^{12} - \phi_{13}^e L^{13} - \phi_{15}^e L^{15}) y_t = (1 + \theta_1^e L) a_t \qquad (1.19)$$

with $y_t \stackrel{\text{def}}{=} \Delta x_t$, $\phi_1^e \stackrel{\text{def}}{=} \phi_1$, $\phi_3^e \stackrel{\text{def}}{=} \phi_3$, $\phi_{12}^e \stackrel{\text{def}}{=} \phi_{s,1}$, $\phi_{13}^e \stackrel{\text{def}}{=} -\phi_1\phi_{s,1}$, $\phi_{15}^e \stackrel{\text{def}}{=} -\phi_3\phi_{s,1}$, and $\theta_1^e \stackrel{\text{def}}{=} \theta_1$.

The superscript e denotes the coefficients of the expanded model. The output of the quantlet is thus self-explaining: the series y_t is just the differenced original series x_t and the other two outputs are the vector ϕ^e (with $\phi_0^e \stackrel{\text{def}}{=} 1$) and the vector θ^e (with $\theta_0^e \stackrel{\text{def}}{=} 1$). The first vector has the dimension $(sP + p + 1) \times 1$ and second one has the dimension $(sQ + q + 1) \times 1$. The scalar k gives the number of coefficients in the original model. For the above given example, we have $k = 4$, whereas the number of coefficients of the expanded model is 6. Later on, we need k for the calculation of some regression diagnostics.

1.4.4 The conditional sum of squares

```
{S,dia} = msarimacond(y,phiconv,thetaconv,mu{,k})
```
calculates the conditional sum of squares for given vectors of coefficients

The sum of squares is a criterion that can be used to identify the coefficients of the best model. The output of the quantlet is the conditional sum of squares for a given model specification (Box and Jenkins, 1976, Chapter 7).

For an ARMA(p,q) model this sum is just

$$S(\psi') \stackrel{\text{def}}{=} \sum_{t=p+1}^{T} \left(a_t(\psi')\right)^2 = \sum_{t=p+1}^{T} \left(\phi^e(L) y_t - \mu - \theta_{-1}^e(L) a_t\right)^2 . \qquad (1.20)$$

Here T denotes the number of observations y_t. Recall that the first entries of the lag-polynomials are for $L^0 = 1$. $\theta_{-1}^e(L)$ denotes the MA-polynomial without the first entry. The first q residuals are set to zero.

The arguments of the quantlet are given by the output of `msarimaconvert`. `mu` is the mean of the series $\{y_t\}_{t=1}^T$. k is the number of coefficients in the original model and will be used to calculate some regression diagnostics. This argument is optional. If you do not specify k, the number of coefficients in the expanded model is used instead.

Furthermore, the quantlet `msarimacond` gives the list `dia` that contains several diagnostics. After the maximization of the conditional sum of squares, one can use these diagnostics to compare different specifications. In the ongoing k denotes the number of ARMA parameters that are different from zero. In our example we have $k = 2$ for both specifications.

The first element of the list—that is `dia.s2`—is the estimated variance of the residuals

$$\hat{\sigma}_a^2 = \frac{S}{T-p-k} \; . \tag{1.21}$$

The second element `dia.R2` gives the coefficient of determination

$$R^2 = 1 - \frac{S}{(T-p-1)\hat{\sigma}_y^2} \; . \tag{1.22}$$

The variance of the dependent variables y_t is calculated for the observations starting with $t = p+1$. It is possible in our context that R^2 becomes negative.

The adjusted coefficient of determination $\bar{R}^2$ is calculated as

$$\bar{R}^2 = 1 - (1-R^2)\frac{T-p-1}{T-p-k} \; . \tag{1.23}$$

It is the third argument of the list and is labeled `dia.aR2`.

The fourth element `dia.logl` gives the values of the log likelihood function evaluated at $\tilde{\psi}$. Given the likelihood function (1.9), $\tilde{\sigma}^2$ is a function of $\tilde{\psi}'$. To take this into account, we plug (1.12) into (1.9) and obtain

$$l(\tilde{\psi}) = -\frac{T-p}{2}\left[1 + \ln 2\pi + \ln\left\{\frac{S(\tilde{\psi}')}{T-p}\right\}\right] \; . \tag{1.24}$$

This expression is the value of the log likelihood function.

The fifth element `dia.AIC` gives the Akaike Information Criteria (AIC)

$$\text{AIC} = -\frac{2\{l(\tilde{\psi}) - k\}}{T-p} \; . \tag{1.25}$$

The sixth element `dia.SIC` gives the Schwarz Information Criteria (SIC)

$$\text{SIC} = -\frac{2\{l(\tilde{\psi}) - k\ln(T-p)\}}{T-p} \; . \tag{1.26}$$

For both criteria see Shumway and Stoffer (2000). These criteria can be used for model selection (Durbin and Koopman (2001)). Eventually, the last element `dia.a` gives the $(T-p) \times 1$ vector of the residuals a_t.

Now we can come back to our example of the airline data: recall that we have identified two possible specifications for this data set. The first specification is a SARIMA(0,1,1)×(12,0,1,1) with $\psi_1 = (\theta_1, \theta_{s,1}, \sigma^2)$. The second is a SARIMA(0,1,12)×(12,0,1,0) with $\psi_2 = (\theta_1, \theta_{12}, \sigma^2)$.

We maximize the conditional sum of squares for both specifications using the BFGS algorithm. Given the estimates $\tilde{\psi}$, the standard errors are obtained by means of the Hessian matrix for the log likelihood. The Hessian is calculated for this function using numerical methods. The square roots of the diagonal elements of the inverse Hessian are the standard errors of the estimates $\tilde{\psi}'$.

The results of the first specification are presented in the following table:

```
"===================================================================="
" Estimation results for the SARIMA(0,1,1)x(12,0,1,1) specification "
"===================================================================="
" Convergence achieved after 11 iterations                           "
" 131 observations included                                          "
"                                                                    "
"      Variable       Coefficient        t-stat      p-value         "
"      ----------------------------------------------------          "
"      theta_1          -0.3776          -4.3206       0.00          "
"      theta_s,1        -0.5728          -8.2073       0.00          "
"      ----------------------------------------------------          "
"      Sum of squared resids  0.1819       s2        0.0014          "
"      R2                     0.3343       adj. R2   0.3292          "
"      AIC                   -3.7110       SIC      -3.6672          "
"===================================================================="
```

XEGmsarima9.xpl

The results of the second specification are given in the next table:

```
"===================================================================="
" Estimation results for the SARIMA(0,1,12)x(12,0,1,0) specification"
"===================================================================="
" Convergence achieved after 10 iterations                          "
" 131 observations included                                         "
"                                                                   "
"      Variable       Coefficient        t-stat      p-value        "
"      -----------------------------------------------------        "
"      theta_1          -0.2464         -3.6852        0.00         "
"      theta_12         -0.5080         -7.9028        0.00         "
"      -----------------------------------------------------        "
"      Sum of squared resids  0.1917        s2        0.0015        "
"      R2                     0.2984        adj. R2   0.2930        "
"      AIC                   -3.6585        SIC      -3.6146        "
"===================================================================="
```

XEGmsarima10.xpl

It is obvious that both specifications deliver good results. However, the sum of squared resids is smaller for the specification with a seasonal MA term. Additionally, both information criteria indicate that this specification is slightly better.

1.4.5 The extended ACF

`eacf(y,p,q)`
displays a table with the extended ACF for time series y_t

After estimating the unknown parameters $\tilde{\psi}$ for competing specifications, one should have a look at the residual series $\{\tilde{a}_t\}_{t=p+1}^{T}$. They should behave like a white noise process and should exhibit no autocorrelation. In order to check for autocorrelation, one could use the ACF and PACF. However, the extended autocorrelation function (EACF) is also a convenient tool for inspecting time series (Peña, Tiao and Tsay, 2001, Chapter 3) that show no seasonal pattern.

In general, the EACF allows for the identification of ARIMA models (differencing is not necessary). The quantlet `eacf` generates a table of the sample

EACF for a time series. You have to specify the maximal number of AR lags (p) and MA lags (q). Every row of the output table gives the ACF up to q lags for the residuals of an AR regression with $k \leqslant p$ lags. Furthermore, the simplified EACF is tabulated. If an autocorrelation is significant according to Bartlett's formula the entry is 1. Otherwise the entry is 0. Bartlett's formula for an MA(q) is given as

$$\text{Var}[\hat{\rho}(q+1)] = \frac{1}{T}\left[1 + 2\sum_{j=1}^{q}\hat{\rho}(j)^2\right]$$

where T is the number of observations (Peña, Tiao and Tsay, 2001). For identification, look for the vertex of a triangle of zeros. You can immediately read off the order of the series from the table.

We use the EACF to explore the behavior of the residuals of both specifications. The next table shows the EACF of the residuals that come from the SARIMA(0,1,1) ×(12,0,1,1) specification.

EACF

q=	0	1	2	3	4	5	6	7	8	9	10
p=0	+0.01	+0.03	-0.12	-0.10	+0.08	+0.08	-0.05	-0.02	+0.11	-0.05	+0.02
p=1	-0.30	+0.00	-0.09	-0.13	+0.09	+0.08	-0.07	+0.00	+0.10	+0.00	-0.03
p=2	+0.21	+0.09	-0.05	-0.06	-0.01	+0.02	+0.04	+0.01	+0.09	-0.01	-0.01
p=3	-0.50	+0.43	-0.26	-0.06	-0.03	-0.01	+0.04	-0.04	+0.06	-0.02	+0.03
p=4	+0.50	+0.48	-0.18	+0.11	+0.01	-0.02	+0.03	+0.00	+0.07	+0.02	+0.04
p=5	-0.50	+0.47	+0.04	+0.05	+0.12	-0.02	+0.02	-0.04	+0.06	-0.02	+0.00
p=6	+0.48	+0.43	-0.02	+0.11	+0.11	+0.08	-0.01	-0.07	+0.07	+0.00	-0.03

q=	0	1	2	3	4	5	6	7	8	9	10
p=0	0	0	0	0	0	0	0	0	0	0	0
p=1	1	0	0	0	0	0	0	0	0	0	0
p=2	1	0	0	0	0	0	0	0	0	0	0
p=3	1	1	1	0	0	0	0	0	0	0	0
p=4	1	1	0	0	0	0	0	0	0	0	0
p=5	1	1	0	0	0	0	0	0	0	0	0
p=6	1	1	0	0	0	0	0	0	0	0	0

XEGmsarima11.xpl

It is obvious, that the vertex of zeros is at position $q = 0$ and $p = 0$. Thus we conclude that the residuals are white noise. Notice, that the first line in the above table at $p = 0$ just gives the ACF of the residual series. According to Bartlett's formula, all autocorrelation coefficients are not significantly different from zero.

The next table gives the EACF of the residuals of the SARIMA(0,1,12)×(12,0,1,0) specification.

EACF

q=	0	1	2	3	4	5	6	7	8	9	10
p=0	-0.12	+0.03	-0.15	-0.08	+0.08	+0.04	-0.07	-0.02	+0.12	-0.06	+0.03
p=1	+0.13	-0.02	-0.12	-0.14	+0.08	+0.05	-0.07	-0.01	+0.11	-0.02	-0.03
p=2	+0.10	-0.36	-0.13	-0.05	-0.05	-0.02	+0.03	+0.00	+0.09	-0.02	-0.06
p=3	-0.49	+0.40	-0.25	+0.04	-0.02	-0.02	+0.01	-0.01	+0.04	+0.00	+0.02
p=4	+0.42	+0.32	-0.27	-0.23	-0.01	-0.03	+0.02	+0.00	+0.04	+0.01	+0.03
p=5	-0.48	+0.25	-0.28	+0.05	-0.06	-0.04	+0.03	-0.03	+0.03	+0.00	+0.04
p=6	+0.37	+0.39	-0.05	-0.12	+0.07	+0.13	-0.05	-0.02	+0.05	+0.01	+0.03

q=	0	1	2	3	4	5	6	7	8	9	10
p=0	0	0	0	0	0	0	0	0	0	0	0
p=1	0	0	0	0	0	0	0	0	0	0	0
p=2	0	1	0	0	0	0	0	0	0	0	0
p=3	1	1	1	0	0	0	0	0	0	0	0
p=4	1	1	1	1	0	0	0	0	0	0	0
p=5	1	1	1	0	0	0	0	0	0	0	0
p=6	1	1	0	0	0	0	0	0	0	0	0

XEGmsarima12.xpl

We can conclude that the vertex of zeros is at position $q = 0$ and $p = 0$. Thus the residuals are white noise. According to Bartlett's formula, once again all autocorrelation coefficients are not significantly different from zero.

1.4.6 The exact likelihood

```
{gkalfilOut,loglike} = gkalfilter(Y,mu,Sig,ca,Ta,Ra,
                                  da,Za,Ha,l)
```
Kalman filters a SSF and gives the value of the log likelihood

As already mentioned in the introductory part of this section, the Kalman filter can be used to evaluate the exact log likelihood function. For the estimation of the unknown parameters the evaluated log likelihood function 1.15 is required. The second element of the quantlet provides the value of the exact log likelihood function.

We now shortly describe the procedure of the Kalman filter and the implementation with `gkalfilter`. Good references for the Kalman filter are—in addition to Harvey (1989)—Hamilton (1994), Gourieroux and Monfort (1997) and Shumway and Stoffer (2000). The first argument is an array with the observed time series. The vector `mu` specifies the initial conditions of the filtering procedure with corresponding covariance matrix `Sig`. Due to the fact that our SSF (1.13) contains no constants, we set `ca` and `da` to zero. Furthermore, we have no disturbance in the measurement equation—that is the equation for y_t in (1.13)—so we also set `Ha` to zero. The covariance matrix for the disturbance in the state equation is given as

$$R = \sigma^2 \begin{bmatrix} 1 & \theta_1 & 0 & \dots & 0 & \theta_{s,1} & \theta_1\theta_{s,1} \\ \theta_1 & \theta_1^2 & 0 & \dots & 0 & \theta_1\theta_{s,1} & \theta_1^2\theta_{s,1} \\ 0 & 0 & 0 & \dots & 0 & 0 & 0 \\ \vdots & \vdots & \vdots & \dots & \vdots & \vdots & \vdots \\ \theta_{s,1} & \theta_1\theta_{s,1} & 0 & \dots & 0 & \theta_{s,1}^2 & \theta_1\theta_{s,1}^2 \\ \theta_1\theta_{s,1} & \theta_1^2\theta_{s,1} & 0 & \dots & 0 & \theta_1\theta_{s,1}^2 & \theta_1^2\theta_{s,1}^2 \end{bmatrix}$$

Eventually, `Za` is an array for the matrix Z given in the measurement equation.

The state space form of the airline model is given in (1.13). It is a well known fact that the product of eigenvalues for a square matrix is equal to its determinant. The determinant of the transition matrix T for our model—given in equation (1.14)—is zero and so all eigenvalues are also zero. Thus our system is stable (Harvey, 1989, p. 114). In that case, we should set the initial values to the unconditional mean and variance of the state vector (Koopman, Shephard and Doornik, 1999). We easily obtain for our model (1.13) that

$$\mu \stackrel{\text{def}}{=} \mathrm{E}[\alpha_0] = 0 \tag{1.27}$$

and

$$\Sigma \stackrel{\text{def}}{=} \mathrm{Var}[\alpha_0] = T\Sigma T^\top + R\,.$$

A way to solve for the elements of Σ is

$$\mathrm{vec}(\Sigma) = (I - T \otimes T)^{-1}\mathrm{vec}(R)\,. \tag{1.28}$$

Here, vec denotes the vec-operator that places the columns of a matrix below each other and $\otimes$ denotes the Kronecker product.

For the estimation, we use the demeaned series of the growth rates g_t of the airline data. The standard errors of the estimates are given by the square roots of the diagonal elements of the inverse Hessian evaluated at $\tilde{\psi}'$. The following table shows the results:

```
"===================================================================="
" Estimation results for the SARIMA(0,1,1)x(12,0,1,0) specification "
" Exact Log Likelihood function is maximized                         "
"===================================================================="
" Convergence achieved after 12 iterations                           "
" 131 observations included                                          "
"                                                                    "
"      Variable      Coefficient       t-stat      p-value           "
"      -----------------------------------------------------         "
"      theta_1          -0.3998        -4.4726        0.00           "
"      theta_s,1        -0.5545        -7.5763        0.00           "
"      sigma2            0.0014         8.0632        0.00           "
"      -----------------------------------------------------         "
"      AIC              -3.6886        SIC          -3.5111          "
"===================================================================="
```

XEGmsarima13.xpl

The estimators are only slightly different from the estimators we have calculated with the conditional likelihood. The variance $\tilde{\sigma}^2$ is identical.

Bibliography

Box, G. E. P. and Jenkins, G. M. (1976). *Time Series Analysis, Forecasting and Control*, Revised Edition, Prentice Hall, Englewood Cliffs.

Chatfield, C. (2001). *Time-Series Forecasting*, Chapman & Hall/CRC, Boca Raton.

Diebold, F. X. (1998). *Elements of Forecasting*, South-Western College Publishing, Cincinnati.

Durbin, J. and Koopman, S J. (2001). *Time Series Analysis by State Space Methods*, Oxford Statistical Science Series 24, Oxford University Press, Oxford.

Franses, P. H. (1998). *Time Series Models for Business and Economic Forecasting*, Cambridge University Press, Cambridge.

Ghysels, E. and Osborn, D. R. (2000). *The Econometric Analysis of Seasonal Time Series*, Cambridge University Press, Cambridge.

Gourieroux, C. and Monfort, A. (1997). *Time Series and Dynamic Models*, Cambridge University Press, Cambridge.

Hamilton, J. D. (1994). *Time Series Analysis*, Princeton University Press, Princeton, New Jersey.

Harvey, A. C. (1989). *Forecasting, Structural Time Series Models and the Kalman Filter*, Cambridge University Press, Cambridge.

Harvey, A. C. (1993). *Time Series Models*, Harvester/Wheatsheaf, New York, London.

Koopman, S. J., Shepard, N. and Doornik, J. A. (1999). Statistical Algorithms for Models in State Space Using SsfPack 2.2, *Econometrics Journal* **2**: 107–160.

Mills, T. C. (1990). *Time series techniques for economists*, Cambridge University Press, Cambridge.

Peña, D., Tiao, G. C. and Tsay, R. S. (2001). *A Course in Time Series Analysis*, Wiley, New York.

Shumway, R. H. and Stoffer, D. S. (2000). *Time Series Analysis and Its Applications*, Springer, New York, Berlin.