NAG C Library Chapter Introduction

g13 – Time Series Analysis

Contents

1 Scope of the Chapter

This chapter provides facilities for investigating and modelling the statistical structure of series of observations collected at equally spaced points in time. The models may then be used to forecast the series.

The chapter covers the following models and approaches.

1. Univariate time series analysis, including autocorrelation functions and autoregressive moving average (ARMA) models.
2. Univariate spectral analysis.
3. Transfer function (multi-input) modelling, in which one time series is dependent on other time series.
4. Bivarate spectral methods including coherency, gain and input response functions.
5. Descriptive methods for multivariate time series.
6. Kalman filter models.
7. GARCH models for volatility.

2 Background to the Problems

2.1 Univariate Analysis

Let the given time series be $x_1, x_2, \ldots, x_n$, where n is its length. The structure which is intended to be investigated, and which may be most evident to the eye in a graph of the series, can be broadly described as

(a) trends – linear or possibly higher-order polynomial;

(b) seasonal patterns, associated with fixed integer seasonal periods. The presence of such seasonality and the period will normally be known *a priori*. The pattern may be fixed, or slowly varying from one season to another;

(c) cycles, or waves of stable amplitude and period p (from peak to peak). The period is not necessarily integer, the corresponding absolute frequency (cycles/time unit) being $f = 1/p$ and angular frequency $\omega = 2\pi f$. The cycle may be of pure sinusoidal form like $\sin(\omega t)$, or the presence of higher harmonic terms may be indicated, e.g., by asymmetry in the wave form;

(d) quasi-cycles, i.e., waves of fluctuating period and amplitude; and

(e) irregular statistical fluctuations and swings about the overall mean or trend.

Trends, seasonal patterns, and cycles might be regarded as **deterministic** components following fixed mathematical equations, and the quasi-cycles and other statistical fluctuations as **stochastic** and describable by short-term correlation structure. For a finite data set it is not always easy to discriminate between these two types, and a common description using the class of autoregressive integrated moving-average (ARIMA) models is now widely used. The form of these models is that of difference equations (or recurrence relations) relating present and past values of the series. The user is referred to Box and Jenkins (1976) for a thorough account of these models and how to use them. We follow their notation and outline the recommended steps in ARIMA model building for which functions are available.

2.1.1 Transformations

If the variance of the observations in the series is not constant across the range of observations it may be useful to apply a variance-stabilizing transformation to the series. A common situation is for the variance to increase with the magnitude of the observations and in this case typical transformations used are the log or square root transformation. A range–mean or standard deviation–mean plot provides a quick and easy way of detecting non-constant variance and of choosing, if required, a suitable transformation. This is a plot of the range or standard deviation of successive groups of observations against their means.

2.1.2 Differencing operations

These may be used to simplify the structure of a time series.

First-order differencing, i.e., forming the new series

$$\nabla x_t = x_t - x_{t-1}$$

will remove a linear trend. First-order seasonal differencing

$$\nabla_s x_t = x_t - x_{t-s}$$

eliminates a fixed seasonal pattern.

These operations reflect the fact that it is often appropriate to model a time series in terms of changes from one value to another. Differencing is also therefore appropriate when the series has something of the nature of a random walk, which is by definition the accumulation of independent changes.

Differencing may be applied repeatedly to a series, giving

$$w_t = \nabla^d \nabla_s^D x_t$$

where d and D are the orders of differencing. The derived series w_t will be shorter, of length $N = n - d - s \times D$, and extend for $t = 1 + d + s \times D, \ldots, n$.

2.1.3 Sample autocorrelations

Given that a series has (possibly as a result of simplifying by differencing operations) a homogeneous appearance throughout its length, fluctuating with approximately constant variance about an overall mean level, it is appropriate to assume that its statistical properties are **stationary**. For most purposes the correlations ρ_k between terms x_t, x_{t+k} or w_t, w_{t+k} separated by lag k give an adequate description of the statistical structure and are estimated by the sample autocorrelation function (acf) r_k, for $k = 1, 2, \ldots$.

As described by Box and Jenkins (1976), these may be used to indicate which particular ARIMA model may be appropriate.

2.1.4 Partial autocorrelations

The information in the autocorrelations ρ_k may be presented in a different light by deriving from them the coefficients of the partial autocorrelation function (pacf) $\phi_{k,k}$, for $k = 1, 2, \ldots$. $\phi_{k,k}$ measures the correlation between x_t and x_{t+k} conditional upon the intermediate values $x_{t+1}, x_{t+2}, \ldots, x_{t+k-1}$. The corresponding sample values $\hat{\phi}_{k,k}$ give further assistance in the selection of ARIMA models.

Both acf and pacf may be rapidly computed, particularly in comparison with the time taken to estimate ARIMA models.

2.1.5 Finite lag predictor coefficients and error variances

The partial autocorrelation coefficient $\phi_{k,k}$ is determined as the final parameter in the minimum variance predictor of x_t in terms of $x_{t-1}, x_{t-2}, \ldots, x_{t-k}$,

$$x_t = \phi_{k,1} x_{t-1} + \phi_{k,2} x_{t-2} + \cdots + \phi_{k,k} x_{t-k} + e_{k,t}$$

where $e_{k,t}$ is the prediction error, and the first subscript k of $\phi_{k,i}$ and $e_{k,t}$ emphasizes the fact that the parameters will alter as k increases. Moderately good estimates $\hat{\phi}_{k,i}$ of $\phi_{k,i}$ are obtained from the sample acf, and after calculating the pacf up to lag L, the successive values $v_1, v_2, \ldots, v_L$ of the prediction error variance estimates, $v_k = \text{var}(e_{k,t})$, are available, together with the final values of the coefficients $\hat{\phi}_{k,1}, \hat{\phi}_{k,2}, \ldots, \hat{\phi}_{k,L}$. If x_t has non-zero mean, $\bar{x}$, it is adequate to use $x_t - \bar{x}$ in place of x_t in the prediction equation.

Although Box and Jenkins (1976) do not place great emphasis on these prediction coefficients, their use is advocated for example by Akaike (1971), who recommends selecting an optimal order of the predictor as the lag for which the final prediction error (FPE) criterion $(1 + k/n)(1 - k/n)^{-1} v_k$ is a minimum.

2.1.6 ARIMA models

The correlation structure in stationary time series may often be represented by a model with a small number of parameters belonging to the autoregressive moving-average (ARMA) class. If the stationary series w_t has been derived by differencing from the original series x_t, then x_t is said to follow an ARIMA

model. Taking $w_t = \nabla^d x_t$, the (non-seasonal) ARIMA (p, d, q) model with p autoregressive parameters $\phi_1, \phi_2, \ldots, \phi_p$ and q moving-average parameters $\theta_1, \theta_2, \ldots, \theta_q$, represents the structure of w_t by the equation

$$w_t = \phi_1 w_{t-1} + \cdots + \phi_p w_{t-p} + a_t - \theta_1 a_{t-1} - \cdots - \theta_q a_{t-q}, \tag{1}$$

where a_t is an uncorrelated series (white noise) with mean 0 and constant variance σ_a^2. If w_t has a non-zero mean c, then this is allowed for by replacing $w_t, w_{t-1}, \ldots$ by $w_t - c, w_{t-1} - c, \ldots$ in the model. Although c is often estimated by the sample mean of w_t this is not always optimal.

A series generated by this model will only be stationary provided restrictions are placed on $\phi_1, \phi_2, \ldots, \phi_p$ to avoid unstable growth of w_t. These are called **stationarity** constraints. The series a_t may also be usefully interpreted as the linear **innovations** in x_t (and in w_t), i.e., the error if x_t were to be predicted using the information in all past values $x_{t-1}, x_{t-2}, \ldots$, provided also that $\theta_1, \theta_2, \ldots, \theta_q$ satisfy **invertibility** constraints. This allows the series a_t to be regenerated by rewriting the model equation as

$$a_t = w_t - \phi_1 w_{t-1} - \cdots - \phi_p w_{t-p} + \theta_1 a_{t-1} + \cdots + \theta_q a_{t-q}. \tag{2}$$

For a series with short-term correlation only, i.e., r_k is not significant beyond some low lag q (see Box and Jenkins (1976) for the statistical test), then the pure moving-average model MA(q) is appropriate, with no autoregressive parameters, i.e., $p = 0$.

Autoregressive parameters are appropriate when the acf pattern decays geometrically, or with a damped sinusoidal pattern which is associated with quasi-periodic behaviour in the series. If the sample pacf $\hat{\phi}_{k,k}$ is significant only up to some low lag p, then a pure autoregressive model AR(p) is appropriate, with $q = 0$. Otherwise moving-average terms will need to be introduced, as well as autoregressive terms.

The seasonal ARIMA (p, d, q, P, D, Q, s) model allows for correlation at lags which are multiples of the seasonal period s. Taking $w_t = \nabla^d \nabla_s^D x_t$, the series is represented in a two-stage manner via an intermediate series e_t:

$$w_t = \Phi_1 w_{t-s} + \cdots + \Phi_P w_{t-s\times P} + e_t - \Theta_1 e_{t-s} - \cdots - \Theta_Q e_{t-s\times Q} \tag{3}$$

$$e_t = \phi_1 e_{t-1} + \cdots + \phi_p e_{t-p} + a_t - \theta_1 a_{t-1} - \cdots - \theta_q a_{t-q} \tag{4}$$

where Φ_i, Θ_i are the seasonal parameters and P, Q are the corresponding orders. Again, w_t may be replaced by $w_t - c$.

2.1.7 ARIMA model estimation

In theory, the parameters of an ARIMA model are determined by a sufficient number of autocorrelations $\rho_1, \rho_2, \ldots$. Using the sample values $r_1, r_2, \ldots$ in their place it is usually (but not always) possible to solve for the corresponding ARIMA parameters.

These are rapidly computed but are not fully efficient estimates, particularly if moving-average parameters are present. They do provide useful **preliminary** values for an efficient but relatively slow iterative method of estimation. This is based on the least-squares principle by which parameters are chosen to minimize the sum of squares of the innovations a_t, which are regenerated from the data using (2), or the reverse of (3) and (4) in the case of seasonal models.

Lack of knowledge of terms on the right-hand side of (2), when $t = 1, 2, \ldots, \max(p, q)$, is overcome by introducing q unknown series values $w_0, w_1, \ldots, w_{1-q}$ which are estimated as nuisance parameters, and using correction for transient errors due to the autoregressive terms. If the data $w_1, w_2, \ldots, w_N = w$ is viewed as a single sample from a multivariate Normal density whose covariance matrix V is a function of the ARIMA model parameters, then the exact likelihood of the parameters is

$$-\tfrac{1}{2}\log|V| - \tfrac{1}{2} w^T V^{-1} w.$$

The least-squares criterion as outlined above is equivalent to using the quadratic form

$$QF = w^T V^{-1} w$$

as an objective function to be minimized. Neglecting the term $-\frac{1}{2}\log|V|$ yields estimates which differ very little from the exact likelihood except in small samples, or in seasonal models with a small number of

whole seasons contained in the data. In these cases bias in moving-average parameters may cause them to stick at the boundary of their constraint region, resulting in failure of the estimation method.

Approximate standard errors of the parameter estimates and the correlations between them are available after estimation.

The model residuals, $\hat{a}_t$, are the innovations resulting from the estimation and are usually examined for the presence of autocorrelation as a check on the adequacy of the model.

2.1.8 ARIMA model forecasting

An ARIMA model is particularly suited to extrapolation of a time series. The model equations are simply used for $t = n+1, n+2, \ldots$ replacing the unknown future values of a_t by zero. This produces future values of w_t, and if differencing has been used this process is reversed (the so-called integration part of ARIMA models) to construct future values of x_t.

Forecast error limits are easily deduced.

This process requires knowledge only of the model orders and parameters together with a limited set of the terms $a_{t-i}, e_{t-i}, w_{t-i}, x_{t-i}$ which appear on the right-hand side of the models (3) and (4) (and the differencing equations) when $t = n$. It does not require knowledge of the whole series.

We call this the state set. It is conveniently constituted after model estimation. Moreover, if new observations $x_{n+1}, x_{n+2}, \ldots$ come to hand, then the model equations can easily be used to update the state set before constructing forecasts from the end of the new observations. This is particularly useful when forecasts are constructed on a regular basis. The new innovations $a_{n+1}, a_{n+2}, \ldots$ may be compared with the residual standard deviation, σ_a, of the model used for forecasting, as a check that the model is continuing to forecast adequately.

2.2 Univariate Spectral Analysis

In describing a time series using spectral analysis the fundamental components are taken to be sinusoidal waves of the form $R\cos(\omega t + \phi)$, which for a given angular frequency ω, $0 \leq \omega \leq \pi$, is specified by its amplitude $R > 0$ and phase ϕ, $0 \leq \phi < 2\pi$. Thus in a time series of n observations it is not possible to distinguish more than $n/2$ independent sinusoidal components. The frequency range $0 \leq \omega \leq \pi$ is limited to a shortest wavelength of two sampling units because any wave of higher frequency is indistinguishable upon sampling (or is aliased with) a wave within this range. Spectral analysis follows the idea that for a series made up of a **finite** number of sine waves the amplitude of any component at frequency ω is given to order $1/n$ by

$$R^2 = \left(\frac{1}{n^2}\right)\left|\sum_{t=1}^{n} x_t e^{i\omega t}\right|^2.$$

2.2.1 The sample spectrum

For a series $x_1, x_2, \ldots, x_n$ this is defined as

$$f^*(\omega) = \left(\frac{1}{2n\pi}\right)\left|\sum_{t=1}^{n} x_t e^{i\omega t}\right|^2,$$

the scaling factor now being chosen in order that

$$2\int_0^{\pi} f^*(\omega)d\omega = \sigma_x^2,$$

i.e., the spectrum indicates how the sample variance (σ_x^2) of the series is distributed over components in the frequency range $0 \leq \omega \leq \pi$.

It may be demonstrated that $f^*(\omega)$ is equivalently defined in terms of the sample autocorrelation function (acf) r_k of the series as

$$f^*(\omega) = \left(\frac{1}{2\pi}\right)\left(c_0 + 2\sum_{k=1}^{n-1} c_k \cos k\omega\right),$$

where $c_k = \sigma_x^2 r_k$ are the sample autocovariance coefficients.

If the series x_t does contain a **deterministic** sinusoidal component of amplitude R, this will be revealed in the sample spectrum as a sharp peak of approximate width π/n and height $(n/2\pi)R^2$. This is called the discrete part of the spectrum, the variance R^2 associated with this component being in effect concentrated at a single frequency.

If the series x_t has no deterministic components, i.e., is purely **stochastic** being stationary with acf r_k, then with increasing sample size the expected value of $f^*(\omega)$ converges to the theoretical spectrum – the **continuous** part

$$f(\omega) = \left(\frac{1}{2\pi}\right)\left(\gamma_0 + 2\sum_{k=1}^{\infty} \gamma_k \cos(\omega k)\right),$$

where γ_k are the theoretical autocovariances.

The sample spectrum does **not** however converge to this value but at each frequency point fluctuates about the theoretical spectrum with an exponential distribution, being independent at frequencies separated by an interval of $2\pi/n$ or more. Various devices are therefore employed to smooth the sample spectrum and reduce its variability. Much of the strength of spectral analysis derives from the fact that the error limits are multiplicative so that features may still show up as significant in a part of the spectrum which has a generally low level, whereas they are completely masked by other components in the original series. The spectrum can help to distinguish deterministic cyclical components from the stochastic quasi-cycle components which produce a broader peak in the spectrum. (The deterministic components can be removed by regression and the remaining part represented by an ARIMA model.)

A large discrete component in a spectrum can distort the continuous part over a large frequency range surrounding the corresponding peak. This may be alleviated at the cost of slightly broadening the peak by tapering a portion of the data at each end of the series with weights which decay smoothly to zero. It is usual to correct for the mean of the series and for any linear trend by simple regression, since they would similarly distort the spectrum.

2.2.2 Spectral smoothing by lag window

The estimate is calculated directly from the sample covariances c_k as

$$f(\omega) = \left(\frac{1}{2\pi}\right)\left(c_0 + 2\sum_{k=1}^{M-1} w_k c_k \cos k\omega\right),$$

the smoothing being induced by the lag window weights w_k which extend up to a **truncation lag** M which is generally much less than n. The smaller the value of M, the greater the degree of smoothing, the spectrum estimates being independent only at a wider frequency separation indicated by the **bandwidth** b which is proportional to $1/M$. It is wise, however, to calculate the spectrum at intervals appreciably less than this. Although greater smoothing narrows the error limits, it can also distort the spectrum, particularly by flattening peaks and filling in troughs.

2.2.3 Direct spectral smoothing

The unsmoothed sample spectrum is calculated for a fine division of frequencies, then averaged over intervals centred on each frequency point for which the smoothed spectrum is required. This is usually at a coarser frequency division. The bandwidth corresponds to the width of the averaging interval.

2.3 Linear Lagged Relationships Between Time Series

We now consider the context in which one time series, called the dependent or output series $y_1, y_2, \ldots, y_n$, is believed to depend on one or more explanatory or input series, e.g., $x_1, x_2, \ldots, x_n$. This dependency may follow a simple linear regression, e.g.,

$$y_t = vx_t + n_t$$

or more generally may involve lagged values of the input

$$y_t = v_0x_t + v_1x_{t-1} + v_2x_{t-2} + \cdots + n_t.$$

The sequence $v_0, v_1, v_2,$ is called the **impulse response function** (IRF) of the relationship. The term n_t represents that part of y_t which cannot be explained by the input, and it is assumed to follow a univariate ARIMA model. We call n_t the (output) noise component of y_t, and it includes any constant term in the relationship. It is assumed that the input series, x_t, and the noise component, n_t, are independent.

The part of y_t which is explained by the input is called the input component z_t:

$$z_t = v_0x_t + v_1x_{t-1} + v_2x_{t-2} + \cdots$$

so $y_t = z_t + n_t$.

The eventual aim is to model both these components of y_t on the basis of observations of $y_1, y_2, \ldots, y_n$ and $x_1, x_2, \ldots, x_n$. In applications to forecasting or control both components are important. In general there may be more than one input series, e.g., $x_{1,t}$ and $x_{2,t}$, which are assumed to be independent and corresponding components $z_{1,t}$ and $z_{2,t}$, so

$$y_t = z_{1,t} + z_{2,t} + n_t.$$

2.3.1 Transfer function models

In a similar manner to that in which the structure of a univariate series may be represented by a finite-parameter ARIMA model, the structure of an input component may be represented by a **transfer function** (TF) model with delay time b, p autoregressive-like parameters $\delta_1, \delta_2, \ldots, \delta_p$ and $q+1$ moving-average-like parameters $\omega_0, \omega_1, \ldots, \omega_q$:

$$z_t = \delta_1 z_{t-1} + \delta_2 z_{t-2} + \cdots + \delta_p z_{t-p} + \omega_0 x_{t-b} - \omega_1 x_{t-b-1} - \cdots - \omega_q x_{t-b-q}. \tag{5}$$

If $p > 0$ this represents an IRF which is infinite in extent and decays with geometric and/or sinusoidal behaviour. The parameters $\delta_1, \delta_2, \ldots, \delta_p$ are constrained to satisfy a stability condition identical to the stationarity condition of autoregressive models. There is no constraint on $\omega_0, \omega_1, \ldots, \omega_q$.

2.3.2 Cross-correlations

An important tool for investigating how an input series x_t affects an output series y_t is the sample **cross-correlation function** (CCF) $r_{xy}(k)$, for $k = 0, 1, 2, \ldots$ between the series. If x_t and y_t are (jointly) stationary time series this is an estimator of the theoretical quantity

$$\rho_{xy}(k) = \text{corr}(x_t, y_{t+k}).$$

The sequence $r_{yx}(k)$, for $k = 0, 1, 2, \ldots$, is distinct from $r_{xy}(k)$, though it is possible to interpret

$$r_{yx}(k) = r_{xy}(-k).$$

When the series y_t and x_t are believed to be related by a transfer function model, the CCF is determined by the IRF $v_0, v_1, v_2, \ldots$ and the autocorrelation function of the input x_t.

In the **particular** case when x_t is an uncorrelated series or **white noise** (and is uncorrelated with any other inputs):

$$\rho_{xy}(k) \propto v_k$$

and the sample CCF can provide an estimate of v_k:

$$\tilde{v}_k = (s_y/s_x)r_{xy}(k)$$

where s_y and s_x are the sample standard deviations of y_t and x_t, respectively.

In theory the IRF coefficients $v_b, \ldots, v_{b+p+q}$ determine the parameters in the TF model, and using $\tilde{v}_k$ to estimate $\tilde{v}_k$ it is possible to solve for **preliminary** estimates of $\delta_1, \delta_2, \ldots, \delta_p$, $\omega_0, \omega_1, \ldots, \omega_q$.

2.3.3 Prewhitening or filtering by an ARIMA model

In general an input series x_t is not white noise, but may be represented by an ARIMA model with innovations or residuals a_t which are white noise. If precisely the same operations by which a_t is generated from x_t are applied to the output y_t to produce a series b_t, then the transfer function relationship between y_t and x_t is preserved between b_t and a_t. It is then possible to estimate

$$\tilde{v}_k = (s_b/s_a)r_{ab}(k).$$

The procedure of generating a_t from x_t (and b_t from y_t) is called prewhitening or filtering by an ARIMA model. Although a_t is necessarily white noise, this is not generally true of b_t.

2.3.4 Multi-input model estimation

The term multi-input model is used for the situation when one output series y_t is related to one or more input series $x_{j,t}$, as described in Section 2.3. If for a given input the relationship is a simple linear regression, it is called a simple input; otherwise it is a transfer function input. The error or noise term follows an ARIMA model.

Given that the orders of all the transfer function models and the ARIMA model of a multi-input model have been specified, the various parameters in those models may be (simultaneously) estimated.

The procedure used is closely related to the least-squares principle applied to the innovations in the ARIMA noise model.

The innovations are derived for any proposed set of parameter values by calculating the response of each input to the transfer functions and then evaluating the noise n_t as the difference between this response (combined for all the inputs) and the output. The innovations are derived from the noise using the ARIMA model in the same manner as for a univariate series, and as described in Section 2.1.5.

In estimating the parameters, consideration has to be given to the lagged terms in the various model equations which are associated with times prior to the observation period, and are therefore unknown. The function descriptions provide the necessary detail as to how this problem is treated.

Also, as described in Section 2.1.6 the sum of squares criterion

$$S = \sum a_t^2$$

is related to the quadratic form in the exact log-likelihood of the parameters:

$$-\tfrac{1}{2}\log|V| - \tfrac{1}{2}w^T V^{-1} w.$$

Here w is the vector of appropriately differenced noise terms, and

$$w^T V^{-1} w = S/\sigma_a^2,$$

where σ_a^2 is the innovation variance parameter.

The least-squares criterion is therefore identical to minimization of the quadratic form, but is not identical to exact likelihood. Because V may be expressed as $M\sigma_a^2$, where M is a function of the ARIMA model parameters, substitution of σ_a^2 by its maximum likelihood estimator yields a concentrated (or profile) likelihood which is a function of

$$|M|^{1/N} S.$$

N is the length of the differenced noise series w, and $|M| = \det M$.

Use of the above quantity, called the deviance, D, as an objective function is preferable to the use of S alone, on the grounds that it is equivalent to exact likelihood, and yields estimates with better properties. However, there is an appreciable computational penalty in calculating D, and in large samples it differs very little from S, except in the important case of seasonal ARIMA models where the number of whole seasons within the data length must also be large.

The user is given the option of taking the objective function to be either S or D, or a third possibility, the marginal likelihood. This is similar to exact likelihood but can counteract bias in the ARIMA model due to the fitting of a large number of simple inputs.

Approximate standard errors of the parameter estimates and the correlations between them are available after estimation.

The model residuals $\hat{a}_t$ are the innovations resulting from the estimation, and they are usually examined for the presence of either autocorrelation or cross-correlation with the inputs. Absence of such correlation provides some confirmation of the adequacy of the model.

2.3.5 Multi-input model forecasting

A multi-input model may be used to forecast the output series provided future values (possibly forecasts) of the input series are supplied.

Construction of the forecasts requires knowledge only of the model orders and parameters, together with a limited set of the most recent variables which appear in the model equations. This is called the state set. It is conveniently constituted after model estimation. Moreover, if new observations $y_{n+1}, y_{n+2}, \ldots$ of the output series and $x_{n+1}, x_{n+2}, \ldots$ of (all) the independent input series become available, then the model equations can easily be used to update the state set before constructing forecasts from the end of the new observations. The new innovations $a_{n+1}, a_{n+2}, \ldots$ generated in this updating may be used to monitor the continuing adequacy of the model.

2.3.6 Transfer function model filtering

In many time series applications it is desired to calculate the response (or output) of a transfer function model for a given input series.

Smoothing, detrending, and seasonal adjustment are typical applications. The user must specify the orders and parameters of a transfer function model for the purpose being considered. This may then be applied to the input series.

Again, problems may arise due to ignorance of the input series values prior to the observation period. The transient errors which can arise from this cause may be substantially reduced by using 'backforecasts' of these unknown observations.

2.4 Multivariate Time Series

Multi-input modelling represents one time series as dependent on several input time series. In some circumstances it may be more appropriate to consider the relationships between all the series simultaneously. The basic tools used in examining multiple time series are cross-correlation matrices, which are the multivariate extension of the autocorrelation function, and three multivariate versions of the partial autocorrelation function: multiple squared partial autocorrelations, partial lagcorrelation matrices and partical autoregression matrices.

It is assumed that the time series have been differenced if necessary, and that they are jointly stationary. The lagged correlations between all possible pairs of series, i.e.,

$$\rho_{ijl} = \text{corr}(x_{i,t}, x_{j,t+l})$$

are then taken to provide an adequate description of the statistical relationships between the series. These quantities are estimated by sample auto- and cross-correlations r_{ijl}. For each l these may be viewed as elements of a (lagged) autocorrelation matrix.

Thus consider the **vector process** x_t (with elements x_{it}) and lagged autocovariance matrices Γ_l with elements of $\sigma_i \sigma_j \rho_{ijl}$ where $\sigma_i^2 = \text{var}(x_{i,t})$. Correspondingly, Γ_l is estimated by the matrix C_l with elements $s_i s_j r_{ijl}$ where s_i^2 is the sample variance of x_{it}.

The correlation matrices provide a description of the joint statistical properties of the series. It is also possible to calculate matrix quantities which are closely analogous to the partial autocorrelations of univariate series (see Section 2.1.3). Wei (1990) discusses both the partial autoregression matrices proposed by Tiao and Box (1981) and partial lag correlation matrices.

In the univariate case the partial autocorrelation function (pacf) between x_t and x_{t+l} is the correlation coefficient between the two after removing the linear dependence on each of the intervening variables $x_{t+1}, x_{t+2}, \ldots, x_{t+l-1}$. This partial autocorrelation may also be obtained as the last regression coefficient associated with x_t when regressing x_{t+l} on its l lagged variables $x_{t+l-1}, x_{t+l-2}, \ldots, x_t$. Tiao and Box

(1981) extended this method to the multivariate case to define the partial autoregression matrix. Heyse and Wei (1985) also extended the univariate definition of the pacf to derive the correlation matrix between the vectors x_t and x_{t+l} after removing the linear dependence on each of the intervening vectors $x_{t+1}, x_{t+2}, \ldots, x_{t+l-1}$, the partial lag correlation matrix.

Note that the partial lag correlation matrix is a correlation coefficient matrix since each of its elements is a properly normalised correlation coefficient. This is not true of the partial autoregression matrices (except in the univariate case for which the two types of matrix are the same). The partial lag correlation matrix at lag 1 also reduces to the regular correlation matrix at lag 1; this is not true of the partial autoregression matrices (again except in the univariate case).

The partial autoregression matrices may be found by solving a multivariate version of the Yule–Walker equations to find the autoregression matrices, using the final regression matrix coefficient as the partial autoregression matrix at that particular lag.

The basis of these calculations is a multivariate autoregressive model:

$$x_t = \phi_{l,1} x_{t-1} + \cdots + \phi_{l,l} x_{t-l} + e_{l,t}$$

where $\phi_{l,1}, \phi_{l,2}, \ldots, \phi_{l,l}$ are **matrix coefficients**, and $e_{l,t}$ is the vector of errors in the prediction. These coefficients may be rapidly computed using a recursive technique which requires, and simultaneously furnishes, a backward prediction equation:

$$x_{t-l-1} = \psi_{l,1} x_{t-l} + \psi_{l,2} x_{t-l+1} + \cdots + \psi_{l,l} x_{t-1} + f_{l,t}$$

(in the univariate case $\psi_{l,i} = \phi_{l,i}$).

The forward prediction equation coefficients, $\phi_{l,i}$, are of direct interest, together with the covariance matrix D_l of the prediction errors $e_{l,t}$. The calculation of these quantities for a particular maximum equation lag $l = L$ involves calculation of the same quantities for increasing values of $l = 1, 2, \ldots, L$.

The quantities $v_l = \det D_l / \det \Gamma_0$ may be viewed as generalized variance ratios, and provide a measure of the efficiency of prediction (the smaller the better). The reduction from v_{l-1} to v_l which occurs on extending the order of the predictor to l may be represented as

$$v_l = v_{l-1}(1 - \rho_l^2)$$

where ρ_l^2 is a multiple squared partial autocorrelation coefficient associated with k^2 degrees of freedom.

Sample estimates of all the above quantities may be derived by using the series covariance matrices C_l, for $l = 1, 2, \ldots, L$, in place of Γ_l. The best lag for prediction purposes may be chosen as that which yields the minimum final prediction error (FPE) criterion:

$$FPE(l) = v_l \times \frac{(1 + lk^2/n)}{(1 - lk^2/n)}.$$

An alternative method of estimating the sample partial autoregression matrices is by using multivariate least-squares to fit a series of multivariate autoregressive models of increasing order.

2.5 Cross-spectral Analysis

The relationship between two time series may be investigated in terms of their sinusoidal components at different frequencies. At frequency ω a component of y_t of the form

$$R_y(\omega) \cos(\omega t - \phi_y(\omega))$$

has its amplitude $R_y(\omega)$ and phase lag $\phi_y(\omega)$ estimated by

$$R_y(\omega) e^{i\phi_y(\omega)} = \tfrac{1}{n} \sum_{t=1}^{n} y_t e^{i\omega t}$$

and similarly for x_t. In the univariate analysis only the amplitude was important – in the cross analysis the phase is important.

2.5.1 The sample cross-spectrum

This is defined by

$$f^*_{xy}(\omega) = \frac{1}{2\pi n}\left(\sum_{t=1}^{n} y_t e^{i\omega t}\right)\left(\sum_{t=1}^{n} x_t e^{-i\omega t}\right).$$

It may be demonstrated that this is equivalently defined in terms of the sample CCF, $r_{xy}(k)$, of the series as

$$f^*_{xy}(\omega) = \frac{1}{2\pi}\sum_{-(n-1)}^{(n-1)} c_{xy}(k)e^{i\omega k}$$

where $c_{xy}(k) = s_x s_y r_{xy}(k)$ is the cross-covariance function.

2.5.2 The amplitude and phase spectrum

The cross-spectrum is specified by its real part or cospectrum $cf^*(\omega)$ and imaginary part or quadrature spectrum $qf^*(\omega)$, but for the purpose of interpretation the cross-amplitude spectrum and phase spectrum are useful:

$$A^*(\omega) = |f^*_{xy}(\omega)|,\ \phi^*(\omega) = \arg(f^*_{xy}(\omega)).$$

If the series x_t and y_t contain deterministic sinusoidal components of amplitudes R_y, R_x and phases ϕ_y, ϕ_x at frequency ω, then $A^*(\omega)$ will have a peak of approximate width π/n and height $(n/2\pi)R_y R_x$ at that frequency, with corresponding phase $\phi^*(\omega) = \phi_y - \phi_x$. This supplies no information that cannot be obtained from the two series separately. The statistical relationship between the series is better revealed when the series are purely stochastic and jointly stationary, in which case the expected value of $f^*_{xy}(\omega)$ converges with increasing sample size to the theoretical cross-spectrum

$$f_{xy}(\omega) = \frac{1}{2\pi}\sum_{-\infty}^{\infty} \gamma_{xy}(k)e^{i\omega k}$$

where $\gamma_{xy}(k) = \mathrm{cov}(x_t, y_{t+k})$. The sample spectrum, as in the univariate case, does not, however, converge to the theoretical spectrum without some form of smoothing which either implicitly (using a lag window) or explicitly (using a frequency window) averages the sample spectrum $f^*_{xy(\omega)}$ over wider bands of frequency to obtain a smoothed estimate $\hat{f}_{xy}(\omega)$.

2.5.3 The coherency spectrum

If there is no statistical relationship between the series at a given frequency, then $f_{xy}(\omega) = 0$, and the smoothed estimate $\hat{f}_{xy}(\omega)$, will be close to 0. This is assessed by the squared coherency between the series:

$$\hat{W}(\omega) = \frac{|\hat{f}_{xy}(\omega)|^2}{\hat{f}_{xx}(\omega)\hat{f}_{yy}(\omega)}$$

where $\hat{f}_{xx}(\omega)$ is the corresponding smoothed univariate spectrum estimate for x_t, and similarly for y_t. The coherency can be treated as a squared multiple correlation. It is similarly invariant in theory not only to simple scaling of x_t and y_t, but also to filtering of the two series, and provides a useful test statistic for the relationship between autocorrelated series. Note that without smoothing,

$$|f^*_{xy}(\omega)|^2 = f^*_{xx}(\omega)f^*_{yy}(\omega),$$

so the coherency is 1 at all frequencies, just as a correlation is 1 for a sample of size 1. Thus smoothing is essential for cross-spectrum analysis.

2.5.4 The gain and noise spectrum

If y_t is believed to be related to x_t by a linear lagged relationship as in Section 2.3, i.e.,

$$y_t = v_0 x_t + v_1 x_{t-1} + v_2 x_{t-2} + \cdots + n_t,$$

then the theoretical cross-spectrum is

$$f_{xy}(\omega) = V(\omega) f_{xx}(\omega)$$

where

$$V(\omega) = G(\omega) e^{i\phi(\omega)} = \sum_{k=0}^{\infty} v_k e^{ik\omega}$$

is called the frequency response of the relationship.

Thus if x_t were a sinusoidal wave at frequency ω (and n_t were absent), y_t would be similar but multiplied in amplitude by $G(\omega)$ and shifted in phase by $\phi(\omega)$. Furthermore, the theoretical univariate spectrum

$$f_{yy}(\omega) = G(\omega)^2 f_{xx}(\omega) + f_n(\omega)$$

where n_t, with spectrum $f_n(\omega)$, is assumed independent of the input x_t.

Cross-spectral analysis thus furnishes estimates of the gain

$$\hat{G}(\omega) = |\hat{f}_{xy}(\omega)| / \hat{f}_{xx}(\omega)$$

and the phase

$$\hat{\phi}(\omega) = \arg\Big(\hat{f}_{xy}(\omega)\Big).$$

From these representations of the estimated frequency response $\hat{V}(\omega)$, parametric TF models may be recognised and selected. The noise spectrum may also be estimated as

$$\hat{f}_{y|x}(\omega) = \hat{f}_{yy}(\omega)\big(1 - \hat{W}(\omega)\big)$$

– a formula which reflects the fact that in essence a regression is being performed of the sinusoidal components of y_t on those of x_t over each frequency band.

Interpretation of the frequency response may be aided by extracting from $\hat{V}(\omega)$ estimates of the IRF $\hat{v}_k$. It is assumed that there is no anticipatory response between y_t and x_t, i.e., no coefficients v_k with $k = -1, -2$ are needed (their presence might indicate feedback between the series).

2.5.5 Cross-spectrum smoothing by lag window

The estimate of the cross-spectrum is calculated from the sample cross-variances as

$$\hat{f}_{xy}(\omega) = \frac{1}{2\pi} \sum_{-M+S}^{M+S} w_{k-S} c_{xy}(k) e^{i\omega k}.$$

The lag window w_k extends up to a truncation lag M as in the univariate case, but its centre is shifted by an alignment lag S usually chosen to coincide with the peak cross-correlation. This is equivalent to an alignment of the series for peak cross-correlation at lag 0, and reduces bias in the phase estimation.

The selection of the truncation lag M, which fixes the bandwidth of the estimate, is based on the same criteria as for univariate series, and the same choice of M and window shape should be used as in univariate spectrum estimation to obtain valid estimates of the coherency, gain etc., and test statistics.

2.5.6 Direct smoothing of the cross-spectrum

The computations are exactly as for smoothing of the univariate spectrum except that allowance is made for an implicit alignment shift S between the series.

2.6 Kalman Filters

Kalman filtering provides a method for the analysis of multi-dimensional time series. The underlying model is:

$$X_{t+1} = A_t X_t + B_t W_t$$

$$Y_t = C_t X_t + V_t$$

where X_t is the unobserved state vector, Y_t is the observed measurement vector, W_i is the process noise, V_i is the measurement noise, A_i is the system state transition matrix, B_i is the system input weight matrix, C_i is the system output weight matrix and D_i is the control matrix (where the subscript i refers to the value of the appropriate quantity at time i).

The vectors X_i, Y_i, U_i and W_i are of dimension n, p, r and m, respectively. The matrices A_i, B_i, C_i and D_i are of dimension n by n, n by m, p by n and n by r, respectively.

Here the process noise and the measurement noise sequences are assumed to be uncorrelated and have zero mean. This implies that:

$$E\{W_i\} = 0 \quad E\{V_i\} = 0 \text{ and } E\{W_i V_i^T\} = 0,$$

and the covariance matrices are

$$E\{W_i W_i^T\} = Q_i \quad E\{V_i V_i^T\} = R_i$$

where the operator E denotes the 'expectation value', and Q_i and R_i are positive definite matrices. At instant i, Q_i is the process noise covariance matrix whilst R_i is the measurement noise covariance matrix. [Note that in the case of the information filter the first condition is relaxed, i.e., the mean of the process noise may be non-zero.]

If the system matrices A_i, B_i, C_i, D_i and also the covariance matrices Q_i, R_i are known then Kalman filtering can be used to compute the minimum variance estimate of the stochastic variable X_i estimated from the observed values Y_1 to Y_j

$$\hat{X}_{i|j} = \hat{X}_{i|Y_1 \cdots Y_j}. \tag{6}$$

When $j = i$ the above estimate is called the filtered estimate, and when $j = i - 1$ it is known as the one-step predicted estimate, or simply the predicted estimate.

Kalman filtering uses a recursive method which involves computing the state covariance matrices $P_{i|i}$ and/or $P_{i|i-1}$ and the estimates $\hat{X}_{i|i}$ and/or $\hat{X}_{i|i-1}$ from their previous values for $i = 1, 2, \ldots$.

If the covariance of the initial state X_0 (represented by $P_{0|-1}$) is known, and the mean of the initial state X (represented by $X_{0|-1}$) is given then the following recurrence relations provide the required estimates.

$$H_i = R_i + C_i P_{i|i-1} C_i^T \tag{7}$$

$$K_i = P_{i|i-1} C_i^T H_i^{-1} \tag{8}$$

$$P_{i|i} = [I - K_i C_i] P_{i|i-1} \tag{9}$$

$$\hat{X}_{i|i} = \hat{X}_{i|i-1} + K_i \nu_i \tag{10}$$

where the one step ahead prediction error is given by $\nu_i = Y_i - C_i \hat{X}_{i|i-1}$

$$P_{i+1|i} = A_i P_{i|i} A_i^T + B_i Q_i B_i^T \tag{11}$$

$$\hat{X}_{i+1|i} = A_i \hat{X}_{i|i} + D_i U_i \tag{12}$$

where K_i is referred to as the Kalman gain matrix and H_i contains the covariance matrix of the prediction errors ν_i. It can be seen that equations (7), (9) and (11) define the recursion involving the covariance matrices $P_{i|i-1}$, $P_{i|i}$ and $P_{i+1|i}$. These matrices are positive semidefinite and can therefore be factorised into their Cholesky ('square root') factors. Equations (9) and (10) yielding $\hat{X}_{i|i}$ and $P_{i|i}$ from $\hat{X}_{i|i-1}$ and $P_{i|i-1}$ are termed measurement-update equations, while equations (11) and (12) yielding $\hat{X}_{i+1|i}$ and $P_{i|i+1}$ from $\hat{X}_{i|i}$ and $P_{i|i}$ are known as time-update equations.

2.6.1 The information filter

An alternative set of Kalman filter equations can be constructed which use the inverse of the covariance matrices. These matrices (e.g., $P_{i+1|i}^{-1}$) are also positive semidefinite and are termed information matrices.

Although the information filter has the disadvantage that it requires the inverses A_i^{-1} and R_i^{-1} to be computed, it is preferable to the covariance filter in situations where there is no (very little) information concerning the initial state of the system. In these circumstances the covariance filter will fail because the initial state covariance matrix $P_{0|-1}$ is infinite (very large), whilst the corresponding information filter initial state $P_{0|-1}^{-1} = 0$ (very small) incurs no such difficulties.

The information filter recursion (with $D_i = 0$) is described by the following equations

$$P_{i+1|i}^{-1} = [I - N_i B_i^T] M_i \tag{13}$$

$$P_{i+1|i+1}^{-1} = P_{i+1|i}^{-1} + C_{i+1}^T R_{i+1}^{-1} C_{i+1} \tag{14}$$

$$\text{where } M_i = \left(A_i^{-1}\right)^T P_{i|i}^{-1} A_i^{-1}$$

$$\text{and } N_i = M_i B_i [Q_i^{-1} + B_i^T M_i B_i]^{-1}$$

$$\hat{a}_{i+1|i} = [I - N_i B_i^T]\left(A_i^{-1}\right)^T \hat{a}_{i|i} \tag{15}$$

$$\hat{a}_{i+1|i+1} = \hat{a}_{i+1|i} + C_{i+1}^T R_{i+1}^{-1} Y_{i+1} \tag{16}$$

$$\text{where } \hat{a}_{i+1|i} = P_{i+1|i}^{-1} \hat{X}_{i+1|i} \tag{17}$$

$$\text{and } \hat{a}_{i+1|i+1} = P_{i+1|i+1}^{-1} \hat{X}_{i+1|i+1}. \tag{18}$$

2.6.2 Square root filters

The use of the Kalman filter equations previously given can result in covariance matrices which are not positive semidefinite. This may happen if some of the measurements are very accurate and numerical computation then involves ill-conditioned quantities. Square root filtering is a technique which overcomes this difficulty by propagating the covariance matrices in Cholesky (square root) form. This has the advantage that, despite computational errors, the product of the Cholesky factors will always yield a positive definite covariance matrix. The numerical conditioning of the Cholesky square root is also generally much better than that of its corresponding covariance matrix. Since the condition number of the Cholesky factor is the square root of the condition number of the covariance matrix, square root filtering can avoid numerical difficulties with only half as many significant digits as the full matrix Kalman filters outlined above.

2.6.3 The square root covariance filter

The time-varying square root covariance Kalman filter (nag_kalman_sqrt_filt_cov_var (g13eac)) provided by this chapter requires the construction of the following block matrix pre-array and block matrix post-array.

$$\underset{\text{(Pre-array)}}{\begin{pmatrix} R_i^{1/2} & C_i S_i & 0 \\ 0 & A_i S_i & B_i Q_i^{1/2} \end{pmatrix}} \quad U_1 = \quad \underset{\text{(Post-array)}}{\begin{pmatrix} H_i^{1/2} & 0 & 0 \\ G_i & S_{i+1} & 0 \end{pmatrix}} \tag{19}$$

where U_1 is an orthogonal matrix that triangularizes the pre-array and the matrices $P_{i|i-1}$, Q_i, H_i and R_i have been Cholesky factorised as follows:

$$P_{i|i-1} = S_i S_i^T \quad Q_i = Q_i^{1/2}\left(Q_i^{1/2}\right)^T \quad R_i = R_i^{1/2}\left(R_i^{1/2}\right)^T \quad H_i = H_i^{1/2}\left(H_i^{1/2}\right)^T$$

where the left factors are lower triangular, and the Kalman filter gain matrix K_i is related to G_i by

$$A_iK_i = G_i\left(H^{1/2}\right)^{-1}.$$

The new state estimate can then be calculated using

$$\hat{X}_{i+1|i} = A_i\hat{X}_{i|i-1} + AK_i(Y_i - C_i\hat{X}_{i|i-1}) + D_iU_i \tag{20}$$

where the product of the matrices A_i and K_i is represented as AK_i.

That this method is computationally equivalent to equations (7)–(12) can be demonstrated by 'squaring' each side of equation (19) (post-multiplying each side by its transpose) and then equating block matrix elements on either side. It can similarly be shown that transposition of columns 2 and 3 of the pre-array, as occurs in function nag_kalman_sqrt_filt_cov_invar (g13ebc), does not affect the elements in the resultant post-array.

2.6.4 The square root information filter

The time-varying square root information Kalman filter (nag_kalman_sqrt_filt_info_var (g13ecc)) provided by this chapter requires the construction of the following block matrix pre-array and block matrix post-array.

$$U_2\begin{pmatrix} Q_i^{-1/2} & 0 & 0 \\ S_i^{-1}A_i^{-1}B_i & S_i^{-1}A_i^{-1} & S_i^{-1}\hat{X}_{i|i} \\ 0 & R_{i+1}^{-1/2}C_{i+1} & R_{i+1}^{-1/2}Y_{i+1} \end{pmatrix} = \begin{pmatrix} F_{i+1}^{-1/2} & * & * \\ 0 & S_{i+1}^{-1} & \xi_{i+1|i+1} \\ 0 & 0 & E_{i+1} \end{pmatrix} \tag{21}$$

(Pre-array) (Post-array)

where the asterisk represents elements that are not required, U_2 is an orthogonal transformation triangularizing the pre-array and F_{i+1}, the matrix containing the innovations in the process noise, is given by

$$F_{i+1}^{-1} = Q_i^{-1} + B_i^TM_iB_i.$$

The matrices $P_{i|i}^{-1}$, Q_i^{-1}, F_{i+1}^{-1} and R_i^{-1} have been Cholesky factorised as follows:

$$P_{i|i}^{-1} = \left(S_i^{-1}\right)^TS_i^{-1}$$

$$Q_i^{-1} = \left(Q_i^{-1/2}\right)^TQ_i^{-1/2}$$

$$R_i^{-1} = \left(R_i^{-1/2}\right)^TR_i^{-1/2}$$

$$F_{i+1}^{-1} = \left(F_{i+1}^{-1/2}\right)^TF_{i+1}^{-1/2}$$

where the right factors are upper triangular.

The new state estimate is computed via

$$X_{i+1|i+1} = S_{i+1}\xi_{i+1|i+1}. \tag{22}$$

That this method is computationally equivalent to equations (13)–(18) can be demonstrated by transposing (21), 'squaring' the right-hand side to eliminate the orthogonal matrix U_2 and then, after performing a block Cholesky decomposition, equating block matrix elements on either side. It can similarly be shown that transposition of rows 2 and 3 of the pre-array, as occurs in function nag_kalman_sqrt_filt_info_invar (g13edc), does not affect the elements in the resultant post-array.

2.6.5 Time invariant condensed square root filters

When the system matrices A, B, C are time invariant, it can be advantageous to perform initial unitary transformations to 'condense' them (create as many zeros as possible) and thereby significantly reduce the number of floating-point operations required by the algorithm. Essentially this entails creating an

appropriate unitary transformation matrix U and solving for the new state vector $X_t = UX$ in the transformed reference frame. After the required number of Kalman filter iterations have been performed the back transformation $X = U^T X_t$ provides the estimated state vector in the original reference frame. It can be shown from equations (7)–(12) that the transformed system matrices for the covariance filter are given by $\{UAU^T, UB, CU^T\}$, which are in agreement with the arguments required by nag_kalman_sqrt_filt_cov_invar (g13ebc). It can similarly be shown, from equations (13)–(18), that the system matrices describing the corresponding transformed information filter are $\{UA^{-1}U^T,\ UB,\ CU^T\}$. These correspond to the arguments used by nag_kalman_sqrt_filt_info_invar (g13edc) ($UA^{-1}U^T$, $UA^{-1}B$, CU^T), where the second matrix is input as the product of $UA^{-1}U^T$ and UB. It should be noted that in the transformed frame the covariance matrix $P'_{i|i}$ is related to the original covariance matrix via the similarity transformation $P'_{i|i} = UP_{i|i}U^T \left(\left(P'_{i|i}\right)^{-1} = U(P'^{-1}_{i|i})U^T \right)$. This means that, for square root Kalman filter functions, the appropriate Cholesky factor of $P'_{i|i}$ must be input.

The condensed matrix forms used by the functions in this chapter are (nag_kalman_sqrt_filt_cov_invar (g13ebc)) the lower observer Hessenberg form where the compound matrix

$$\left(\begin{array}{c} UAU^T \\ \hline CU^T \end{array}\right)$$

is lower trapezoidal and (nag_kalman_sqrt_filt_info_invar (g13edc)) the upper controller Hessenberg form where the compound matrix $(UB|UAU^T)$ is upper trapezoidal.

Both nag_kalman_sqrt_filt_cov_invar (g13ebc) and nag_kalman_sqrt_filt_info_invar (g13edc) contain the block matrix

$$\left(\begin{array}{c|c} & CU^T \\ \hline UB & UAU^T \end{array}\right)$$

within their pre-array, and the structure of this matrix (for $n = 6$, $m = 3$ and $p = 2$) is illustrated below for both Hessenberg forms

Lower observer Hessenberg

$$\left(\begin{array}{ccc|cccccc} & & & x & 0 & 0 & 0 & 0 & 0 \\ & & & x & x & 0 & 0 & 0 & 0 \\ x & x & x & x & x & x & 0 & 0 & 0 \\ x & x & x & x & x & x & x & 0 & 0 \\ x & x & x & x & x & x & x & x & 0 \\ x & x & x & x & x & x & x & x & x \\ x & x & x & x & x & x & x & x & x \\ x & x & x & x & x & x & x & x & x \end{array}\right).$$

Upper controller Hessenberg

$$\left(\begin{array}{cc|cccccc} & & x & x & x & x & x & x \\ & & x & x & x & x & x & x \\ & & x & x & x & x & x & x \\ \hline x & x & x & x & x & x & x & x \\ 0 & x & x & x & x & x & x & x \\ 0 & 0 & x & x & x & x & x & x \\ 0 & 0 & 0 & x & x & x & x & x \\ 0 & 0 & 0 & 0 & x & x & x & x \\ 0 & 0 & 0 & 0 & 0 & x & x & x \end{array}\right).$$

2.6.6 Model fitting and forecasting

If the state space model contains unknown parameters, θ, these can be estimated using maximum likelihood. Assuming that W_t and V_t are normal variates the log-likelihood for observations $Y_t, t = 1, 2, \ldots, n$ is given by

$$\text{constant} - \frac{1}{2}\sum_{t=1}^{n}\ln(\det(H_t)) - \frac{1}{2}\sum_{t=1}^{t} r_t^T H_t^{-1} r_t.$$

Optimal estimates for the unknown model parameters θ can then be obtained by using a suitable optimizer function to maximize the likelihood function.

Once the model has been fitted forecasting can be performed by using the one-step-ahead prediction equations. The one-step-ahead prediction equations can also be used to 'jump over' any missing values in the series.

2.6.7 Kalman filter and time series models

Many commonly used time series models can be written as state space models. A univariate ARMA(p, q) model can be cast into the following state space form:

$$\begin{aligned} x_t &= Ax_{t-1} + B\epsilon_t \\ w_t &= Cx_t \end{aligned}$$

where $r = \max(p, q+1)$, the first element of the state vector x_t is w_t,

$$A = \begin{pmatrix} \phi_1 & 1 & & & & \\ \phi_2 & & 1 & & & \\ \cdot & & & \cdot & & \\ \cdot & & & & \cdot & \\ \phi_{r-1} & & & & & 1 \\ \phi_r & 0 & 0 & . & . & 0 \end{pmatrix}, \quad B = \begin{pmatrix} 1 \\ -\theta_1 \\ -\theta_2 \\ \cdot \\ \cdot \\ -\theta_{r-1} \end{pmatrix} \quad \text{and} \quad C^T = \begin{pmatrix} 1 \\ 0 \\ 0 \\ \cdot \\ \cdot \\ 0 \end{pmatrix}.$$

The representation for a k-variate ARMA(p, q) series (VARMA) is very similar to that given above, except now the state vector is of length kr and the ϕ and θ are now $k \times k$ matrices and the 1s in A, B and C are now the identity matrix of order k. If $p < r$ or $q + 1 < r$ then the appropriate ϕ or θ matrices are set to zero, respectively.

Since the compound matrix

$$\begin{pmatrix} C \\ A \end{pmatrix}$$

is already in lower observer Hessenberg form (i.e., it is lower trapezoidal with zeros in the top right-hand triangle) the invariant Kalman filter algorithm can be used directly without the need to generate a transformation matrix U^*.

2.7 GARCH Models

2.7.1 ARCH models and their generalisations

Rather than modelling the mean (for example using regression models) or the autocorrelation (by using ARMA models) there are circumstances in which the variance of a time series needs to be modelled. This is common in financial data modelling where the variance (or standard deviation) is known as volatility. The ability to forecast volatility is a vital part in deciding the risk attached to financial decisions like portfolio selection. The basic model for relating the variance at time t to the variance at previous times is the autoregressive conditional heteroskedastic (ARCH) model. The standard ARCH model is defined as

$$y_t|\psi_{t-1} \sim N(0, h_t),$$

$$h_t = \alpha_0 + \sum_{i=1}^{q} \alpha_i \epsilon_{t-1}^2,$$

where ψ_t is the information up to time t and h_t is the conditional variance.

In a similar way to that in which autoregressive models were generalised to ARMA models the ARCH models have been generalised to a GARCH model; see Engle (1982), Bollerslev (1986) and Hamilton (1994).

$$h_t = \alpha_0 + \sum_{i=1}^{q} \alpha_i \epsilon_{t-1}^2 + \sum_{j=1}^{p} \beta h_{t-j}.$$

This can be combined with a regression model:

$$y_t = b_0 + \sum_{i=1}^{k} b_i x_{it} + \epsilon_t,$$

where $\epsilon_t | \psi_{t-1} \sim N(0, h_t)$ and where x_{it}, for $i = 1, \ldots, k$ are the exogenous variables.

The above models assume that the change in variance, h_t, is symmetric with respect to the shocks, that is, that a large negative value of ϵ_{t-1} has the same effect as a large positive value of ϵ_{t-1}. A frequently observed effect is that a large negative value ϵ_{t-1} often leads to a greater variance than a large positive value. The following three asymmetric models represent this effect in different ways using the parameter λ as a measure of the asymmetry.

Type I AGARCH(p, q)

$$h_t = \alpha_0 + \sum_{i=1}^{q} \alpha_i (\epsilon_{t-i} + \gamma)^2 + \sum_{j=1}^{p} \beta_i h_{t-j}.$$

Type II AGARCH(p, q)

$$h_t = \alpha_0 + \sum_{i=1}^{q} \alpha_i (|\epsilon_{t-i}| + \gamma \epsilon_{t-i})^2 + \sum_{j=1}^{p} \beta_i h_{t-j}.$$

GJR-GARCH(p, q), or Glosten, Jagannathan and Runkle GARCH (Glosten *et al.* (1993))

$$h_t = \alpha_0 + \sum_{i=1}^{q} (\alpha_i + \gamma S_{t-1}) \epsilon_{t-1}^2 + \sum_{j=1}^{p} \beta_i h_{t-j},$$

where $S_t = 1$ if $\epsilon_t < 0$ and $S_t = 0$ if $\epsilon_t \geq 0$.

The first assumes that the effects of the shocks are symmetric about γ rather than zero, so that for $\gamma < 0$ the effect of negative shocks is increased and the effect of positive shocks is decreased. Both the Type II AGARCH and the GJR GARCH (Glosten *et al.* (1993)) models introduce asymmetry by increasing the value of the coefficient of ϵ_{t-1}^2 for negative values of ϵ_{t-1}. In the case of the Type II AGARCH the effect is multiplicative while for the GJR GARCH the effect is additive.

Coefficient	$\epsilon_{t-1} < 0$	$\epsilon_{t-1} > 0$
Type II AGARCH	$\alpha_i (1 - \gamma)^2$	$\alpha_i (1 + \gamma)^2$
GJR GARCH	$\alpha_i + \gamma$	α_i

(Note that in the case of GJR GARCH, γ needs to be positive to inflate variance after negative shocks while for Type I and Type II AGARCH, γ needs to be negative.)

$$\ln(h_t) = \alpha_0 + \sum_{i=1}^{q} \alpha_i z_{t-i} + \sum_{i=1}^{q} \phi_i (|z_{t-i}| - E[|z_{t-i}|]) + \sum_{j=1}^{p} \beta_i \ln(h_{t-j}),$$

where $z_t = \dfrac{\epsilon_t}{\sqrt{h_t}}$ and $E[|z_{t-i}|]$ denotes the expected value of $|z_{t-i}|$.

Note that the ϕ_i terms represent a symmetric contribution to the variance while the α_i terms give an asymmetric contribution.

2.7.2 Fitting GARCH models

The models are fitted by maximising the conditional log-likelihood. For the Normal distribution the conditional log-likelihood is

$$\frac{1}{2}\sum_{i=1}^{T}\left(\log(h_i)+\frac{\epsilon_i^2}{h_i}\right).$$

For the Student's t-distribution the function is more complex. An approximation to the standard errors of the parameter estimates is computed from the Fisher information matrix.

3 Recommendations on Choice and Use of Available Functions

3.1 ARMA-type Models

ARMA-type modelling usually follows the methodology made popular by Box and Jenkins. It consists of four steps: identification, model fitting, model checking and forecasting. The availability of functions for each of these four steps is given below for the three types of modelling situation considered: univariate, input-output and multivariate.

3.1.1 Univariate series

(a) Model identification

The function nag_tsa_mean_range (g13auc) may be used in obtaining either a range–mean or standard deviation–mean plot for a series of observations, which may be useful in detecting the need for a variance-stabilising transformation. nag_tsa_mean_range (g13auc) computes the range or standard deviation and the mean for successive groups of observations that may then be used to produce a scatter plot of range against mean or of standard deviation against mean.

The function nag_tsa_diff (g13aac) may be used to difference a time series. The $N = n - d - s\times D$ values of the differenced time series which extends for $t = 1 + d + s\times D, \ldots, n$ are stored in the first N elements of the output array.

The function nag_tsa_auto_corr (g13abc) may be used for direct computation of the autocorrelations. It requires the time series as input, after optional differencing by nag_tsa_diff (g13aac).

An alternative is to use nag_tsa_spectrum_univar_cov (g13cac), which uses the FFT to carry out the convolution for computing the autocovariances. Circumstances in which this is recommended are

(i) if the main aim is to calculate the smoothed sample spectrum,

(ii) if the series length and maximum lag for the autocorrelations are both very large, in which case appreciable computing time may be saved.

For more precise recommendations, see Gentleman and Sande (1966). In this case the autocorrelations r_k need to be obtained from the autocovariances c_k by $r_k = c_k/c_0$.

The function nag_tsa_auto_corr_part (g13acc) computes the partial autocorrelation function and prediction error variance estimates from an input autocorrelation function. Note that nag_tsa_multi_part_lag_corr (g13dnc), which is designed for multivariate time series, may also be used to compute the partial autocorrelation function together with χ^2 statistics and their significance levels.

Finite lag predictor coefficients are also computed by the function nag_tsa_auto_corr_part (g13acc). It may have to be used twice, firstly with a large value for the maximum lag L in order to locate the optimum FPE lag, then again with L reset to this lag.

The function nag_tsa_arma_roots (g13dxc) may be used to check that the autoregressive part of the model is stationary and that the moving-average part is invertible.

(b) Model estimation

ARIMA models may be fitted using the function nag_tsa_multi_inp_model_estim (g13bec). This function can fit both simple ARIMA models as well as more complex multi-input models. There is a choice of using least-squares or maximum likelihood estimation.

(c) Model checking

The function nag_tsa_resid_corr (g13asc) calculates the correlations in the residuals from a model fitted by nag_tsa_multi_inp_model_estim (g13bec). In addition the standard errors and correlations of the residual autocorrelations are computed along with a portmanteau test for model adequacy.

(d) Forecasting using an ARIMA model

The function nag_tsa_multi_inp_model_forecast (g13bjc) can be used to compute forecasts using a specified ARIMA model using the observed values of the series.

3.1.2 Multi-input/transfer function modelling

(a) Model identification

Normally use nag_tsa_cross_corr (g13bcc) for direct computation of cross-correlations, from which cross-covariances may be obtained by multiplying by $s_y s_x$, and impulse response estimates (after prewhitening) by multiplying by s_y/s_x, where s_y, s_x are the sample standard deviations of the series.

An alternative is to use nag_tsa_spectrum_bivar_cov (g13ccc), which exploits the FFT to carry out the convolution for computing cross-covariances. The criteria for this are the same as given in Section 3.1.1 for calculation of autocorrelations. The impulse response function may also be computed by spectral methods without prewhitening using nag_tsa_noise_spectrum_bivar (g13cgc).

nag_tsa_arma_filter (g13bac) may be used to prewhiten or filter a series by an ARIMA model.

nag_tsa_transf_filter (g13bbc) may be used to filter a time series using a transfer function model.

(b) Estimation of multi-input model parameters

The function nag_tsa_transf_prelim_fit (g13bdc) is used to obtain preliminary estimates of transfer function model parameters. The model orders and an estimate of the impulse response function (see Section 3.2.1) are required.

The simultaneous estimation of the transfer function model parameters for the inputs, and ARIMA model parameters for the output, is carried out by nag_tsa_multi_inp_model_estim (g13bec).

This function requires values of the output and input series, and the orders of all the models. Any differencing implied by the model is carried out internally.

The function also requires the maximum number of iterations to be specified, and returns the state set for use in forecasting.

(c) Multi-input model checking

The function nag_tsa_resid_corr (g13asc), primarily designed for univariate time series, can be used to test the residuals from an input-output model.

(d) Forecasting using a multi-input model

The function nag_tsa_multi_inp_model_forecast (g13bjc) can be used to compute forecasts for a speicified multi-input model using the observed values of the series. Forecast for the input series have to be provided.

(e) Filtering a time series using a transfer function model

The function for this purpose is nag_tsa_transf_filter (g13bbc).

3.1.3 Multivariate series

The function nag_tsa_multi_diff (g13dlc) may be used to difference the series. The user must supply the differencing parameters for each component of the multivariate series. The order of differencing for each individual component does not have to be the same. The function may also be used to apply a log or square root transformation to the components of the series.

The function nag_tsa_multi_cross_corr (g13dmc) may be used to calculate the sample cross-correlation or cross-covariance matrices. It requires a set of time series as input. The user may request either the cross-covariances or cross-correlations.

The function nag_tsa_multi_part_lag_corr (g13dnc) computes the partial lag correlation matrices from the sample cross-correlation matrices computed by nag_tsa_multi_cross_corr (g13dmc), and the function nag_tsa_multi_part_regsn (g13dpc) computes the least-squares estimates of the partial autoregression matrices and their standard errors. Both functions compute a series of χ^2 statistic that aid the determination of the order of a suitable autoregressive model. nag_tsa_multi_auto_corr_part (g13dbc) may also be used in the identification of the order of an autoregressive model. The function computes multiple squared partial autocorrelations and predictive error variance ratios from the sample cross-correlations or cross-covariances computed by nag_tsa_multi_cross_corr (g13dmc).

The function nag_tsa_arma_roots (g13dxc) may be used to check that the autoregressive part of the model is stationary and that the moving-average part is invertible.

3.2 Spectral Methods

3.2.1 Univariate spectral estimation

Two functions are available, nag_tsa_spectrum_univar_cov (g13cac) carrying out smoothing using a lag window and nag_tsa_spectrum_univar (g13cbc) carrying out direct frequency domain smoothing. Both can take as input the original series, but nag_tsa_spectrum_univar_cov (g13cac) alone can use the sample autocovariances as alternative input. This has some computational advantage if a variety of spectral estimates needs to be examined for the same series using different amounts of smoothing.

However, the real choice in most cases will be which of the four shapes of lag window in nag_tsa_spectrum_univar_cov (g13cac) to use, or whether to use the trapezium frequency window of nag_tsa_spectrum_univar (g13cbc). The references may be consulted for advice on this, but the two most recommended lag windows are the Tukey and Parzen. The Tukey window has a very small risk of supplying negative spectrum estimates; otherwise, for the same bandwidth, both give very similar results, though the Parzen window requires a higher truncation lag (more acf values).

The frequency window smoothing procedure of nag_tsa_spectrum_univar (g13cbc) with a trapezium shape parameter $p \simeq \frac{1}{2}$ generally gives similar results for the same bandwidth as lag window methods with a slight advantage of somewhat less distortion around sharp peaks, but suffering a rather less smooth appearance in fine detail.

3.2.2 Cross-spectrum estimation

Two functions are available for the main step in cross-spectral analysis. To compute the cospectrum and quadrature spectrum estimates using smoothing by a lag window, nag_tsa_spectrum_bivar_cov (g13ccc) should be used. It takes as input either the original series or cross-covariances which may be computed in a previous call of the same function or possibly using results from nag_tsa_cross_corr (g13bcc). As in the univariate case, this gives some advantage if estimates for the same series are to be computed with different amounts of smoothing.

The choice of window shape will be determined as the same as that which has already been used in univariate spectrum estimation for the series.

For direct frequency domain smoothing, nag_tsa_spectrum_bivar (g13cdc) should be used, with similar consideration for the univariate estimation in choice of degree of smoothing.

The cross-amplitude and squared coherency spectrum estimates are calculated, together with upper and lower confidence bounds, using nag_tsa_cross_spectrum_bivar (g13cec). For input the cross-spectral estimates from either nag_tsa_spectrum_bivar_cov (g13ccc) or nag_tsa_spectrum_bivar (g13cdc) and corresponding univariate spectra from either nag_tsa_spectrum_univar_cov (g13cac) or nag_tsa_spectrum_univar (g13cbc) are required.

The gain and phase spectrum estimates are calculated together with upper and lower confidence bounds using nag_tsa_gain_phase_bivar (g13cfc). The required input is as for nag_tsa_cross_spectrum_bivar (g13cec) above.

The noise spectrum estimates and impulse response function estimates are calculated together with multiplying factors for confidence limits on the former, and the standard error for the latter, using nag_tsa_noise_spectrum_bivar (g13cgc). The required input is again the same as for nag_tsa_cross_spectrum_bivar (g13cec) above.

3.3 Kalman Filtering

There are four main functions available for Kalman filtering covering both the covariance and information filters with time-varying ro time-invariant filter. For covariance filters the functions are nag_kalman_sqrt_filt_cov_var (g13eac) for time-varying filter and nag_kalman_sqrt_filt_cov_invar (g13ebc) for time-invariant filter while the equivalent for the information filter they are nag_tsa_cross_spectrum_bivar (g13cec) and nag_kalman_sqrt_filt_info_invar (g13edc) respectively. In addition, for use with the time invariant-filters, the function nag_trans_hessenberg_observer (g13ewc) provides the required transformation to lower or upper Hessenberg form while nag_trans_hessenberg_controller (g13exc) provides the transformation to lower or upper controller Hessenberg form.

3.4 GARCH Models

The main choice in selecting a type of GARCH model is whether the data is symmetric or asymmetric and if asymmetric what form of asymmetry should be included in the model.

A symmetric ARCH or GARCH model can be fitted by nag_estimate_agarchI (g13fac) and the volatility forecast by nag_forecast_agarchI (g13fbc). For asymmetric data the choice is between the type of asymmetry as described in Section 2.7.

GARCH Type	Fit	Forecast
Type I	nag_estimate_agarchI (g13fac)	nag_forecast_agarchI (g13fbc)
Type II	nag_estimate_agarchII (g13fcc)	nag_forecast_agarchII (g13fdc)
GJR	nag_estimate_garchGJR (g13fec)	nag_forecast_garchGJR (g13ffc)

All functions allow the option of including regressor variables in the model.

3.5 Time Series Simulation

There are functions available in Chapter g05 for generating a realisation of a time series from a specified model: nag_rngs_arma_time_series (g05pac) for univariate time series and nag_rngs_varma_time_series (g05pcc) for multivariate time series. There is also a suite of functions for simulating GARCH models: nag_generate_agarchI (g05hkc), nag_generate_agarchII (g05hlc) and nag_generate_garchGJR (g05hmc).

3.6 Summary of Recommendations

ARMA modelling
- ACF ..... nag_tsa_auto_corr (g13abc)
- Diagnostic checking ..... nag_tsa_resid_corr (g13asc)
- Differencing ..... nag_tsa_diff (g13aac)
- Fitting ..... nag_tsa_multi_inp_model_estim (g13bec)
- Mean/range ..... nag_tsa_mean_range (g13auc)
- PACF ..... nag_tsa_auto_corr_part (g13acc)

Bivariate spectral analysis
- Barltett, Tukey, Parzen windows ..... nag_tsa_spectrum_bivar_cov (g13ccc)
- Direct smoothing ..... nag_tsa_spectrum_bivar (g13cdc)
- Other representations ..... nag_tsa_cross_spectrum_bivar (g13cec)
- Other representations ..... nag_tsa_gain_phase_bivar (g13cfc)
- Other representations ..... nag_tsa_noise_spectrum_bivar (g13cgc)

GARCH
- Asymmetric ARCH/GARCH
  - Fitting ..... nag_estimate_agarchI (g13fac)
  - Fitting ..... nag_estimate_agarchII (g13fcc)
  - Fitting ..... nag_estimate_garchGJR (g13fec)
  - Forecasting ..... nag_forecast_agarchI (g13fbc)
  - Forecasting ..... nag_forecast_agarchII (g13fdc)

4 Functions Withdrawn or Scheduled for Withdrawal

None.

5 References

Akaike H (1971) Autoregressive model fitting for control *Ann. Inst. Statist. Math.* **23** 163–180

Bollerslev T (1986) Generalised autoregressive conditional heteroskedasticity *Journal of Econometrics* **31** 307–327

Box G E P and Jenkins G M (1976) *Time Series Analysis: Forecasting and Control* (Revised Edition) Holden–Day

Engle R (1982) Autoregressive conditional heteroskedasticity with estimates of the variance of United Kingdom inflation *Econometrica* **50** 987–1008

Gentleman W S and Sande G (1966) Fast Fourier transforms for fun and profit *Proc. Joint Computer Conference, AFIPS* **29** 563–578

Glosten L, Jagannathan R and Runkle D (1993) Relationship between the expected value and the volatility of nominal excess return on stocks *Journal of Finance* **48** 1779–1801

Hamilton J (1994) *Time Series Analysis* Princeton University Press

Heyse J F and Wei W W S (1985) The partial lag autocorrelation function *Technical Report No. 32* Department of Statistics, Temple University, Philadelphia

Tiao G C and Box G E P (1981) Modelling multiple time series with applications *J. Am. Stat. Assoc.* **76** 802–816

Wei W W S (1990) *Time Series Analysis: Univariate and Multivariate Methods* Addison–Wesley