

Characteristics of the ARMA models

MS-C2128 Prediction and Time Series Analysis

Fall term 2020

Week 5: Predicting using ARMA models, Kalman filter and Dynamic regression

- 1 Prediction
 - 1 Predicting using ARMA models
 - 2 Exponential smoothing
- 2 Kalman filter
- 3 Dynamic regression

- 1 Prediction
- 2 Kalman filter
- 3 Dynamic regression

Predicting using ARMA models: Idea

- Consider a time series $x_t, t = \dots, -2, -1, 0, 1, 2, \dots$. We wish to predict the value x_{t+s} based on the *observed* values up to the time point t .
- The prediction $\hat{x}_{t+s|t}$ is thus some function of the values $\dots, x_{t-2}, x_{t-1}, x_t$.
- The goal is to find a function of $\dots, x_{t-2}, x_{t-1}, x_t$ such that the obtained value is as close as possible to the true value x_{t+s} .
- **An optimal prediction** (in the sense of the mean square error) is obtained from the conditional expected value

$$\hat{x}_{t+s|t} := E [x_{t+s} \mid x_t, x_{t-1}, \dots]$$

Predicting using ARMA models: Idea

Assume that the process $(x_t)_{t \in T}$ has an MA representation

$$x_t = \sum_{j=0}^{\infty} \psi_j L^j \epsilon_t, \quad \psi_0 = 1, \quad \sum_{j=0}^{\infty} |\psi_j| < \infty,$$

and assume that the process $(\epsilon_t)_{t \in T}$ has been observed until the time point t (that is, we have the observations $\epsilon_t, \epsilon_{t-1}, \epsilon_{t-2}, \dots$).

Predicting using ARMA models: Idea

- Now

$$x_{t+s} = \epsilon_{t+s} + \psi_1 \epsilon_{t+s-1} + \dots + \psi_{s-1} \epsilon_{t+1} + \psi_s \epsilon_t + \psi_{s+1} \epsilon_{t-1} + \dots$$

and the **optimal prediction** at the time point t (in the sense of the mean square error) is

$$\hat{x}_{t+s|t} := E [x_{t+s} \mid \epsilon_t, \epsilon_{t-1}, \dots] = \psi_s \epsilon_t + \psi_{s+1} \epsilon_{t-1} + \psi_{s+2} \epsilon_{t-2} + \dots$$

- The mean square error of the prediction is

$$E [(x_{t+s} - \hat{x}_{t+s|t})^2] = (1 + \psi_1^2 + \psi_2^2 + \dots + \psi_{s-1}^2) \sigma^2.$$

Predicting using ARMA models

Instead of observing the stationary process $(\epsilon_t)_{t \in T}$, one usually observes the process $(x_t)_{t \in T}$,

$$(1 - \phi_1 L - \dots - \phi_p L^p) x_t = (1 + \theta_1 L + \dots + \theta_q L^q) \epsilon_t$$

and it is often not reasonable to use the $MA(\infty)$ representation for predicting. However, invertibility of the process guarantees that it is irrelevant whether one observes $(x_t)_{t \in T}$ or $(\epsilon_t)_{t \in T}$.

- If the process is invertible, the s -step prediction at time point t is

$$\hat{x}_{t+s|t} = \begin{cases} \phi_1 \hat{x}_{t+s-1|t} + \phi_2 \hat{x}_{t+s-2|t} + \dots + \phi_p \hat{x}_{t+s-p|t} \\ \quad + \theta_s \epsilon_t + \theta_{s+1} \epsilon_{t-1} + \dots + \theta_q \epsilon_{t+s-q} & s = 1, 2, \dots, q, \\ \phi_1 \hat{x}_{t+s-1|t} + \phi_2 \hat{x}_{t+s-2|t} + \dots + \phi_p \hat{x}_{t+s-p|t} & s \geq q + 1, \end{cases}$$

where $\hat{x}_{\tau|t} = x_{\tau}$, when $\tau \leq t$ and the terms ϵ_t can be calculated recursively using the formula $\epsilon_t = x_t - \hat{x}_{t|t-1}$.

Remark

- If one predicts far to the future, that is $s > q$, the prediction does not take the moving average into account as there is not enough data to calculate it.
- Above, it was assumed that we observe infinitely long history of the values x_t . In practice, one can set the unobserved (historical) values to be equal to 0. This works well assuming that there are still plenty of observations and that the coefficients $|\theta_j|$ and $|\psi_i|$ are reasonably small.

Predicting using ARMA models: Optimality

- If the series x_t is a realization from ARMA(p, q) process with known parameters $\phi_1, \phi_2, \dots, \phi_p, \theta_1, \theta_2, \dots, \theta_q$, then the prediction $\hat{x}_{t+s|t}$ is optimal in the sense of the mean square error

$$\text{MSE}(\hat{x}_{t+s|t}) = E [(x_{t+s} - \hat{x}_{t+s|t})^2].$$

- In practice, the parameters of the ARMA(p, q) are not known and have to be estimated. Then the mean square error depends on estimations errors and the prediction is not, strictly speaking, optimal anymore. However, it can usually be thought to be close to optimal.

Predicting using ARMA models

- Consider $\hat{x}_{t+s|t}$ as a function of s .
 - ① Under pure AR models or ARMA models, $\hat{x}_{t+s|t} \rightarrow 0$ exponentially, as $s \rightarrow \infty$.
 - ② Under MA(q) model, $\hat{x}_{t+s|t}$ is equal to 0 for $s > q$.
- Thus ARMA models are suitable for short term forecasting, but not for long term forecasting!

Exponential smoothing

- Ad-hoc forecasting method, that does not rely on modeling.
- Applied widely.
 - Easy to calculate in practice.
 - Gives reasonably good predictions in practice, but is rarely optimal.

Simple exponential smoothing

- The value of x_{t+1} is predicted using a weighted sum of the previous observation $x_t, x_{t-1}, x_{t-2}, \dots$,

$$\hat{x}_{t+1|t} = \sum_{i=0}^{\infty} w_i x_{t-i}$$

- The weights $w_i = \alpha(1 - \alpha)^i$, $0 < \alpha < 1$, decrease exponentially. (That is why the method is called exponential smoothing.)
- The parameter α is called the smoothing factor.
- One can update the prediction using the formula

$$\hat{x}_{t+1|t} = \alpha x_t + (1 - \alpha)\hat{x}_{t|t-1} = \alpha \hat{\epsilon}_t + \hat{x}_{t|t-1},$$

where $\hat{\epsilon}_t = x_t - \hat{x}_{t|t-1}$ is the prediction error at time point t .

$$\hat{x}_{t+1|t} = \alpha x_t + (1 - \alpha)\hat{x}_{t|t-1}$$

- One can show that simple exponential smoothing gives optimal predictions in the case that x_t is an ARIMA(0,1,1) process:
 - Dx_t is a MA(1) process

$$Dx_t = x_t - x_{t-1} = \epsilon_t + \theta_1 \epsilon_{t-1}, \quad (\epsilon)_{t \in T} \sim WN(0, \sigma^2)$$

- Choose $\alpha = \theta_1 + 1$

Simple exponential smoothing

$$\hat{x}_{t+1|t} = \alpha x_t + (1 - \alpha)\hat{x}_{t|t-1}$$

- One can also show that simple exponential smoothing gives optimal predictions in the case that x_t is a noisy random walk, that is in the case:

$$x_t = m_t + \epsilon_t, \text{ where}$$

$$m_t = m_{t-1} + \eta_t, \text{ is a random walk}$$

$$(\epsilon_t)_{t \in T} \sim IID(0, \sigma_1^2), \quad (\eta_t)_{t \in T} \sim IID(0, \sigma_2^2)$$

- Optimal α depends on the signal-noise-ratio $\frac{\text{var}(\epsilon_t)}{\text{var}(\eta_t)}$.
- The proof relies on using the Kalman filter. (We talk about Kalman filter later.)
- Level m_t (of the random walk) can be estimated from the observations x_t :

$$m_t = \alpha x_t + (1 - \alpha)m_{t-1} \quad \text{and} \quad \hat{x}_{t+1|t} = m_t.$$

Double exponential smoothing

Simple exponential smoothing does not provide good predictions when there is a trend in the data. Double exponential smoothing is an extension of the simple exponential smoothing to the cases when there is a trend.

Double exponential smoothing

In double exponential smoothing, on top of predicting, the level m and the trend β are updated:

$$\begin{aligned}\hat{x}_{t+l|t} &= m_t + l\beta_t \\ m_t &= \alpha_1 x_t + (1 - \alpha_1)(m_{t-1} + \beta_{t-1}) \\ \beta_t &= \alpha_2(m_t - m_{t-1}) + (1 - \alpha_2)\beta_{t-1}.\end{aligned}$$

- One can write

$$\begin{aligned}m_t &= m_{t-1} + \beta_{t-1} + \alpha_1 \hat{\epsilon}_t \\ \beta_t &= \beta_{t-1} + \alpha_1 \alpha_2 \hat{\epsilon}_t,\end{aligned}$$

where $\hat{\epsilon}_t = x_t - \hat{x}_{t|t-1}$.

- The parameters α are called the data smoothing factors.
- The parameters β are called the trend smoothing factors.
- For well selected α_i , this method is optimal under ARIMA(0,2,2) models.

Exponential smoothing: Comments

- Exponential smoothing is often applied using fixed smoothing parameters.
- Sometimes smoothing parameters are estimated from the data. That usually provides a better fit.
- Applying SARIMA models is recommended, if possible.

Content

- 1 Prediction
- 2 Kalman filter
- 3 Dynamic regression

Kalman filter is a tool for predicting $\mathbf{x}_{t+1} = (x_{1(t+1)}, \dots, x_{k(t+1)})$, when we only observe a noisy version, $\mathbf{y}_t = (y_{1t}, \dots, y_{dt})$, of \mathbf{x}_t .

State space representation

Consider a situation where we wish to predict $\mathbf{x}_{t+1} = (x_{1(t+1)}, \dots, x_{k(t+1)})$, but we only observe $\mathbf{y}_t = (y_{1t}, \dots, y_{dt})$.

Assume that this system has the following **state space representation**

$$\begin{aligned}\mathbf{x}_{t+1} &= \mathbf{F}\mathbf{x}_t + \mathbf{v}_{t+1} \\ \mathbf{y}_t &= \mathbf{H}^\top \mathbf{x}_t + \mathbf{w}_t\end{aligned}$$

where \mathbf{y}_t and \mathbf{w}_t are d -variate random vectors, \mathbf{x}_{t+1} and \mathbf{v}_{t+1} are k -variate random vectors and

$$\text{cov}(\mathbf{v}_t, \mathbf{v}_s) = \begin{cases} \mathbf{Q}, & t = s \\ \mathbf{0}, & t \neq s \end{cases} \quad \text{and} \quad \text{cov}(\mathbf{w}_t, \mathbf{w}_s) = \begin{cases} \mathbf{R}, & t = s \\ \mathbf{0}, & t \neq s. \end{cases}$$

Assume also that \mathbf{v}_t , \mathbf{w}_t , and \mathbf{x}_t are mutually independent for all $t \geq 1$.

Kalman filter provides a prediction for $\mathbf{x}_{t+1} = (x_{1(t+1)}, \dots, x_{k(t+1)})$, based on the observations $\mathbf{y}_t = (y_{1t}, \dots, y_{dt})$.

- Prediction for \mathbf{x}_{t+1} at time point t is the conditional expected value

$$\hat{\mathbf{x}}_{t+1|t} := E[\mathbf{x}_{t+1} \mid \mathbf{Y}_t], \quad \mathbf{Y}_t := (\mathbf{y}_t^\top, \dots, \mathbf{y}_1^\top).$$

Kalman filter calculates the predictions $\hat{\mathbf{x}}_{1|()}, \hat{\mathbf{x}}_{2|1}, \dots, \hat{\mathbf{x}}_{T|T-1}$ recursively and every prediction is associated with the mean square error matrix

$$\mathbf{P}_{t+1|t} := E [(\mathbf{x}_{t+1} - \hat{\mathbf{x}}_{t+1|t})(\mathbf{x}_{t+1} - \hat{\mathbf{x}}_{t+1|t})^\top]$$

Kalman filter: Algorithm

- 1 The initial values (has to be selected):

$$\hat{\mathbf{x}}_{1|0} = E[\mathbf{x}_1]$$

$$\mathbf{P}_{1|0} = E [(\mathbf{x}_1 - E[\mathbf{x}_1])(\mathbf{x}_1 - E[\mathbf{x}_1])^\top]$$

- 2 Recursion formulae for the prediction $\hat{\mathbf{x}}_{t+1|t}$ and for the matrix $\mathbf{P}_{t+1|t}$ are

$$\hat{\mathbf{x}}_{t+1|t} = \mathbf{F}\hat{\mathbf{x}}_{t|t-1} + \mathbf{F}\mathbf{P}_{t|t-1}\mathbf{H}(\mathbf{H}^\top\mathbf{P}_{t|t-1}\mathbf{H} + \mathbf{R})^{-1}(\mathbf{y}_t - \mathbf{H}^\top\hat{\mathbf{x}}_{t|t-1})$$

$$\mathbf{P}_{t+1|t} = (\mathbf{F} - \mathbf{K}_t\mathbf{H}^\top)\mathbf{F}\mathbf{P}_{t|t-1}(\mathbf{F} - \mathbf{H}\mathbf{K}_t^\top) + \mathbf{K}_t\mathbf{R}\mathbf{K}_t^\top + \mathbf{Q},$$

where \mathbf{K}_t is the **Kalman gain**

$$\mathbf{K}_t := \mathbf{F}\mathbf{P}_{t|t-1}\mathbf{H}(\mathbf{H}^\top\mathbf{P}_{t|t-1}\mathbf{H} + \mathbf{R})^{-1}.$$

- 3 Prediction $\hat{\mathbf{y}}_{t+1|t}$ is obtained by applying the formula

$$\hat{\mathbf{y}}_{t+1|t} = \mathbf{H}^\top\hat{\mathbf{x}}_{t+1|t} \quad \text{ja}$$

$$E [(\mathbf{y}_{t+1} - \hat{\mathbf{y}}_{t+1|t})(\mathbf{y}_{t+1} - \hat{\mathbf{y}}_{t+1|t})^\top] = \mathbf{H}^\top\mathbf{P}_{t+1|t}\mathbf{H} + \mathbf{R}.$$

Kalman filter: Generalization

- It is possible to apply the Kalman filter approach also when one wishes to drop the linearity assumption. Then, the state space representation is given as

$$\begin{aligned}\mathbf{x}_{t+1} &= \mathbf{f}_t(\mathbf{z}_t, \mathbf{x}_t) + \mathbf{v}_{t+1} \\ \mathbf{y}_t &= \mathbf{h}_t(\mathbf{x}_t) + \mathbf{w}_t,\end{aligned}$$

where \mathbf{x}_{t+1} , \mathbf{y}_t , \mathbf{v}_{t+1} and \mathbf{w}_t are as above, \mathbf{z}_t is exogenous, independent of all the other variables, and \mathbf{f}_t and \mathbf{h}_t are functions that depend on time t , state \mathbf{x}_t and on the input \mathbf{z}_t .

- In this case, it is not as straightforward to obtain predictions as in the linear case, but the predictions are usually very good.

State space representation: MA(1) process

Example

Consider the MA(1) process

$$y_t = \epsilon_t + \theta_1 \epsilon_{t-1}.$$

Define the **state vector** \mathbf{x}_t and the noise \mathbf{v}_{t+1} by setting

$$\mathbf{x}_t = \begin{bmatrix} \epsilon_t \\ \epsilon_{t-1} \end{bmatrix} \quad \text{and} \quad \mathbf{v}_{t+1} = \begin{bmatrix} \epsilon_{t+1} \\ 0 \end{bmatrix}.$$

Then

$$\mathbf{x}_{t+1} = \mathbf{F}\mathbf{x}_t + \mathbf{v}_{t+1}, \quad \mathbf{F} = \begin{bmatrix} 0 & 0 \\ 1 & 0 \end{bmatrix}$$

and

$$y_t = \mathbf{H}^\top \mathbf{x}_t, \quad \text{where} \quad \mathbf{H}^\top = [1 \quad \theta].$$

Example: GPS navigation

- Assume that there are m_1 satellites that measure the pseudo-distances and their differences between an object and the satellite at time points t . Assume also that there are m_2 support stations that measure the distances between the object and the station. The vector $\mathbf{y}_t = (y_{1t}, \dots, y_{dt})^\top$, $d = 2m_1 + m_2$, contains the measured distances.
- State vector \mathbf{x}_t contains the location coordinates ξ_t and the speed \mathbf{v}_t of the object.

$$\mathbf{x}_t = \begin{bmatrix} \xi_t \\ \mathbf{v}_t \end{bmatrix}.$$

Example: GPS navigation

- Now, the corresponding state space model for the location (and the speed) of the object is

$$\begin{aligned}\mathbf{x}_{t+1} &= \mathbf{F}\mathbf{x}_t + \mathbf{u}_t \\ \mathbf{y}_t &= \mathbf{h}(\mathbf{x}_t) + \mathbf{w}_t,\end{aligned}$$

where $\mathbf{h}(\mathbf{x}_t)$ is a well chosen *nonlinear* function and \mathbf{u}_t and \mathbf{w}_t are the noise.

- 1 Prediction
- 2 Kalman filter
- 3 Dynamic regression**

In statistics, time series models can be divided into two groups:

- 1 **Autoprojective time series models** are models that involve only the time series to be forecasted (e.g. ARMA models).
- 2 **Dynamic regression models** are models that may involve the time series to be forecasted and the history of another time series as well.

Static linear regression

Consider the linear time series regression model

$$y_t = \alpha + \beta x_t + \epsilon_t, \quad t \in T,$$

where the response variable y_t and the explanatory variable x_t come from time series and where

- 1 $E[\epsilon_t | x_t] = 0, t \in T$
- 2 $\text{var}(\epsilon_t | x_t) = \sigma^2, t \in T$
- 3 $\text{cor}(\epsilon_t, \epsilon_s | x_t, x_s) = 0, t \neq s.$

Then the conditional expected value of the response variable y_t conditional on $x_t = x$ is

$$E[y_t | x_t = x] = \alpha + \beta x.$$

Assume that the explanatory variable x_s is the same constant x for all $s \leq t$:

$$x_s = x, \quad s \leq t$$

Then the conditional expected value of y_s is a constant for all $s \leq t$:

$$E[y_s | x_s = x] = \alpha + \beta x = y.$$

Static linear regression

$$y_t = \alpha + \beta x_t + \epsilon_t, \quad t \in T,$$

- Assume that, at time point $t + 1$, the explanatory variable grows by one unit:

$$x_{t+1} = x + 1.$$

- Then the conditional expected value of the response variable y_{t+1} is

$$E[y_{t+1} \mid x_{t+1} = x + 1] = \alpha + \beta(x + 1) = y + \beta.$$

- The conditional expected value remains the same until the value of the explanatory variable changes.

Static linear regression

$$y_t = \alpha + \beta x_t + \epsilon_t, \quad t \in T,$$

- The parameter β models the instant change in the conditional expected value of the response variable y_t as the value of x_t changes by one unit.
- The model is static:
 - The conditional expected value of the response variable does not change unless the value of the explanatory variable changes.
 - The conditional expected value of the response variable changes immediately, without any lags, when the value of the explanatory variable changes.

Dynamic regression: Idea

- In dynamic regression, the conditional expected value of the response variable may also change slowly, or progressively, when the value of the explanatory variable changes.
- A very simple example of a dynamic regression model is the distributed lag model. The distributed lag model is a model for time series data in which a regression equation is used to predict the current values of the response variable based on both, the current and the lagged (past) values of an explanatory variable.

Distributed lag model

A simple **distributed lag model** is given by

$$y_t = \alpha + \beta_0 x_t + \beta_1 x_{t-1} + \dots + \beta_p x_{t-p} + \epsilon_t, \quad t = p + 1, p + 2, \dots,$$

where the residual ϵ_t is white noise on condition x_t, \dots, x_{t-p} .

- In this model, there are $p + 1$ explanatory variables. Those are the values of x_s at time points $t, t - 1, \dots, t - p$.
- The response variable y_t , at time point t depends on
 - The value of the variable x_s at time point t .
 - The lagged (past) values of the variable x_s .
- The conditional expected value of the response variable y_t is

$$E[y_t \mid x_t, x_{t-1}, \dots, x_{t-p}] = \alpha + \beta_0 x_t + \beta_1 x_{t-1} + \dots + \beta_p x_{t-p}.$$

Distributed lag model

$$y_t = \alpha + \beta_0 x_t + \beta_1 x_{t-1} + \dots + \beta_p x_{t-p} + \epsilon_t, \quad t = p + 1, p + 2, \dots,$$

Remark

The distributed lag models are stationary with respect to time in the sense that the regression parameters depend only on the corresponding lags, not on the actual time points:

- The regression parameters can be seen as the derivatives of y with respect to x with different lags:

$$\beta_s = \frac{\partial y_t}{\partial x_{t-s}} = \frac{\partial y_{t+s}}{\partial x_t}.$$

- Without this condition, the parameters β_0, \dots, β_p would depend on time, and not only on the length s of the corresponding time interval.

Distributed lag model

$$y_t = \alpha + \beta_0 x_t + \beta_1 x_{t-1} + \dots + \beta_p x_{t-p} + \epsilon_t, \quad t = p + 1, p + 2, \dots,$$

- Assume that the explanatory variable x_s is the same constant x for all $s \in \{t - p, \dots, t\}$,

$$x_s = x, \quad s \in \{t - p, \dots, t\}.$$

- Then the conditional expected value of y_t

$$E[y_t \mid x_t, x_{t-1}, \dots, x_{t-p}] = \alpha + \beta_0 x + \beta_1 x + \dots + \beta_p x = \alpha + \beta x = y,$$

where $\beta = \beta_0 + \beta_1 + \dots + \beta_p$.

Distributed lag model

Assume that

$$x_s = \begin{cases} x & \text{when } s \in \{t - p, \dots, t\} \\ x + 1 & \text{when } s = t + 1, t + 2, \dots, \end{cases}$$

Then

$$E[y_t \mid x_t, \dots, x_{t-p}] = \alpha + \beta_0 x + \beta_1 x + \dots + \beta_p x = \alpha + \beta x = y$$

$$E[y_{t+1} \mid x_{t+1}, \dots, x_{t-p+1}] = \alpha + \beta_0(x + 1) + \beta_1 x + \dots + \beta_p x = y + \beta_0$$

$$E[y_{t+2} \mid x_{t+2}, \dots, x_{t-p+2}] = \alpha + \beta_0(x + 1) + \beta_1(x + 1) + \beta_2 x + \dots + \beta_p x$$

\vdots

$$E[y_{t+p+1} \mid x_{t+p}, \dots, x_t] = \alpha + \beta_0(x + 1) + \dots + \beta_{p-1}(x + 1) + \beta_p x$$

$$E[y_{t+p+1} \mid x_{t+p+1}, \dots, x_{t+1}] = \alpha + \beta_0(x + 1) + \dots + \beta_p(x + 1) = y + \beta$$

Distributed lag model

Assume that the value x_s is the same constant $x + 1$ and does not change after the time point $t + 1$. Then the conditional expected value of y_s

$$E[y_s \mid x_s, \dots, x_{s-p}] = \alpha + \beta_0(x + 1) + \dots + \beta_p(x + 1) = y + \beta$$

remains the same for $s \geq t + p + 1$.

Distributed lag model

$$y_t = \alpha + \beta_0 x_t + \beta_1 x_{t-1} + \dots + \beta_p x_{t-p} + \epsilon_t, \quad t = p + 1, p + 2, \dots,$$

The regression parameters of the model can be interpreted as follows.

- (i) The parameter β_0 models the instant change in the conditional expected value of the response variable y_t as the value of x_t changes by one unit.
- (ii) The sum

$$\beta = \beta_0 + \beta_1 + \dots + \beta_p$$

models the long term change in the conditional expected value of the response variable y_t as the value of x_t changes by one unit.

Distributed lag model: Parameter estimation

The parameters of the distributed lag model

$$y_t = \alpha + \beta_0 x_t + \beta_1 x_{t-1} + \dots + \beta_p x_{t-p} + \epsilon_t, \quad t = p + 1, p + 2, \dots,$$

can be estimated using classic linear regression.

- Problems:

- If the autocorrelations of the process $(x_t)_{t \in T}$ are non-zero, the explanatory variables of the regression model are multicollinear.
- If we have n observations of the series $(x_t)_{t \in T}$, we have $n - p$ observations for estimating $p + 2$ parameters.

Distributed lag model: General form

In distributed lag models, y_t might depend on the entire history of the explanatory process, and the explanatory process may be k -variate. That is

$$y_t = \alpha + \sum_{i=0}^{\infty} \beta_{1i} X_{1(t-i)} + \sum_{i=0}^{\infty} \beta_{2i} X_{2(t-i)} + \dots + \sum_{i=0}^{\infty} \beta_{ki} X_{k(t-i)} + \epsilon_t,$$

where $(\epsilon_t)_{t \in T} \sim WN(0, \sigma^2)$.

- If infinitely many of the regression parameters β_{ji} differ from zero (for example, $\beta_{ji} = \beta_{j0} \delta_j^i$, $|\delta_j| < 1$), then a change in the value of the explanatory variable has an effect on the conditional expected value of y_t infinitely long after that.

Simple ARMAX model

An ARMAX model is a combination of an ARMA model and a linear regression model, where the response process y_t depends on

- the history of the autoregressive part of the process
- noise
- an exogenous variable x_t .

$$y_t = \epsilon_t + \sum_{i=1}^p \phi_i y_{t-i} + \sum_{i=1}^q \theta_i \epsilon_{t-i} + \sum_{i=0}^b \eta_i x_{t-i},$$

where

- $(\epsilon_t)_{t \in T} \sim WN(0, \sigma^2)$,
- ϕ_1, \dots, ϕ_p are the parameters of the AR part,
- $\theta_1, \dots, \theta_q$ are the parameters of the MA part, and
- η_0, \dots, η_b are the parameters related to the exogenous variable x_t .

As when applying ARMA models, also when applying ARMAX models, one can model seasonal changes. In ARMAX models, the number of exogenous variables is not limited to one. An ARMAX model, in general form, can be given as

$$\Phi(L)y_t = \Theta(L)\epsilon_t + \sum_{j=1}^k H_j(L)x_{jt},$$

where $\Phi(L)$, $\Theta(L)$ and $H_j(L)$ are similar lag polynomials as in the case of ARMA models and the $x_j = (x_{j1}, x_{j2}, \dots, x_{jt})$ are the observed explanatory variables.

An ARIMAX model is like an ARIMA model, but it contains one or more exogenous variables.

- The differences $D^h y_t$ (for some order h) of an ARIMAX model form an ARMAX model and the parameters for the differences can be estimated similarly as in the case of ARIMA models.

R: `arima()`, `arimax()`.

- ARIMAX models are reasonably general, but they are not suitable for long term forecasting. Moreover, ARIMAX models are linear on the explanatory variables.
- In general, we may consider models

$$f(\mathbf{y}_t) = g((\mathbf{y}_s)_{\{s \leq t-1\}}, (\mathbf{x}_s)_{\{s \leq t\}}) + \epsilon_t,$$

where the d -variate $\mathbf{y}_t = (y_{1t}, \dots, y_{dt})$ depends on its own history and on the history of the k -variate time series $(\mathbf{x}_t)_{t \in T}$, $\mathbf{x}_t = (x_{1t}, \dots, x_{kt})$, $t \in T$, through some nonlinear functions f and g .

- In practice, we often have to rely on these general models, but developing theory for them is very difficult.
 - One has to know (or estimate) the functions f and g and then proceed case by case.

References:

- 1 Brockwell, P., Davis, R. (2009): Time Series – Theory and Methods, Springer
- 2 Hamilton, J. (1994): Time Series Analysis, Princeton University Press
- 3 Ali-Löytty, S. (2004): Kalmanin suodatin ja sen laajennukset paikannuksessa, Diplomityö, TTY

Next week:

- Guest lecturer Marko Voutilainen
- Summary