

You have from the 14 of March to the 23 of March to finish the exam. Make sure you show all the crucial steps that led to the final answer of each question. Interpretation is part of the examination. No need to show the R codes, just the relevant output to answer the questions. Good luck!

1. (25 pts) Consider the following ARMA(2,1) time series model where ε_t follows a white noise process:

$$y_t = 2 - 0.8y_{t-1} - 0.1y_{t-2} + \varepsilon_t + 0.8\varepsilon_{t-1} \quad (1)$$

- (a) (10 pts) Find the set of homogeneous solutions. Explain the shape of this solution. Is it stable?
 - (b) (5 pts) Is (1) stationary and invertible? Explain.
 - (c) (5 pts) Given $y_0 = 1$ and $y_1 = 0.2$, use the iteration method to find the solution to y_3 . Is the solution of the form of a constant plus past shocks? Does y_3 depend on ε_2 ? Explain why or why not.
 - (d) (5 pts) Compute the ACF for the first and second lag. How does the moving average component affect the ACF for the first two lags?
2. (25 pts) Suppose we have the following VAR with two variables, y_t and z_t :

$$\begin{aligned} y_t &= 0.1 + 0.1y_{t-1} + 0.4z_{t-1} + e_{1t} \\ z_t &= 0.8 - 0.3y_{t-1} + e_{2t} \end{aligned}$$

and that the residual variance-covariance matrix is given by:

$$\Sigma = \begin{bmatrix} 0.5 & 0.4 \\ 0.4 & 2 \end{bmatrix}$$

- (a) (2 pts) Find the two steps-ahead conditional forecast of y_t , $E_t[y_{t+2}]$. Discuss how z_t affects $E_t[y_{t+2}]$.
 - (b) (13 pts) Assume z_t does not affect contemporaneously y_t (recursive identification approach). Show your matrix B after imposing this restriction and find all coefficients of the structural VAR (including σ_y and σ_z).
 - (c) (5 pts) Under the restriction imposed in b), compute the IRF for both variables to a one standard deviation shock in z_t at impact and step 1. Interpret them. What if the size of the shock doubles. What are the responses at impact and step 1?
 - (d) (5 pts) Compute the FEVD of y_t at impact and step 1. How do your findings in terms of FEVD complement your conclusions in c)?
3. (25 pts) ARMA in practice. Use the library `quantmod` to get Amazon's stock price data. Use `getSymbols("AMZN", src="yahoo", from="2007-12-01", to="2020-03-13")`.

- (a) (2 pts) Plot the daily close price data. Does it look stationary? Plot its ACF, does your conclusion change?
 - (b) (5 pts) Compute the Amazon daily stock returns. Take log of the close date and take the first difference. Plot the daily returns, do they look stationary? Plot its ACF, does it change your conclusion? Now formally test for the unit root of the close price in levels and for the daily returns. Are they stationary?
 - (c) (5 pts) Fit an ARMA(2,1) with an intercept to the daily Amazon stock returns. What is the unconditional mean of the Amazon daily stock return in percentage?
 - (d) (3 pts) Check the model adequacy with the ACF plot of the residuals and the Ljung-Box test with 10 lags.
 - (e) (5 pts) Apply the `as.ts()` to your xts return series. Re-estimate the ARMA(2,1) to this new formatted series. Load the Forecast library and forecast the stock return for the next 5 days. Would you recommend buying the Amazon stock? Why or why not.
 - (f) (5 pts) How do the AR and MA components estimated coefficients relate to your forecasts in e)?
4. (25 pts) VAR in practice. Use the library `readxl` to load the dataset `data_midterm_2020.xlsx`. This dataset consists of quarterly Portuguese data on 5 variables in volume: GDP, household private consumption, government consumption, investment and taxes on products.
- (a) (5 pts) Get the data ready for analysis. 1) Remove the first column of the data; 2) apply the `ts()` function to the data in order to get the data in time series format; 3) put the data in logs and multiply by 100; 4) Take the first difference to all variables; 5) put the variables of the data in the following order: G, TAX, C, I, GDP.
 - (b) (5 pts) Fit a VAR model to the transformed data. Choose the lags based on the Schwarz information criterion (SC). Check if the VAR model is adequate (use an ADF test on all 5 residuals series).
 - (c) (10 pts) Compute the IRFs (20 steps) and plot them. What do you learn from them in terms of how a government expenditure shock affects the economy?
 - (d) (5 pts) Compute the FEVD (20 steps). How relevant, in relative terms, are G shocks in explaining variation in GDP ?
 - (e) (5 pts, Bonus points, challenge question for the bold!) Create a new dataset in memory that removes the last observation of the dataset you were working on (removes observation 99, note that you may need to apply the `ts()` function to the new dataset again). Re-estimate a VAR model to this subsample and forecast one step-ahead GDP. Now fit an ARMA model with the `auto.arima()` function to GDP series alone and make also a one step-ahead forecast. Compare both forecasts (VAR and ARMA) to the original observation you removed. Which forecast is better? Explain.