

ASSIGNMENT 5. AUTOREGRESSIVE MOVING AVERAGE (ARMA) MODELING

1. Run `geosa5.m`, selecting one series from either V1 or V2 for analysis. Run the script on either the full series length or some sub-period. Use the FPE (Final Prediction Error) method to pick the best ARMA model.
2. (Caption to Fig. 1) Three-part plot: time series plot, acf & pacf, spectrum. Describe how the patterns of decay in the acf and pacf support or refute an AR(1) generating process
3. (Caption to Fig. 2) Acf of residuals for ARMA model selected by FPE. Using the information in the figure, discuss whether persistence is of *practical* importance for the original time series and whether the selected model adequately effectively removes the persistence. Using `Results.model`—returned to the workspace after running `geosa5--` and the Matlab **present** function, tell whether the model parameters are significantly different from zero at the 95% significance level ($\alpha = 0.05$).
4. (Caption to Fig. 3) Zoomed original and whitened time series. Use an arrow to point out a year (or observation number) in which whitening makes an especially large difference to the series. Refer back to the equation in Fig. 2 to explain why, specifically by referring to one or more of the estimated model coefficients. If whitening makes no discernible difference, explain why in terms of the model and its estimated coefficients.
5. (Caption to Fig. 4) Spectra of original and whitened series. Are there noticeable differences in the distribution of variance over frequency for the two series? Why or why not? What is the ratio of the area under the whitened-series spectra to the area under the original-series spectra? Clue: you can compute the answer from the statistics annotated on Figure 2 and the definition of the spectrum.

Running `geosa5.m`

1. `>geosa5`
2. Message box: message introducing `geosa5.m`; click OK to remove message and move on
3. Click OK to acknowledge the message on limitations on order of candidate models
4. Respond to input dialog with the name of your data file; click OK
5. Menu: select either V1 or V2 as the source structure for your time series
6. Menu: click on the time series to be analyzed. You can only select one series, and when you do, an asterisk appears in its box. If you click again on another series, that series becomes your selection. When satisfied with your choice, click “Accept Selection”
7. Input dialog: select either the default (full length) or any sub-period for analysis

8. Click OK to acknowledge the message box with information on your selection
9. Input dialog: select either the default (full length) or any sub-period for analysis
10. Input dialog: select either the default (most recent 30 observations) or any sub-period of maximum length 30 observations for zoomed time series plots of original and whitened series. The purpose of zooming here is to allow a detailed look at the observation-to-observation fluctuations. The full time series plot may be too condensed to pick out the details. Fig 1 appears, with three sets of axes and graphics. A menu also appears (see next step).

top: time series plot

lower left: acf and pacf out to lag 20. Includes 95% CI.

lower right: spectrum with default lag-window of M equal to 1/10 the sample length

11. Menu: choose the model structure (AR or ARMA) for modeling the series. If you choose AR or ARMA, you will need to specify the model order. If you choose to let the final prediction error (FPE) criterion pick the model, both the model structure and order are selected automatically.

With each choice, Figure 2 is generated or revised to show the plot of acf of the residuals from fitting the model. Annotated on the plot are the model equation, the results of the Portmanteau-Q test, and the percentage variance due to persistence. This last quantity is computed as $(\text{var}(x) - \text{var}(e)) / \text{var}(x)$, where $\text{var}(x)$ is the variance of the original data, and $\text{var}(e)$ is the variance of the AR or ARMA residuals. Note that for a “practically unimportant” model, the residual series will be almost identical to the original series (mean subtracted), the numerator will approach zero, and the percentage of variance due to persistence will approach zero.

You may try various options at this point and check out the results before choosing “Satisfied”. Note that the program bombs if you pick “satisfied” first. Try various models and structures. But for your plot to hand in, use the FPE option. Then click “Satisfied”.

Zoomed time series plots of the original series and residuals then appear in Figure 3, and spectra of the original series (with 95% CI) and the series whitened by the selected ARMA model in Fig 4. A menu also appears...

12. Menu: interactively change the lag window as needed to vary the smoothness and resolution of your spectra. When you are happy, click “Accept spectrum”. The final plots appear and geosa5 finishes.
13. The closing message refers to a structure Results that remains in the workspace. Results.what describes the fields in Results. Results.model contains the parameters and associated information on the model. You can use

```
>present(Results.model)
```

to get a screen listing of the model equation, including the model parameters and their standard deviations. To be significantly different from zero, these parameters should be more than TWO standard deviations from zero.

PROGRAMMING NOTES

Script geosa5 makes extensive use of functions in MATLAB's System Identification toolbox.

Selected Matlab functions called :

iddata -- create DATA OBJECT to be used for Identification routines

ar -- computes AR-models of signals using various approaches

armax -- computes the prediction error estimate of an ARMA model.

predict -- computes the k-step ahead prediction from AR or ARMA model

present -- presents a parametric model on the screen

detrend -- removes a linear trend from a vector

Selected user-written Matlab functions called

acf -- autocorrelation function and approximate 95% confidence bands

pacf -- compute partial autocorrelation function

portmant -- compute Portmanteau Q statistic