Carnegie Mellon University
Tepper School of Business

45-912 Business Forecasting with Time Series Models Spring 2019
(6 units)

Professor: Fallaw Sowell Email: fs0v@andrew.cmu.edu
Office: TEP 4213 Phone: (412)-268-3769

Class Meetings:

Section E4	Monday	6:30 - 8:30pm	room TEP 2118

DELIVERABLES

Students will be able to take a new time series and determine its trend and seasonal characteristics. After accounting for trend and seasonal characteristics, the students will be able to estimate an ARMA model. Students will be able to estimate Multivariate VAR models. For these estimated models the students will be able to make forecasts and summarize the uncertainty inherent in the forecasts. Issues associated with linear regression with time series data will be addressed.

OBJECTIVE AND OVERVIEW

This course is an introduction to basic time series models. The course uses R to create forecasts. Forecasts are constructed from the estimated summary statistics and parameters of time series models: univariate ARIMA and multivariate VAR . Students should be able to interpret the uncertainty in the forecasts and in the estimated parameters. Diagnostic statistics and model selection criteria are presented.

GRADING AND ASSIGNMENTS

Your course grade will be determined by a set of assignments and quizzes. The assignments are designed to give practice and experience with the techniques and models presented in the course. The final grade will be split as follows:

20%	in class quizzes
20%	Individual Assignments
60%	Group Assignments

Individual Assignments: Simulated data I will give you.

This set of assignments is designed to help students learn how to model time series. The focus is on obtaining the correct models.

You will have a personal data set that I created. Each data set is a collection of 5 series denoted Y_t, X_t, W_t, V_t and U_t. Each series has 500 observations. The odd numbered data sets are quarterly observations (1894:Q1 to 2018:Q4) and the even numbered data sets are monthly observations (July 1977 to Feb 2019).

You will apply the tests and procedures taught in the course to these series. I created these data so that the concepts taught in the course are highlighted. The homework assignments will demonstrate your understanding of the techniques. The assignments are due at the beginning of class and should be on one side of a single sheet of paper.

Sheet 1. Due April 1. Submit the summary data for the data set and plot the Y_t series in R. This is a demonstration that you are using R and should not be performed with excel.

Sheet 2. Due April 24. Submit the appropriate (transformed to stationarity) ARIMA models for the X_t and the V_t series.

Group Assignments: Actual data the group selects. Teams of two to four students.

These assignments focuses on forecasting and presenting the forecasts in a business setting.

For this set of assignments you will write a paper during the mini. You will apply the different estimation techniques presented in the course to your data series. You will report the parameter estimates and forecasts.

Your paper will grow over the mini and will be due periodically. A brief description of your data, an R plot of the series and summary statistics are due by 5:00 pm on Sunday March 31 (no more than two pages, one page is likely sufficient). A midterm paper is due by 5:00 pm on Sunday April 14 and is worth half of the Group Assignments grade.

The midterm paper will be the first part of the final paper and should include:

1. A general description of the data.
2. A plot of the series.
3. Summary statistics.
4. Any problems with the series. (Are there any outliers or unusual events? Typically this is not an issue for data obtained from the Government. This is more an issue for students working with data from their current employers.)
5. Related to the previous item, explanations of any simplifying assumptions used in modeling the series. (How are outliers treated? Again, typically, this is not an issue for data obtained from the Government. This is more an issue for students working with data from their current employers.)
6. At least one set of forecasts and confidence intervals for the series from a model presented in the first three weeks of the course. (I expect that most students will use exponential smoothing for the forecasts.) The paper should be at the level of a business document. An appendix should include technical details of the analysis that supports your business document.

The final paper will be due by 9:00 pm on Monday May 6. The final paper will include at least one additional set of forecasts and confidence intervals from the most appropriate model presented in the course. Most students will find other data that help explain their series and will present forecasts from a VAR model. Your paper should include at least 3 other series that were considered in building the appropriate VAR (the week 5 topic) model.

Submitting your midterm and final requires:

- Email me a single zip file that contains, your paper in a pdf file, the data in an excel and the R programs you used in preparing your paper. The R program should run directly from the directory with no adjustments for paths.

SCHEDULE FOR SPRING 2019

WEEK	DATES	TOPICS
0	March 23	Intro to Forecasting, Overview, Smoothing
1	March 25	Conditioning and Intro to basic models
	March 31	GROUP ASSIGNMENT: DATA DESCRIPTION, PLOT & SUMMARY
2	April 1	Basic ARMA models
	April 1	INDIVIDUAL ASSIGNMENT: SHEET 1.
3	April 8	Transformations and Selecting ARIMA models
	APRIL 14	GROUP ASSIGNMENT: MIDTERM PAPER. Due by 5:00 PM
4	April 15	Seasonal ARIMA Models
5	April 22	Multivariate VAR Models
	APRIL 22	INDIVIDUAL ASSIGNMENT: SHEET 2.
6	April 29	Regression with time series
	MAY 6	GROUP ASSIGNMENT: FINAL PAPER. Due by 9:00 PM

OPTIONAL HELP SESSIONS

Optional recitations. Historically, these are basically designed to provide assistance with R. The (R help sessions) recitations schedule is:

Date		Time	Room
Saturday	March 30	10:00 am - noon	Room TPR 4219
Saturday	April 6	9:00 am - 11:00am	Room TPR 4219
Saturday	April 13	10:00 am - noon	Room TPR 4219
Saturday	April 20	10:00 am - noon	Room TPR 4219
Saturday	April 27	10:00 am - noon	Room TPR 4219
Saturday	May 4	10:00 am - noon	Room TPR 4219

COURSE SUMMARY

Week 0:

This lecture will start with an overview of the course structure, how the lectures will be organized and how the course grade will be determined.

The game of "Texas Hold'em Poker" will be used to highlight the features of forecasting in a business setting.

Next we will introduce the basic idea of a time series, deterministic and stochastic. The students should be able to give examples.

We will then explain the basic "Approaches" to time series analysis and forecasting. This is followed by a basic introduction to R.

The first forecasting technique you will study is Exponential Smoothing. We will start with a basic model, then add a trend and end with a series that has a trend and a seasonal cycle.

Week 1:

Make the connection between the introductory statistics courses and time series models. We will review the bivariate conditional distributions with normal errors. The main point is that the needed information is the variance-covariance. You only need a covariance structure to determine the conditional expectation and this is your forecast.

Introduce basic time series models. For most of these simple models the forecasts can be determine based on the student's knowledge from prior courses.

The issue of model selection is presented. Given all the possible variables, how do we select the best model?

Introduction to the lag operator. Write the AR model with lag operator notation. Note that an AR(1) can be written as an MA(∞) with only one parameter.

We will learn about stationarity. This is needed to give some structure to the time series model.

Introduce the Wold Representation Theorem and think of the ARMA(p, q) model as an approximation to the Wold representation for a weakly stationary time series.

Week 2:

In this lecture you will learn how to identify the appropriate ARMA(p, q) model for a given series.

The Ljung-Box Q-statistic is introduced as a test for white noise.

The structure of the autocovariance function and the partial autocorrelation function for different models is presented. These give guidance in the selection of the appropriate ARMA(p, q) model.

Week 3:

Start with different transformations that are needed to obtain a weakly stationary series. Weak stationarity is needed to apply the Wold Decomposition theorem presented in Week 2. The log transformation and the first difference are considered. The US GDP series and the SP500 data are modelled.

Week 4:

Models with seasonal variation are presented. The nonfarm payroll data is modelled.

Week 5:

In this lecture you will extend the ARIMA model into a multivariate VAR model.

Week 6:

In this lecture you learn how to correctly run a linear regression with time series data. This is an alternative approach to the VAR models presented in week 5.

RECORDINGS

No student may record any classroom activity without the express written consent of Professor Fallaw Sowell. If a student believes that he/she is disabled and needs to record or tape classroom activities, he/she should contact the Office of Equal Opportunity Services, Disability Resources to request an appropriate accommodation.