Carnegie Mellon University
Tepper School of Business

45-912 Business Forecasting with Time Series Models Spring 2016
(6 units)

Professor: Fallaw Sowell Email: fs0v@andrew.cmu.edu
Office: GSIA 313 (Old Building) Phone: (412)-268-3769

Class Meetings:
Section A Tuesday & Thursday 3:30 - 5:15pm rm 153 Posner Hall

TEXTBOOKS

1. Introduction to Time Series Analysis and Forecasting with Applications of SAS and SPSS, by Robert Yaffee with Monnie McGee, ISBN 0-12-767870-0.
2. SAS for Forecasting Time Series, second edition, by John C. Brocklebank and David A. Dickey, ISBN 0-471-9566-8.
3. I am not requiring a textbook for SAS. If you would like an introductory book, I would recommend The Little SAS Book, forth edition, by Lora D. Delwiche and Susan J. Slaughter, ISBN 9781599947259.
4. SAS documentation is available at
   http://support.sas.com/documentation/94/index.html

The CMU bookstore has been asked to stock the three books.

DELIVERABLES

Students will be able to take a new time series and determine its trend and seasonal characteristics. They will be able to determine if the series has conditional heteroskedasticity.

After accounting for trend and seasonal characteristics, the students will be able to estimate an ARMA model and when appropriate, an ARCH or GARCH model. For these estimated models the students will be able to make forecasts and summarize the uncertainty inherent in the forecasts.

OBJECTIVE AND OVERVIEW

This course is an introduction to basic time series models. The course uses SAS to create forecasts. Forecasts are constructed from the estimated summary statistics and parameters of time series models: mainly ARIMA but also ARCH and GARCH. 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. 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:

30%	Individual Assignments
70%	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. The Y_t time series is your data of interest. The other four time series are thought to perhaps help explain the Y_t series. Each series has 500 observations. The odd numbered data sets are quarterly observations (1891:Q1 to 2015:Q4) and the even numbered data sets are monthly observations (July1974 to Feb2016).

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 March 29. Submit the summary data for the data set and plot the Y_t series in SAS. This is a demonstration that you are using SAS and should not be performed with excel.

Sheet 2. Due April 12. Submit which series needs to be logged and which series need to be first differenced. Determine the transformations needed for each of your 5 series to be weakly stationary. Submit the LOGTEST SBC test values and the Augmented Dickey Fuller Statistics for the most appropriate models.

Sheet 3. Due April 19. Submit the appropriate ARIMA models for the X_t and the V_t series.

Sheets 4. Due April 26. Submit the appropriate ARIMA models for the X_t, W_t, V_t and U_t series. Two of the series X_t, W_t, V_t and U_t will help explain the Y_t series. Submit the transfer function model for the Y_t series. (This should be submitted on no more than three single sided sheets of paper.)

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, a SAS plot of the series and summary statistics are due by 5:00 pm on Saturday March 26 (no more than two pages, one page is likely sufficient). A midterm paper is due by 5:00 pm on Saturday April 9 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 5:00 pm on Sunday May 1. 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 transfer function model. Your paper should include at least 4 other series that were considered in building the appropriate transfer function (the week 5 topic) model.

Submitting your midterm and final requires:

- Provide me with a hard copy of your paper.
- Email me a single zip file that contains the data and the SAS programs you used in preparing your paper. Your data should be in a sasuser file.

SCHEDULE FOR SPRING 2016

WEEK	DATES	TOPICS	Yaffee	Brocklebank	SAS PROC's
1	March 17-22	Intro to Forecasting, Overview, Exponential Smoothing	Chapters 1 & 2		PROC's FORECAST, GPLOT, SUMMARY
2	March 24-29	Intro to ARIMA models	Chapters 3 & 4	Chapters 1 & 2	PROC ARIMA
	MARCH 26	GROUP ASSIGNMENT: DATA DESCRIPTION, PLOT & SUMMARY			
	MARCH 29	INDIVIDUAL ASSIGNMENT: SHEET 1.			
3	March 31 - April 5	Selecting ARIMA models	Chapters 3 & 4	Chapter 3	PROC ARIMA
4	April 7-12	Seasonal ARIMA Models	Chapter 5	Chapter 4	PROC ARIMA
	APRIL 9	GROUP ASSIGNMENT: MIDTERM PAPER. Due at my office by 5:00 PM			
	APRIL 12	INDIVIDUAL ASSIGNMENT: SHEET 2.			
5	April 14-19	Multivariate Models Transfer Functions	Chapter 9	Chapter 4	PROC ARIMA
	APRIL 19	INDIVIDUAL ASSIGNMENT: SHEET 3.			
6A	April 21-	Modeling Unique Events	Chapter 5	Chapter 4	PROC ARIMA
6B	April 28	Regression with time series errors and Models w/ Conditional Variance	Chapter 10	Chapters 2 & 5	PROC AUTOREG
	APRIL 26	INDIVIDUAL ASSIGNMENT: SHEET 4.			
	MAY 1	GROUP ASSIGNMENT: FINAL PAPER. Due at my office by 5:00 PM			

OPTIONAL HELP SESSIONS

This course has optional recitations. These will be run by my TA, Camilo Botia and are basically designed to provide assistance with SAS.

The (SAS help sessions) recitations schedule is:

Date		Time	Room
Saturday	March 26	10:00am-noon	343 Posner Hall
Saturday	April 2	10:00am-noon	343 Posner Hall
Saturday	April 9	10:00am-noon	343 Posner Hall
Saturday	April 16	10:00am-noon	343 Posner Hall
Saturday	April 23	10:00am-noon	343 Posner Hall
Saturday	April 30	10:00am-noon	343 Posner Hall

COURSE SUMMARY

Week 1:

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 SAS.

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 2:

This lecture makes 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 3:

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

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 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 4:

In this lecture you will learn how to have SAS forecast once an ARMA model is selected. Models with seasonal variation are then presented. Tests for seasonal unit roots are presented.

In the SAS programs, we will learn about controlling the SAS plots to focus attention on the appropriate information.

Week 5:

In this lecture you will extend the ARIMA model into a multivariate transfer function. This is how other variables are used to help forecast a time series.

Week 6A:

The students will learn how to account for one-time events that affect our data, e.g. a strike, a war, a terrorist event, etc.

In time series analysis this is sometimes called intervention analysis. The basic idea is to use dummy variables to capture unique events. The coefficients on the dummy variables model the unique event so that the other terms in the model can capture the model's basic features.

Week 6B:

In this lecture you learn how to estimate and make forecasts for models with conditional variance, e.g., ARCH and GARCH models. Conditional variance models are widely used to capture the behavior of financial data.

Another way to think about these models is how to correctly run a linear regression with time series data.

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.