

Carnegie Mellon University
Tepper School of Business

45-921
(6 units)

Business Forecasting with Time Series Models

Spring 2009

Professor: Fallaw Sowell
Office: GSIA 313 (Old Building)

Email: fs0v@andrew.cmu.edu
Phone: (412)-268-3769

Class Meetings:

Section A	Tuesday & Thursday	3:30pm - 5:20pm	rm 153 Posner Hall
Section E	Thursday	6:30pm - 9:30pm	rm 151 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 0-471-95666-8.
3. I am not requiring a textbook for SAS. If you would like an introductory book, I would recommend The Little SAS Book, third edition, by Lora D. Delwiche and Susan J. Slaughter, ISBN 1-59047-333-7.
4. SAS documentation is available at
<http://support.sas.com/documentation/onlinedoc/91pdf/index.html>

The three books are available at the CMU bookstore.

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

Your course grade will be determined by a paper you write during the mini. You will apply 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 on Saturday March 28 (no more than two pages, one page is likely sufficient). A midterm paper is due on Saturday April 11 and is worth half your course 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. A brief description of the system or mechanism that generated the series.
5. Any problems with the series. (Are there any outliers or unusual events?)
6. Related to the previous item, explanations of any simplifying assumptions used in modeling the series. (How are outliers treated?)
7. 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.

The final paper will be due by 5:00 pm on Sunday May 3. The final paper will include at least one additional set of forecasts and confidence intervals from the most appropriate model presented in the course. (I expect that most students will find other data that help explain their series and will present forecasts from a transfer function 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.

SCHEDULE FOR SPRING 2009

WEEK	DATES	TOPICS	Yaffee	Brocklebank	SAS PROC's
1	March 17-19	Intro to Forecasting, Overview, Exponential Smoothing	Chapters 1 & 2		PROC's FORECAST, MEANS, GPLOT, SUMMARY
2	March 24-26	Intro to ARIMA models	Chapters 3 & 4	Chapters 1 & 2	PROC ARIMA
MARCH 28 ASSIGNMENT: DATA DESCRIPTION, PLOT & SUMMARY					
3	March 31 -April 2	Selecting ARIMA models	Chapters 3 & 4	Chapter 3	PROC ARIMA
4	April 7-9	Seasonal ARIMA Models	Chapter 5	Chapter 4	PROC ARIMA
APRIL 11 ASSIGNMENT: MIDTERM PAPER					
5	April 14-16	Multivariate Models Transfer Functions	Chapter 9	Chapter 4	PROC ARIMA
6	April 21-23	Regression with time series errors and Models w/ Conditional Variance	Chapter 10	Chapters 2 & 5	PROC AUTOREG
7	April 28-30	Modeling Unique Events	Chapter 5	Chapter 4	PROC ARIMA
MAY 3 ASSIGNMENT: Final Paper Due at my office by 5:00 PM					

Optional Help Sessions

This course will have an optional recitation. The recitation will be run by my TA, Stephen Lenkey. These are basically designed to provide assistance with SAS. Of course, Stephen can also answer your basic time series questions.

The (SAS help sessions) recitations schedule is:

	Date	Time	Room
Saturday	March 21	9:00am-11:00am	rm 146 Posner
Saturday	March 28	9:00am-11:00am	rm 146 Posner
Saturday	April 4	9:00am-11:00am	rm 146 Posner
Saturday	April 11	9:00am-11:00am	rm 146 Posner
Saturday	April 18	9:00am-11:00am	rm 146 Posner
Sunday	April 26	9:00am-11:00am	rm 146 Posner

Course Summary

Week 1:

This lecture will start with an overview of the course structure, how the lectures will be organized and how 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 program 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 6:

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.

Week 7:

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.