Department of Resource Economics
University of Massachusetts
Fall 2019

Res Ec 797A: Special Topics in Forecasting (The name of the course will be changed to Applied Univariate and Econometric Time-Series Techniques)

Instructor: Bernard J. Morzuch (Bernie)
221 Stockbridge Hall
Phone: 413-545-5718
email: morzuch@resecon.umass.edu

Office Hours: As you need me or by appointment

Objectives:

I will do the following:

(1) Summarize important time-series methodologies and quantitative forecasting techniques in use and the issues surrounding their use;
(2) Discuss in detail some of the most important (though not necessarily the most widely used) techniques;
(3) Develop the theory necessary to understand the techniques used.

This is not an econometric theory, statistical theory, or forecasting theory course. Forecasting is a pragmatic art. My goal is to have the course reflect this pragmatism.

Readings:

Applied Econometric Time Series by Walter Enders, Second Edition (Wiley), 2004. ISBN 0-471-23065-0.

. This is an old edition. It is very similar to the third and fourth editions.

Introductory Business and Economic Forecasting by Paul Newbold and Theodore Bos. Second Edition (South-Western), 1994. ISBN 0-538-82874-9.

. This text is ideal for presenting and explaining univariate time-series techniques at their most rudimentary level. It is a great background text.
. Reading material is available in Moodle.

Prerequisites:

A graduate-level course in regression, knowledge of matrix algebra, and an intermediate-level ability in statistics are necessary.

Assignments:

There will be 10 or 11 assignments, each spaced approximately one week apart from the next. These must be done in a timely fashion. The purpose of each assignment is to ensure that you are mastering and digesting the topics as they are being covered. Assignments are worth 50% of your final grade.

Exams:

There will be two exams: a midterm (scheduled for Thursday, November 7 from 6:00 to 8:00 p.m.) and a final exam (scheduled for Thursday, December 19 from 10:30 a.m. to 12:30 p.m.). Our last class is Tuesday, December 10. I'd like to change the final to an earlier date.

Each exam is worth 25% of your final grade. The final exam will concentrate on material covered since the midterm exam.

Course Organization:

The material will be presented in a lecture format. Implementation of the appropriate computer software will be explained as we proceed through the lectures. You will not be required to make in-class presentations nor will you have to do a paper for this course.

Computer Software:

Students use spreadsheets for introductory and basic univariate techniques. I use SAS (and alternatively R) for the advanced univariate time-series techniques. I use SAS and Stata for vector autoregression models (VARMs) and vector error-correction models (VECMs). I have not yet graduated to using R for econometric systems of equations.

Course Outline: A Brief Overview

I will begin with background material on forecasting. This includes terminology, followed by simple and common univariate estimation procedures found in Newbold and Bos but not in Enders. We then move to several not-so-simple univariate methods. All of this paves the way for exploring the behavior of univariate time series and, in particular, testing for their stationarity. This semester we have 26 class periods. Approximately half of the course (14 lectures) will be devoted to this background material and to univariate techniques.

For the remaining 12 lectures, we will be covering the topics in Enders presented near the end of this syllabus. When you pick up the text and page through these readings, you may find them a bit frightening. I will attempt to lessen the difficulty by deciphering these pages for you. Enders is close to state-of-the-art regarding time-series econometrics.

During the last lecture, I would like to address what happens to forecasts when we combine forecasts from different methods. We complete the course by comparing the forecasting performance of all of the techniques that we've studied through the semester.

Course Outline: An Overview That Links The Topics

I spend approximately 14 lectures on univariate techniques (simple and advanced) and 12 on econometric techniques. I consider univariate exponential-smoothing techniques to be simple. "Simple" is not to imply that they are undeserving of a graduate-level course. These simple techniques provide a framework for exposing and considering how to handle things like stochastic trend and stochastic seasonality. (These two items are a nemesis in econometric time-series methodology, and the simple univariate methodologies provide background on how they come about and how to think about treating them.) I put Box-Jenkins in the category of advanced univariate methodology. Box-Jenkins (BJ) methodology requires an understanding of difference equations. We spend time on difference equations. The BJ methodology is clumsy and effectively regarded as ancient. However, it is the premiere body of knowledge that rigorously treats the univariate behavior of a variable by way of its moments and their application to the variable's autocorrelation function. When this is done, the meaning of a unit-root process becomes much clearer than if time is not devoted to such coverage. (The ultimate forecasting equation in any BJ specification is effectively a version of a VAR, which comes in the econometrics portion of the course.) Before getting to VAR, we consider the different ways to test for a unit root. Is a process a pure unit root process? Is it a unit root process with drift? Is it a unit root process with drift and trend? What happens to the behavior of the process if it is augmented with lagged values of itself? All of these possibilities are considered when testing if a process might be difference-stationary or trend-stationary. So, up to now, we are focusing strictly on different univariate specifications and ways to test for lack of stationarity, i.e., for unit roots.

Here comes VAR, which is nothing more than a reduced form. It can have any number of augmentations, let's just say "a lot." The arguments in the equation are specified in levels. Although you have evidence from the paragraph above regarding whether or not the variable is stationary (and particularly its order of integration), you disregard this information at this point. An additional, different variable with its own augmentations is introduced into the original VAR. (At this juncture, subject this variable to the same univariate tests and obtain similar information as the variable above.) Switch roles of the two variables in the previous equation to come up with a new (and different) equation, noting that it also is a VAR. Incorporate both equations into one model. The result is a vector autoregression model (VARM) with a lot of augmentations. Call this augmented VARM the unrestricted form. Whittle down the number of augmentations using the procedure suggested by Hsiao (FPE: Final Prediction Error) or the AIC or the SBC. The point is that you are testing an unrestricted VARM against a restricted VARM.

We now apply the logic of the VARM to the development of a vector error-correction model (VECM). Make use of the information that you have from above about the order of integration of each variable. To proceed, they must be integrated of the same order. (If not, you can't do the tests that follow.) Regress the first variable above (in its levels) on the second variable above (in its levels). This is called the cointegrating regression. Retrieve the estimated residuals. Test these residuals for stationarity. If they are stationary, the two variables are cointegrated. Cointegration turns out to be a long-run equilibrium restriction. Now, with a bit of algebra, the variables in the equations in the paragraph above are transformed to first differences. Doing so will result in the error term from the cointegrating regression being an argument in the equation. It is a mechanism for making short-run corrections. It is an error-correcting (EC) mechanism. It is the "EC" contained within "VECM." Remember, this information is a restriction. So, after doing all of this, the model itself becomes another form of a restricted VAR; i.e., a VECM is a restricted VAR.

Individual Topics:

Background Issues and Vocabulary; Basic Forecasting Methods
- Components of a time series: level, trend, seasonality, white noise
- Forecasting vocabulary: univariate; multivariate; data generating process (DGP); random walk
- Performance measures: MSE, MAE, MAPE, and Theil's U
- Forecasting accuracy measures that penalize for increased numbers of parameters
- Simple forecasting methods: naive no-change; naive change; naive seasonal change
- No updating versus updating when making forecasts in the post-sample period
- Regression-based forecasting models
- The simple moving average and seasonality
- The basics of decomposing a series; seasonal adjustment; the seasonal index method
- Getting within-sample and post-sample forecasts when using the decomposition method
- Exponential smoothing methods: simple; Holt; Holt-Winters

Methodology for Discovering the Data Generating Process of a Series
- Difference equations
- First-order difference equations
- General, homogeneous, and particular solutions
- Stability and equilibrium conditions
- Second-order difference equations
- Stability and stationarity
- Theoretical behavior of time-series models
- Autocovariances and autocorrelations for autoregressive and moving average processes
- The theoretical autocorrelation function (ACF) and the partial autocorrelation function (PACF)
- Using the ACF and PACF to make a decision about the data generating process

Box-Jenkins Approach
- Box-Jenkin (ARIMA) Modeling: Identification, Estimation, Diagnostic Testing, Forecasting
- Testing for the stationarity of a *non-seasonal* series and a *seasonal* series
- Using the Q-statistic to test if a series is white noise
- Analyzing the ACF and PACF of a series as a precursor to model development
- Estimating a model; analyzing the model's residuals
- Writing an ARIMA model in backshift operator notation
- Writing an ARIMA model in a form used for forecasting

The Econometric Approach
- The unit root problem
- Spurious regressions: background for unit root tests; Granger and Newbold's finding
- Unit root tests: the appropriate hypotheses and appropriate test statistics
- Model procedure when testing for unit roots
- Dickey-Fuller tests; augmented Dickey-Fuller tests
- Choosing deterministic regressors in augmented models; testing for white noise
- Vector autoregression (VAR): unrestricted and restricted VAR models
- Hsiao's approach for imposing restrictions: final prediction error (FPE)
- Granger causality

- Cointegration: what it is; how to test for it; its potential value
- Error Correction; setting up an error correction model (ECM); potential benefits of an ECM
- Error correction as another form of restricted VAR
- Engle-Granger methodology
- Autoregressive Conditional Heteroscedasticity (ARCH)

Putting Everything Together
- Combining forecasts; how to put together a combined forecast; benefits of combining

Chapter Coverage in Enders:
- Difference equations: Chapter 1
- Stationary time-series models: Chapter 2
- Autoregressive conditional heteroscedasticity (ARCH): Chapter 3, pp. 112-145
- Testing for trends and unit roots: Chapter 4
  - Unit root processes, pp. 156-181
  - Unit root tests, pp. 181-229
- Multiequation time-series models: Chapter 5
  - Intervention analysis, pp. 240-247
  - Transfer function models, pp. 247-264
  - Vector autoregression, pp. 264-283
  - Granger causality, pp. 283-290
- Cointegration and error-correction models: Chapter 6
  - Cointegration, pp. 319-328
  - Error-correction models, pp. 328-339

Additional Texts:

There are several additional texts out there that I like. I am mentioning two favorites directly below and why I like them. Periodically, we will be addressing material presented in these texts, and I will make this available to you.

Franses, P. H., D. van Dijk, and A. Opschoor. Time Series Models for Business and Economic Forecasting, 2e. Cambridge University Press, 2014.

In terms of level of difficulty, this text is between Newbold and Bos' text and Enders' text, leaning more in the direction of Enders. It does a nice job explaining difficult topics without getting bogged down in cumbersome mathematics.

Choi, In. Almost All About Unit Roots: Foundations, Developments, and Applications. Cambridge University Press, 2015.

There has been an explosion in the literature on non-stationarity and unit roots during the past 30 years. This book takes the approach of discussing as many papers as possible in presenting developments in the literature on unit roots. It is meant for someone with background in this area. It is current, difficult, and complete!