ECO 773 Business Forecasting

SPRING 2025 SYLLABUS

Professor:	Dr. Emir Malikov
Class Time:	Tuesday/Thursday @ 5:30—6:45pm
Class Location:	BEH 120
Office Hours:	every Wednesday @ 3—5pm, or by appt
Office Location:	BEH 509
Email:	emir.malikov@unlv.edu

COURSE DESCRIPTION

The varying fortunes of forecasters arise because good forecasts can seem almost magical, while bad forecasts may be dangerous. Forecasting is obviously a difficult activity, and businesses that do it well have a big advantage over those whose forecasts fail. In this course, we will explore model-building techniques for producing forecasts using ***time-series predictive econometrics*** that do not rely much on economics but are, fundamentally, data-driven.

We will first start with rigorous econometric theories behind the forecasting techniques. You will learn how to identify, build, estimate, diagnose and select the best forecasting model to construct optimal out-of-sample forecasts. Whereas our main focus will be on seasonal autoregressive integrated moving-average (SARIMA) models, the course will also cover time-series smoothing, decomposition, and volatility forecasting approaches for univariate—and multivariate, if time permits—predictive analysis. You will also learn how to implement these techniques in *R*.

Throughout the course, the emphasis will be on *understanding* the methodologies, including their underlying assumptions and statistical properties. The latter is particularly important in an age where generative AI excels in complex data analytics. Your understanding of what happens under the hood of econometric models will provide a critical human context to the interpretation of data forecasts. By the end of the semester, you will be equipped with essential forecasting techniques and the working knowledge about how to operationalize them in practice. The ultimate goal of this course is to help you develop a solid understanding of how to extract information from historical data and extrapolate it into the future.

COURSE PREREQUISITE(S)

A successful completion of ECO 770. ECO 740 or equivalent is also expected.

TEXTBOOK & READINGS

Throughout the semester, I will be providing you with lecture slides. They are an amalgam of the material from multiple sources, including some of my own, and are the bare minimum reading that is expected of you. It is imperative that you supplement it with reading the textbook as well.

Required Textbook

"*Time Series Models for Business and Economic Forecasting*," 2nd Ed., by Frances, P.H., van Dijk, D. and Opschoor, A. [FVDO]. An e-textbook is available for rent through Amazon for about $40.

HIGHLY Recommended Supplementary Texts

1. *"Forecasting in Economics, Business, Finance and Beyond*," Edition 2017, by Diebold, F.X., accessible FOR FREE on the author's website. This textbook is less dense than the main textbook.
2. *"Forecasting: Principles and Practice*," 3rd Ed., by Hyndman, R.J. and Athanasopoulos, G., accessible FOR FREE online. This textbook is more practice-oriented, has short lecture videos and *R* code snippets.

PREREQUISITE KNOWLEDGE

A solid ***knowledge of algebra, (at least univariate) calculus, basic matrix algebra, fundamentals of math statistics and probability theory, and ECO 770 is expected***. It is your responsibility to brush up your math/stats skills outside the classroom, should there be such a need. At the very least, be comfortable with the material in Appx. A—D and Ch. 2—6 of the Baby Wooldridge.

STATISTICAL SOFTWARE

The required statistical software for this course is *R* or, to be precise, the *R-Studio* for its better interface. *R* is a language and environment for statistical computing and graphics, available as a free open-source software. It is arguably "the" language of data science and statistics (e.g., Google do their analytics using *R*). You will use it for both the computational problems in the homework assignments as well as the forecasting project. The software is free and can be easily downloaded from the web.

When it comes to learning and mastering *R*, practice is the key. **Learning by doing is the way** here. In fact, I would argue it is the only way, which is why learning to code in *R* will be your own responsibility. I will not "teach" it in class: after all, ours is not a programming course. Having said that, there is an immense amount of tutorials, books, videos and free tutorials available online, so Google or any one of the many generative AIs available online is your best friend here. Also, always remember that virtually any questions you will ever have about coding in *R* had already been asked by someone else and answered on numerous *R*-help forums. Just google it or try ChatGPT or Bard or Gemini or any other AI chatbot! In my own experience, the most effective way to learn *R* is by tinkering with *R* on one's own.

To help you get started with *R*, the first couple of homework assignments will have you complete several **interactive courses on *R* coding**. Each enrolled student will be provided free access to these courses **through DataCamp Learn** until the end of the semester. While select courses will be mandatorily assigned for homework credit, you are welcome—in fact, encouraged—to complete additional data courses available on the platform for your own learning and professional development while you have free access. ***For more details, refer to the "Resources for Learning & Mastering R" page on Canvas.***

Over the course of the semester, I will also be providing you with codes pertaining to the forecasting models that we will be learning in class.

STUDY TIPS

To do well in the course, you MUST study outside the classroom. **This is not an easy class, so be prepared to be intellectually challenged on a continuous basis and put in hours of hard work**. First of all, work through lecture notes/materials and the textbook. Read them multiple times; read them with a pencil in hand. Second, work through the examples contained in each chapter as well as end-of-chapter problems. More generally, I urge you to study on a continuing basis: it will be virtually impossible for you to do well in this class by pulling a few all-nighters before the exam or before the homework due date.

GRADING

Your course grade consists of the following components:

GRADE COMPONENTS

Component	Weight
Homework	30%
Exams	25% each
Forecasting Project	20%

At the end of the course, your score on each of the above components will be converted to a 100-point based scale. Your final score in the class will be a weighted average of these points. I will then map your final score (also on a 100-point based scale) to the grade you will receive in the course using the following table.

GRADE SCALE

Grade	Final Score
A	≥ 94
A–	≥ 90 but < 94
B+	≥ 86 but < 90
B	≥ 80 but < 86
B–	≥ 75 but < 80
C+	≥ 71 but < 75
C	≥ 65 but < 71
C–	≥ 60 but < 65
D	≥ 50 but < 60
F	< 50

When computing the final score, I **round it up** to the nearest whole number. That is, 85.56 is rounded up to 86, and so is 85.01.

Though unlikely, depending on performance of the class, I may adjust this table which you will be notified about. No adjustment will make a student worse off.

Requests to revise grades on exams and/or assignments are accepted in WRITING only. Please compose a brief letter argumentatively explaining (point by point) why you think your grade should be changed, then turn in this letter along with the original assignment back to me. You may contest your grade within one week from the date your graded assignments was returned to you.

All grades will be recorded and continually updated in the *Canvas* gradebook. If you think that one of your grades is missing, you must bring it to my attention within a week of the assignment's due date. I will not chase your missing grades if the assignment was past due more than a week. Therefore, I suggest you check on your grades on a regular basis.

How to Understand the "Total" score in the Canvas Gradebook

You can monitor your performance in the course throughout the semester by looking at your "Total" score in the *Canvas* gradebook. Although, it is crucial that you understand how it is being computed and why and how it changes over the course of the semester. Here are its key features.

The "Total" score in the *Canvas* gradebook is your "as-of-now" course grade that the platform continually computes for you. This as-of-now grade is "dynamic" and it constantly gets updated as more graded deliverables are assigned and completed with time.

This tentative course grade might go up or down every time you complete an assignment or take a test because, when computing the "Total" score, the *Canvas* gradebook assumes that all incomplete and/or un-attempted assignments/exams do not exist and effectively reassigns their weights to the past assignments/exams you have already completed. This means that your past assignments/exams will always be given a larger weight than they will be eventually when your truly final total score is computed at the end of the semester.

Every time you complete a new deliverable, by construction, your updated "Total" score will be higher than it was previously *only if* your score on the latest assignment is greater than the weighted average of your scores on previous assignments. For concreteness, consider the following example. There is a total of three assignments in the course: test A, test B and test C with the corresponding weights being 25%, 35% and 40%, respectively. Say, you have only taken the first two tests so far and scored 62/100 on test A and 80/100 on test B. The *Canvas* gradebook will compute your as-of-now "Total" score to be $62\left(\frac{25}{25+35}\right)+80\left(\frac{35}{25+35}\right)=72.5$. If you score exactly 72.5 on test C, then your final "Total" score will remain what it had been before: $62\left(\frac{25}{25+35+40}\right)+80\left(\frac{35}{25+35+40}\right)+72.5\left(\frac{40}{25+35+40}\right)=72.5$. If you score more/less than 72.5 on test C, then your new "Total" score will also become greater/less than 72.5. For example, if you get 75 on test C, then $62\left(\frac{25}{25+35+40}\right)+80\left(\frac{35}{25+35+40}\right)+75\left(\frac{40}{25+35+40}\right)=73.5>72.5$; or if you get 70 on test C, then $62\left(\frac{25}{25+35+40}\right)+80\left(\frac{35}{25+35+40}\right)+70\left(\frac{40}{25+35+40}\right)=71.5<72.5$.

DELIVERABLES

Homework/Problem Sets

You should expect about 8—10 homework assignments this semester. The exact due dates will be determined by the flow of the course. Problem sets will include analytical material as well as applications of forecasting methods. **They will take a substantial amount of time, so plan ahead!**

Each question (or a sub-question within a problem) will be graded on a scale of 4 (all correct), 3 (mostly correct), 2 (mostly incorrect), 1 (all incorrect) and 0 (left unattempted).

Homework submissions **must be properly scanned as high-resolution PDF** and turned in through *Canvas*. If handwritten, it is imperative that it is legible for, otherwise, I reserve the right to require that it be typeset, including all equations. Albeit handwritten solutions are allowed on default, you are encouraged to typeset them using *Overleaf/LaTex*.

All your answers must be written as complete thoughts in well-written sentences. "Yes" or "no" are not acceptable answers and will be given a grade of zero. Needless to say, all equations must be mathematically sound.

Exams

Expect 2 in-class paper-and-pencil exams (a midterm and a "near-final"). The date of a midterm exam is to be determined depending on the flow of the class; I will announce it at least a week in advance. The second exam will take place during the second-to-last week of the semester.

Forecasting Project

You will work on a forecasting project in which you will use the data (provided by me) to perform preliminary statistical assessment, develop a forecasting model, evaluate forecast performance, and write up your results. This project is individual and, essentially, a take-home final. I will provide you with both the data and concrete guidelines 10 days before the due date. The writeup of the project, along with the *R* script, is due Thursday, May 15 @ 10:00am. No extensions.

CLASS PARTICIPATION

Participation is important for this class and, in fact, is graded, so ***you should attend class*** *unless you are ill, in which case please stay home.* The success of the course depends heavily upon the contributions of each student: it is you who makes the class more fun! I expect you to come to class prepared and willing to contribute to the class discussion. ***If you must miss class, please get notes from another student.*** If any questions remain, please attend office hours.

EXTRA CREDIT

I do NOT allow individual students to do any additional work for extra credit. All students are advised to consult with me throughout the semester to avoid end-of-the-semester "surprises."

OFFICE HOURS

I will hold weekly office hours on Wednesdays between 3 and 5pm (except for the school breaks and observed holidays). If you cannot make it to see me during my regular office hours, please email me for an appointment.

Since access to upper floors of BEH is still restricted, you will need to let me know beforehand via email of your intention to come see me so that I plan to come down (or send someone) to the first floor to get you.

EMAILS

If you cannot see me during my office, feel free to email me with any questions: I will be happy to help. I should be able to reply within one business day. Having said that, please keep in mind that it is nearly impossible to meaningfully answer conceptual or debugging questions via email. In such instances, instead I will ask you to set an appointment with me or see me during my next office hours in order to get your question answered/explained.

CLASSROOM ETIQUETTE

We are all adults, so please be considerate to myself and your fellow classmates:

- Please do not be late for class.
- Please no chatter amongst yourselves. It is distracting, and I will call you out on it.
- Put your cell phones on silent and put them away before the class starts. This also implies no texting, being on social media or surfing web for unrelated stuff on your laptops/tablets during the lecture. If I see an inappropriate use of technology during class, I will call you out publicly.
- The use of laptops/tablets is permitted for course-related purposes only.
- Please wait until the class ends before gathering your belongings.
- If you need to leave early, please sit near the door.

The bottom line is, if you do not want to participate and pay attention in class, then do not come to class.

ANNOUNCEMENTS

All announcements (including those about assignments), supplementary course materials and other relevant information will be posted on the *Canvas* course page. ***It is your responsibility to check for announcements daily.*** Your scores for homework assignments and exams will also be posted on *Canvas*.

UNLV ACADEMIC POLICIES

Please refer to Academic Policies for Students.

COURSE OUTLINE

This course outline is *tentative* and subject to change as the semester progresses.

INTRODUCTION (FVDO Chapters 1-2)

- Forecasting with Time Series Econometrics
- Universal Considerations for Any Forecasting Task
- Features of Time Series Variables
- Logs and the Growth Rates

CYCLES: FUNDAMENTALS & IDENTIFICATION (FVDO Chapter 3)

- White Noise
- Covariance Stationarity
- The Lag Operator
- *AR*(*p*) Processes
  - Stationary *AR*(*p*) Processes
  - The *AR*(*p*) to *MA*(∞) Inversion
  - Differencing to Achieve Stationarity
- The World Representation
  - *ARMA*(*p*,*q*) Processes
- *MA*(*q*) Processes
  - Stationary *MA*(*q*) Processes
  - The *MA*(*q*) to *AR*(∞) Inversion
- Processes with Non-Zero Means
- (Partial) Autocorrelations
  - ACF and PACF
  - Identification of *ARMA*(*p*,*q*) Models via Autocorrelations
  - Empirical ACF and PACF; the Q-statistic

CYCLES: ESTIMATION, DIAGNOSTICS & MODEL SELECTION (FVDO Chapter 3)

- Estimation
  - Estimation of *AR*(*p*) Processes
  - *E*stimation of *ARMA*(*p*,*q*) Processes
- Model Diagnostics
  - Testing for Residual White Noise
  - Testing for Normality
- Model Selection using Information Criteria

CYCLES: FORECASTING WITH *ARMA* MODELS (FVDO Chapter 3)

- Optimality of the Conditional Mean Predictor
- Optimal *ARMA* Forecasts
  - Optimal Forecasts for *MA*(*q*)
  - Optimal Forecasts for *AR*(*p*)
  - Optimal Forecasts for *ARMA*(*p*,*q*)
- Point Forecast Evaluation

  - Absolute Standards for Point Forecasts
  - Relative Standards for Point Forecasts; the Diebold-Mariano Test

TIME SERIES DECOMPOSITION & SMOOTHING (Lecture Notes)

- Non-*ARMA* Forecasters
- Exponential Smoothing
- Time Series Decompositions
  - Moving Average Smoothing
  - Classical Decomposition

TRENDS (FVDO Chapter 4)

- Modeling Trends
  - Deterministic Trend
  - Stochastic Trend
- Unit Root Tests
  - The ADF (Single) Unit Root
  - Extending ADF to Include Deterministic Components
  - Extending ADF to Two Unit Roots

SEASONALITY (FVDO Chapter 5)

- Modeling Seasonality
  - Deterministic Seasonality
  - Stochastic Seasonality
  - Seasonal *ARIMA*(*p*,*d*,*q*)(*P*,*D*,*Q*) Processes
- Seasonal Unit Root Tests
  - The HEGY Tests for Seasonal Unit Roots
  - Extending ADF to Include Deterministic Components
  - Extending ADF to Two Unit Roots

NOISE: CONDITIONAL VOLATILITY DYNAMICS (FVDO Chapter 7)

- Conditional Heteroskedasticity
- Modeling Conditional Variance
  - *ARCH*(*q*) Processes
  - *GARCH*(*p*,*q*) Processes
  - Selected Extensions: *GARCH-M* and *TGARCH*
- Pretesting, Estimation & Diagnostics of *GARCH*
  - Testing for Conditional Volatility
  - Estimation of *GARCH* Models
  - Diagnostics of *GARCH* Models
- Forecasting with *GARCH*

NONLINEAR & MULTIVARIATE MODELS (FVDO Chapters 8-9)

- *TAR* and *STAR* Models
- *VAR* Models