Advanced Time Series Methods
Course Overview and Teaching Plan

Jiajing Sun

School of Economics and Management, University of Chinese Academy of Sciences

Spring 2026

Why this course?

- Many economic and financial variables are observed sequentially: GDP, inflation, interest rates, returns, exchange rates, volatility, and risk measures.
- The objective is not only description, but also **dynamic interpretation**, **forecasting**, **policy analysis**, and **risk measurement**.
- Time ordering creates dependence, and dependence changes both estimation and inference.

Big picture

This course studies how to model, estimate, interpret, and forecast dependent data in a rigorous econometric framework.

What kinds of questions will we study?

- Is a series stable, drifting, or driven by stochastic trends?
- How persistent are shocks, and how quickly do they die out?
- When do variables move together in the long run?
- How should we model time-varying volatility and risk?
- How do we do reliable inference when data are dependent?
- How do filtering and state-space methods extract latent signals?

Two visual motivations

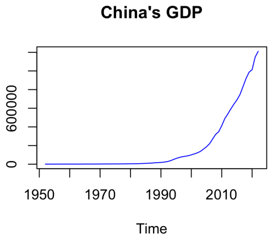

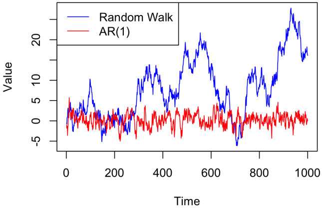

Real macroeconomic series often contain trend, persistence, and structural change.

Stable dynamics and shock-accumulating dynamics can look very different and require different tools.

Course Roadmap

Core textbook and scope

- We use the course textbook as the main reference.
- The core material in this lecture series covers **Chapters 2–9**.
- **Chapter 10 (Machine Learning)** is treated as an optional advanced topic rather than a core block.
- R applications are integrated across the course rather than taught as a separate module.

Core emphasis

Univariate foundations $\rightarrow$ multivariate systems $\rightarrow$ volatility $\rightarrow$ nonparametrics $\rightarrow$ robust inference $\rightarrow$ filtering $\rightarrow$ nonstationarity $\rightarrow$ continuous-time finance.

Course structure in hours

- **Total:** 60 contact hours.
- **Core teaching:** 54 hours.
- **Q&A / tutorial / assignment discussion:** 3 hours.
- **Final exam:** 3 hours.

Class meeting format

Each meeting is 3 contact hours, and each contact hour is 45 minutes.

What you should be able to do by the end

By the end of the course, students should be able to:

1. model univariate and multivariate dependence structures;
2. distinguish stationary, nonstationary, and cointegrated dynamics;
3. analyze volatility and time-varying uncertainty;
4. use robust inference tools for dependent data;
5. understand filtering, state-space, and continuous-time perspectives;
6. connect theory, computation, and empirical interpretation using R.

How the course is organized

Foundations and univariate models $\rightarrow$ multivariate systems $\rightarrow$ volatility and nonparametrics $\rightarrow$ robust inference and filtering $\rightarrow$ nonstationarity and continuous-time topics

- We begin with the language of stationary dependence and ARMA modeling.
- We then expand to VAR, VECM, SVAR, volatility, filtering, and continuous-time intuition.

Course Roadmap

Chapter allocation in the 54 teaching hours

Chapter	**Theme**	**Hours**
Chapter 2	Univariate time series	8
Chapter 3	Multivariate linear time series	10
Chapter 4	Volatility models	6
Chapter 5	Nonparametric methods	4
Chapter 6	HAR inference / robust inference	9
Chapter 7	Filtering	5
Chapter 8	Nonstationary processes	7
Chapter 9	Continuous-time finance	5

What this allocation means in practice

- **Most time** goes to the core econometric backbone: univariate models, multivariate dynamics, and robust inference under dependence.
- **Significant time** is also reserved for nonstationary processes, because trends, unit roots, and long-run relations are central in applied work.
- **Specialized but important blocks** include volatility, filtering, and continuous-time finance.

Course Roadmap

20-lecture map (Lectures 1–10)

Lec.	Main topic	Core chapter
1	Course map; dependence, stationarity, ergodicity, mixing, and Wold decomposition	Ch. 2
2	AR(p), MA(q), ARMA(p, q); lag polynomials; roots; causality; invertibility	Ch. 2
3	ACF, PACF, model identification, estimation, diagnostics, and forecasting	Ch. 2
4	VAR and VMA representations; covariance-stationary multivariate models	Ch. 3
5	VAR estimation, stability, forecasting, and interpretation	Ch. 3
6	Impulse responses, FEVD, and structural interpretation	Ch. 3
7	Cointegration, VECM, and long-run multivariate dynamics	Ch. 3 / 8
8	Volatility basics: ARCH, GARCH, and conditional variance dynamics	Ch. 4
9	Asymmetric / multivariate volatility and forecasting risk	Ch. 4
10	Nonparametric methods for time series	Ch. 5

20-lecture map (Lectures 11–20)

Lec.	Main topic	Core chapter
11	HAC, long-run variance estimation, and robust inference under dependence	Ch. 6
12	Bootstrap and advanced inference topics for time-series data	Ch. 6
13	Further robust inference applications and empirical implementation	Ch. 6
14	Linear filters and signal extraction	Ch. 7
15	State-space models and the Kalman filter	Ch. 7
16	Trends, random walks, unit roots, and difference stationarity	Ch. 8
17	Nonstationary systems, long-run interpretation, and empirical issues	Ch. 8
18	Continuous-time finance and diffusion-based intuition	Ch. 9
19	Q&A / tutorial / assignment discussion / review	Review
20	Final exam	Exam

How to read the lecture map

- Lectures 1–3 build the univariate foundation and the ARMA toolkit.
- Lectures 4–7 move to multivariate systems and long-run relations.
- Lectures 8–10 introduce specialized modeling blocks: volatility and nonparametrics.
- Lectures 11–15 focus on inference, filtering, and signal extraction.
- Lectures 16–18 address nonstationarity and continuous-time ideas.

Course Roadmap

How R is used in this course

- R is integrated throughout the course rather than separated into a standalone programming week.
- In each major block, we will use R for:
  - simulation,
  - visualization,
  - model estimation,
  - diagnostics,
  - interpretation of empirical output.
- The goal is not coding for its own sake, but econometric understanding through implementation.

What I expect from you

- This is a course for **master's and PhD students**, so we will combine intuition with mathematical structure.
- I expect you to understand definitions, derivations, and empirical meaning.
- You should be able to explain major concepts in words, not only reproduce formulas.
- Regular review after each lecture will matter much more than last-minute memorization.

A good weekly routine for this course

1. Read the relevant textbook section before class.
2. During class, focus on the conceptual logic behind the formulas.
3. After class, rewrite the key definitions and derivations in your own words.
4. Reproduce the main R examples and change parameters to build intuition.
5. Bring questions to the tutorial / review session instead of letting confusion accumulate.

Why the sequence is designed this way

- We start with stationary linear models because they are the basic language of time-series econometrics.
- We then generalize to multivariate systems and structural interpretation.
- After that, we add volatility, nonparametrics, robust inference, filtering, and nonstationarity.
- The final continuous-time block broadens the perspective rather than replacing the discrete-time core.

Course Roadmap

The series in one sentence

This lecture series moves from the foundations of dependent-data modeling to a broad toolkit for analyzing persistence, volatility, signal extraction, nonstationarity, and dynamic economic structure.

Welcome to the course

We begin with the foundations of univariate time-series analysis, and then build step by step toward a broader econometric view of dynamic systems.