Financial Econometrics 2026

Course Topics & Subtopics

Based on Stock & Watson, *Introduction to Econometrics*, 4th ed. · SJTU Summer School · Dr. Fatih Kansoy (`fatih.ai`)

Ten lectures, each with a Stata and a Python lab. This is the syllabus from which the **3 quizzes, the midterm, and the final** are drawn. The hierarchy mirrors a textbook contents page — **Lecture → numbered section → subtopics** — with the load-bearing estimators, tests, and real-world cases named so each exam item maps to a topic. *Notation:* β population coefficients, $\widehat{\beta}$ estimates; u the error, $\widehat{u}$ the residual; $\mathbb{E}(Y \mid X)$ the conditional mean; SE the standard error; R^2 the goodness of fit; t, F test statistics; ρ_k the autocorrelation at lag k.

The course at a glance — three blocks

Block	Lectures	The question
A. Foundations	L1–L2	The probability and statistics that underpin everything
B. The regression model	L3–L6	OLS, inference, multiple regressors, and the omitted-variable problem
C. Panel data & time series	L7–L10	Fixed effects, forecasting, dynamic causal effects, and financial time series

Cross-cutting questions asked of every estimate, and reused all term: is it *precise* (the standard error, the t-statistic, the confidence interval), and is it *causal* (the exogeneity condition $\mathbb{E}(u \mid X) = 0$)? Plus: reading regression output; heteroskedasticity-robust, clustered, and HAC standard errors; the t and F machinery; prediction versus causation; and doing every analysis twice, in **Stata** and in **Python**.

Block A — Foundations

Lecture 1 — Economic Questions and Probability (S&W Ch 1–2)

1.1 **What econometrics is for.** The four questions econometrics answers (does a policy work, forecasting, causal effects, testing theory); **correlation is not causation**; the **ideal randomised experiment** as the benchmark; why finance relies on **observational** data; prediction vs causation as the spine of the course.

1.2 **Random variables and distributions.** Discrete vs continuous; the probability distribution (pmf/pdf) and the **cumulative distribution**; a return as a random variable; the Bernoulli and the normal.

1.3 **Expectation, variance, and the shape of returns.** The **expected value** (mean) and the **variance** / standard deviation; **skewness** and **kurtosis**; mean as expected return, variance as risk; fat tails (a preview of L10).

1.4 **Two random variables.** Joint, marginal, and **conditional** distributions; the **conditional expectation** $\mathbb{E}(Y \mid X)$ — the object regression estimates; **covariance and correlation**; statistical independence; the conditional-mean-zero idea.

1.5 **The normal and its relatives.** The normal distribution and **standardisation**; the 68–95–99.7 rule; the **Student** t, **chi-squared**, and F distributions and where each appears later.

1.6 **Random sampling and the sampling distribution. I.i.d.** sampling; the sample average $\bar{Y}$ as a random variable; the **law of large numbers**; the **central limit theorem** (why t-statistics are normal in large samples).

Lecture 2 — Review of Statistics (S&W Ch 3)

2.1 **Estimation.** Estimator vs estimate; the sample mean as an estimator of μ_Y; the three properties — **unbiasedness**, **consistency**, **efficiency**; $\bar{Y}$ is **BLUE**; the standard error of the mean.

2.2 **Hypothesis testing.** Null vs alternative; the t**-statistic**; the p**-value**; the **significance level** and the rejection rule; **Type I vs Type II error**, size and power; one- vs two-sided tests; a real-world Type I/II example.

2.3 **Confidence intervals.** Construction $\widehat{\mu} \pm 1.96\,\text{SE}$; the correct interpretation (the procedure, not the parameter); the duality of tests and intervals.

2.4 **Comparing two means.** The difference-in-means estimator and its standard error; the two-sample t-test; differences-of-means as the seed of regression with a binary regressor (L4); statistical vs economic significance.

Block B — The Regression Model

Lecture 3 — Linear Regression with One Regressor (S&W Ch 4)

3.1 **The population regression line.** $Y_i = \beta_0 + \beta_1 X_i + u_i$; the intercept and **slope**; the error term u; the slope as $\Delta Y/\Delta X$ and as the conditional-mean response.

3.2 **The OLS estimator.** Minimising the **sum of squared residuals**; the normal equations and the **closed form** $\widehat{\beta}_1 = \widehat{\text{cov}}(X,Y)/\widehat{\text{var}}(X)$; **fitted values** and **residuals**; OLS by hand on a small sample.

3.3 **Measures of fit.** TSS = ESS + SSR; the R^2; the **standard error of the regression** (RMSE); what a high or low R^2 does and does not mean.

3.4 **The least-squares assumptions.** (1) $\mathbb{E}(u \mid X) = 0$ — the exogeneity / causal assumption; (2) i.i.d. draws; (3) large outliers are unlikely; what each one buys.

3.5 **The sampling distribution of OLS.** $\widehat{\beta}_1$ is unbiased and consistent; its variance and the role of var(X); the normal large-sample distribution; **prediction vs causation** with one regressor (the CAPM beta example).

Lecture 4 — Regression: Hypothesis Tests and Confidence Intervals (S&W Ch 5)

4.1 **Testing the slope.** The t**-statistic** $t = (\widehat{\beta}_1 - \beta_{1,0})/\text{SE}(\widehat{\beta}_1)$; the p-value; two-sided and one-sided tests; testing a non-zero null (e.g. a CAPM beta of one).

4.2 **Confidence intervals for a slope.** $\widehat{\beta}_1 \pm 1.96\,\text{SE}$; reporting and interpreting; the t vs normal distinction in small samples.

4.3 **A binary regressor.** Regression on a 0/1 dummy **is** the difference in means; reading the intercept and slope; the link back to L2.

4.4 **Heteroskedasticity and standard errors.** Homoskedastic vs **heteroskedastic** errors; why financial data are heteroskedastic; **heteroskedasticity-robust standard errors**, and why they are the default; the cost of getting the SE wrong.

4.5 **The Gauss–Markov theorem.** OLS is BLUE under homoskedasticity; what fails under heteroskedasticity; why we still use OLS with robust SEs.

Lecture 5 — Multiple Regression (S&W Ch 6)

5.1 **Omitted variable bias.** The **two conditions** for OVB (the omitted variable is a determinant of Y *and* correlated with X); the **sign of the bias**; bias = simple minus controlled slope; the canonical examples.

5.2 **The multiple regression model.** $Y = \beta_0 + \beta_1 X_1 + \cdots + \beta_k X_k + u$; the **ceteris paribus** (partial-effect) interpretation; holding other regressors fixed.

5.3 **Measures of fit.** R^2 vs **adjusted** R^2; why adding regressors never lowers R^2; the penalty in $\bar{R}^2$.

5.4 **The OLS assumptions in multiple regression.** The four assumptions; **conditional mean independence** and the role of control variables (a control need not be causal).

5.5 **Multicollinearity.** **Perfect** multicollinearity and the **dummy-variable trap**; **imperfect** multicollinearity and inflated standard errors; the variance inflation factor; collinearity is not bias.

Lecture 6 — Hypothesis Tests in Multiple Regression (S&W Ch 7)

6.1 **Tests on a single coefficient.** The t-test and CI for one β_j in the presence of others; size distortion from testing many coefficients one at a time.

6.2 **Joint hypotheses and the F-statistic.** Why a joint test is not the sum of single tests; the **F-statistic** for q restrictions; the F distribution and critical values; the heteroskedasticity-robust F.

6.3 **The overall F and restrictions across coefficients.** The overall significance of a regression; testing restrictions (e.g. $\beta_1 = \beta_2$); rewriting the regression to test a linear combination; **lincom**.

6.4 **Specification and confidence sets.** The confidence set / ellipse for two coefficients; **model specification** and sensitivity analysis; control variables vs the variable of interest.

Block C — Panel Data and Time Series

Lecture 7 — Panel Data (S&W Ch 10)

7.1 **Panel structure.** Entities observed over time; **between** vs **within** variation; the panel as a tool against omitted variables.

7.2 **Fixed effects.** The before-and-after (two-period) estimator; **entity fixed effects** in three equivalent forms (dummies, within-demeaning, first differences); **time fixed effects**; the **two-way** fixed-effects model.

7.3 **What fixed effects can and cannot do.** They absorb all *time-invariant* entity confounders; the **strict exogeneity** assumption and feedback; what FE cannot remove (time-varying confounders); fixed vs random effects.

7.4 **Clustered standard errors.** Serial correlation within an entity; **clustered standard errors** and why they can more than double the naive SE; the real Guns-panel result (FE plus clustering flips significance).

Lecture 8 — Time Series and Forecasting (S&W Ch 14)

8.1 **The language of time series.** **Lags**, **differences**, and **growth rates**; **autocorrelation** and the ACF; **stationarity**; why ordered, serially-correlated data break the i.i.d. tools.

8.2 **Autoregressive models.** The **AR(1)** and **AR(p)** models; **persistence** $\beta_1 = \rho_1$, the long-run mean, and the half-life; lag-length choice by **BIC vs AIC**; seasonality.

8.3 **Forecasting and its limits.** One-step forecasts; the **RMSFE** and forecast intervals; **in-sample fit vs pseudo out-of-sample** forecasting; why forecasts fail at turning points.

8.4 **The random walk and efficient markets.** The **random walk** and **unit roots**; **stock returns are unforecastable** (the EMH in the data); statistically significant but economically worthless; volatility clustering (a preview of L10); **spurious regression**; the **Dickey–Fuller** test.

8.5 **Predicting with other variables.** The **ADL** model; **Granger predictive power** (prediction, not causation); **HAC standard errors** for serially-correlated errors; the **term spread** as a leading indicator.

Lecture 9 — Estimation of Dynamic Causal Effects (S&W Ch 15)

9.1 **Distributed lags and multipliers.** The **dynamic causal effect** and the **impulse response**; the **distributed lag model**; **dynamic** multipliers (the impact and lagged effects); the **cumulative** multiplier and the **long-run** multiplier.

9.2 **HAC standard errors.** Why distributed-lag errors are serially correlated; the **correction factor** and **Newey–West** (Bartlett) weights; the truncation rule $m = 0.75\,T^{1/3}$; the **robust-vs-HAC decision** (distributed lag needs HAC, an ADL with lagged Y does not); **lincom** for the cumulative multiplier's SE.

9.3 **Exogeneity. Past-and-present** vs **strict** exogeneity; when the OLS multipliers are causal; the **multicollinearity paradox** (insignificant lags, a significant cumulative effect); exogeneity is an economic judgement, not a test — the oil-shock case.

Lecture 10 — Volatility, VAR, and Cointegration (S&W Ch 16)

10.1 **Volatility clustering and ARCH/GARCH.** Returns unforecastable in the *mean* but their *variance* is not (persistent squared returns); the **ARCH** model and the **GARCH(1,1)** recursion $\sigma_t^2 = \omega + \alpha u_{t-1}^2 + \beta\sigma_{t-1}^2$; **persistence** $\alpha + \beta$, the unconditional variance, and volatility forecasting; **Value-at-Risk**; the **leverage effect** (GJR).

10.2 **Vector autoregressions.** The **VAR** system; **Granger causality** inside a system; the parameter count $k(kp + 1)$; **multiperiod forecasts** (iterated vs direct); impulse responses; the GDP-growth / term-spread VAR.

10.3 **Unit roots and cointegration. Orders of integration** I(0)/I(1)/I(2); the **DF-GLS** unit-root test; **cointegration** (two I(1) series sharing a stochastic trend); the **Engle–Granger** two-step; the **error-correction** model; the term structure and pairs trading.

Appendices

Appendix A — Master list of estimators, tests & formulas (exam-targetable)

#	Lec.	Tool / formula
A1	L1	Mean $\mathbb{E}(Y)$, variance $\text{var}(Y)$, $\text{cov}(X,Y)$, corr = $\text{cov}/(\sigma_X\sigma_Y)$; conditional mean $\mathbb{E}(Y \mid X)$
A2	L1	Standardise $Z = (Y-\mu)/\sigma$; CLT: $\bar{Y} \approx N(\mu, \sigma^2/n)$
A3	L2	SE of the mean $s_Y/\sqrt{n}$; $t = (\widehat{\mu} - \mu_0)/\text{SE}$; CI $\widehat{\mu} \pm 1.96\,\text{SE}$
A4	L3	OLS $\widehat{\beta}_1 = \widehat{\text{cov}}(X,Y)/\widehat{\text{var}}(X)$, $\widehat{\beta}_0 = \bar{Y} - \widehat{\beta}_1\bar{X}$
A5	L3	TSS = ESS + SSR; R^2 = ESS/TSS; SER = $\sqrt{\text{SSR}/(n-2)}$
A6	L4	Slope test $t = (\widehat{\beta}_1 - \beta_{1,0})/\text{SE}(\widehat{\beta}_1)$; CI $\widehat{\beta}_1 \pm 1.96$ SE; robust SE
A7	L5	OVB = $\widehat{\beta}_1^{\text{short}} - \widehat{\beta}_1^{\text{long}}$; sign = $\text{sign}(\beta_{\text{omit}}) \cdot \text{sign}(\text{corr}(X, \text{omit}))$; $\bar{R}^2$
A8	L6	F-test for q restrictions; overall F; $F = t^2$ for $q = 1$
A9	L7	Fixed effects (demeaning / first differences); clustered SE
A10	L8	AR(1) $Y_t = \beta_0 + \beta_1 Y_{t-1} + u_t$; long-run mean $\beta_0/(1-\beta_1)$; half-life $\ln(0.5)/\ln\beta_1$; Dickey–Fuller
A11	L9	Distributed lag; cumulative multiplier $\sum_{j=0}^{h}\beta_{j+1}$; HAC $m = 0.75\,T^{1/3}$
A12	L10	GARCH(1,1) $\sigma_t^2 = \omega + \alpha u_{t-1}^2 + \beta\sigma_{t-1}^2$; persistence $\alpha+\beta$; unconditional var $\omega/(1-\alpha-\beta)$; 1-day 99% VaR = $2.33\,\sigma_t$
A13	L10	VAR parameter count $k(kp+1)$; cointegration $z_t = Y_t - \theta X_t \sim$ I(0)

Appendix B — Real-world cases, data & labs (by lecture)

L1 . The distribution of stock returns; expected return vs risk; correlation vs causation.

L2 . Testing whether a strategy beats the market; Type I/II errors in a trading rule.

L3 . OLS by hand; the CAPM beta; earnings and a single regressor (the CPS / Growth data).

L4 . Robust vs homoskedastic SEs; a beta significantly different from one; the Growth–tradeshare slope.

L5 . Omitted-variable bias and the gender pay gap with controls; the size/value factors.

L6 . Joint significance of risk factors; the overall F; restrictions across coefficients.

L7 . The Guns state panel: a shall-issue effect that survives or vanishes under FE and clustering.

L8 . Inflation is forecastable, returns are not (the EMH); the yield-curve recession signal.

L9 . Oil-price shocks → inflation (fast, precise) and → output (slow, endogenous); HAC inference.

L10 . GARCH on the daily S&P 500 (the 2020 COVID crash, VaR); a VAR of GDP growth and the spread; the cointegrated 10-year and 3-month Treasury rates.

Appendix C — Cross-cutting skills the exams test

C1. Reading regression output — coefficients, standard errors, t-statistics, p-values, R^2, and what each column means.

C2. Inference — computing and interpreting a t-statistic, a p-value, and a confidence interval; the size of a test; an F-test of a joint hypothesis.

C3. Prediction versus causation — distinguishing a forecast from a causal effect, and naming the assumption ($\mathbb{E}(u \mid X) = 0$) a causal claim requires.

C4. Omitted-variable bias — diagnosing it, signing it, and reading bias as the change in a slope when a control is added.

C5. The right standard error — choosing among homoskedasticity-robust, clustered (panel), and HAC (time-series) standard errors, and saying why.

C6. Time-series judgement — stationarity, persistence, forecastability, unit roots and spurious regression, and the difference between Granger prediction and causation.

C7. Connecting model to data — matching a real case (the EMH, an oil shock, a GARCH volatility spike, a cointegrated pair) to the relevant tool, in both Stata and Python.

This topic map is the source for the 3 quizzes, the midterm, and the final. Coverage: ***Quiz 1*** *= L1–L3 · **Quiz 2** = L4–L6 · **Quiz 3** = L7–L8 · **Midterm** = L1–L6 (Blocks A–B, S&W Ch 1–7) · **Final** = L7–L10 (Block C, S&W Ch 10, 14–16).*