Financial Econometrics 2026
Shanghai Jiao Tong University
Summer School 2026
Final Examination
With Solutions

Reading time: 10 minutes. **Writing time:** 120 minutes.
Total marks: 100. **Calculators are permitted.**
Instructions: Answer **all** questions. In Section A, circle the letter of the single best answer. In Sections B and C, show your working and write in clear sentences. Standard errors are heteroskedasticity-robust unless a question states that they are clustered or HAC.

Structure:

- **Section A** – Multiple choice: 20 questions, 2.5 marks each (**50 marks**).
- **Section B** – One analytical problem in five parts (**30 marks**).
- **Section C** – One short interpretation (**20 marks**).

Useful facts. $t = \dfrac{\text{estimate} - \text{claimed value}}{SE}$; 95% CI = estimate $\pm\, 1.96\, SE$; AR(1) long-run mean $= \beta_0/(1-\beta_1)$; half-life of a shock $= \ln(0.5)/\ln(\text{persistence})$; 95% one-step forecast interval = forecast $\pm\, 1.96\, SER$; HAC truncation $m = \text{round}\big(0.75\, T^{1/3}\big)$; 5% critical values: ADF (intercept only) -2.86; Engle–Granger -3.41; QLR ($q = 3$) 4.71 (1%: 6.02); GARCH(1,1): $\sigma_t^2 = \omega + \alpha u_{t-1}^2 + \beta \sigma_{t-1}^2$, persistence $= \alpha + \beta$, unconditional variance $= \omega/(1-\alpha-\beta)$; 1-day 99% VaR $= 2.33\, \sigma_t$.

Section A: Multiple Choice

(20 questions × 2.5 marks = 50 marks)

1. A panel regression of bank lending on capital ratios (40 European banks, 12 years) includes entity fixed effects only. Which of the following is still a threat to internal validity?

(A) each bank's home-country legal tradition, fixed over the whole sample.

(B) a change of chief executive at some banks that shifts both capital and lending.

(C) permanent differences in average balance-sheet size across the forty banks.

(D) each bank's charter type, which was set once at its founding date.

Solution. (B). Entity fixed effects absorb every *time-invariant* confounder, observed or not, but nothing that changes within a bank over time. A mid-sample management change moves both regressor and outcome inside the entity, so it survives the fixed effects. (A), (C) and (D) are all fixed traits, exactly what the fixed effects remove.

2. A panel policy coefficient is 0.12 with an ordinary SE of 0.04 and an SE clustered by entity of 0.08. Using the valid standard error, the t-statistic and the 5% verdict are:

(A) $t = 3.0$; significant.

(B) $t = 0.67$; not significant.

(C) $t = 1.5$; significant.

(D) $t = 1.5$; not significant.

Solution. (D). The clustered SE is the valid one in a panel, because an entity's errors are correlated across years: $t = 0.12/0.08 = 1.5 < 1.96$, not significant. (A) uses the ordinary SE, which wrongly treats the entity's years as independent; (B) inverts the ratio (0.08/0.12); (C) pairs the right t with the wrong verdict.

3. Lab 5 uses the Guns panel: 51 US states, identified by `stateid`, observed every year 1977 to 1999. Before running any fixed-effects regression with `xtreg`, the data set must be declared as a panel with:

(A) `xtset stateid year`

(B) `tsset year`

(C) `xtset year stateid`

(D) `egen panelid = group(stateid year)`

Solution. (A). `xtset entity time`, so `xtset stateid year`, exactly as in the lab do-file. (B) declares a single time series and cannot hold 51 states; (C) reverses the arguments, declaring the year as the entity; (D) builds an identifier but declares nothing.

4. In a state-level panel of loan defaults on a regulation dummy with entity fixed effects, the national interest-rate cycle moves defaults in every state in the same years. To absorb this confounder you should add:

(A) standard errors clustered at the state level.

(B) a lagged dependent variable to the regression.

(C) time fixed effects, one for each sample year.

(D) interaction terms between the dummy and state size.

Solution. (C). A shock common to all entities in a given year is exactly what time fixed effects absorb. Clustering (A) fixes the standard errors, not omitted-variable bias; (B) and (D) change the dynamics or the slope but do not remove a common year shock.

5. In Lab 5 the effect of a shall-issue law on the log violent-crime rate is estimated with *state fixed effects* and standard errors *clustered by state* (control variables omitted here). The command is:

(A) `reg lvio shall, vce(cluster stateid)`

(B) `xtreg lvio shall, re vce(cluster stateid)`

(C) `xtreg lvio shall, fe vce(cluster stateid)`

(D) `reg lvio shall, robust`

Solution. (C). `xtreg, fe` estimates the within (fixed-effects) model and `vce(cluster stateid)` clusters by state, as in the lab. (A) is pooled OLS with clustering, no fixed effects; (B) fits random effects, a different estimator; (D) is pooled OLS with only heteroskedasticity-robust SEs.

6. Monthly returns follow the fitted AR(1) $\widehat{r}_t = 0.2 + 0.1\, r_{t-1}$, and the last observed return is $r_T = 2.0$. The one-step-ahead forecast $\widehat{r}_{T+1}$ is:

(A) 2.0.

(B) 0.4.

(C) 2.2.

(D) 0.22.

Solution. (B). $\widehat{r}_{T+1} = 0.2 + 0.1 \times 2.0 = 0.4$. (A) is the random-walk forecast, ignoring the model; (C) adds the intercept to r_T without applying the slope; (D) applies the slope to the intercept instead of to r_T.

7. For an AR model of monthly sales growth you compute, over the same sample:

	AR(1)	AR(2)	AR(3)	AR(4)
BIC	520.0	516.5	514.8	516.2
AIC	512.0	505.5	501.0	499.8

The lag lengths chosen by BIC and AIC are:

(A) BIC chooses three lags; AIC chooses four.

(B) BIC chooses four lags; AIC chooses three.

(C) both criteria choose three lags.

(D) both criteria choose four lags.

Solution. (A). Each criterion is minimised: BIC is smallest at AR(3) (514.8) and AIC at AR(4) (499.8). BIC's penalty per lag, $\ln T/T$, exceeds AIC's $2/T$, so BIC is the more parsimonious criterion and here stops one lag earlier.

8. An AR(1) fitted to monthly excess stock returns over 600 months gives a slope of 0.05 with standard error 0.06 and $R^2 = 0.003$. The best reading of this output is:

(A) the model is too short; keep adding lags until the fit starts to improve.

(B) the positive slope shows momentum that a disciplined trader could exploit.

(C) the low R^2 shows the sample is too small for reliable estimation.

(D) the result is exactly what market efficiency predicts for excess returns.

Solution. (D). An insignificant slope ($t = 0.83$) and a near-zero R^2 are the *correct* answer for excess returns: any reliable pattern would be arbitraged away. (A) is specification mining; (B) ignores that the slope is statistically indistinguishable from zero; (C) is wrong because 600 months is a large sample, precision is not the problem.

9. Testing an AR(2) of the change in inflation for a break at an unknown date, you compute the Chow F at every date in the central 70% of the sample; the largest value is 5.10, in 2009:Q2 ($q = 3$ coefficients may break). The QLR 5% critical value is 4.71 (1%: 6.02); the ordinary F cutoff at 5% would be 2.60. The correct conclusion is:

(A) reject stability at the 1% level, because 5.10 is far above the ordinary cutoff 2.60.

(B) reject stability at the 5% level, because 5.10 exceeds the QLR value 4.71.

(C) do not reject stability, because 5.10 is below the 1% value 6.02.

(D) do not reject stability, because a statistic found by searching over dates cannot be tested.

Solution. (B). The largest of many Chow statistics must be judged against the larger QLR critical values, because searching over dates inflates the ordinary distribution: $5.10 > 4.71$, so reject stability at 5% (but not at 1%, since $5.10 < 6.02$). (A) uses the ordinary cutoff, which is invalid after a search; (C) states the 1% comparison but draws the wrong 5% conclusion; (D) is defeatist, the QLR distribution is built exactly to handle the search.

10. A retailer's monthly sales growth follows the fitted AR(1) $\widehat{g}_t = 0.90 + 0.70\, g_{t-1}$ with $SER = 1.50$; the latest value is $g_T = 5.00$, so the point forecast is $\widehat{g}_{T+1} = 4.40$. The 95% one-step forecast interval is approximately:

(A) $[1.5,\ 7.3]$.

(B) $[2.9,\ 5.9]$.

(C) $[4.2,\ 4.6]$.

(D) $[2.1,\ 7.9]$.

Solution. (A). Forecast $\pm 1.96\, SER = 4.40 \pm 1.96 \times 1.50 = 4.40 \pm 2.94 = [1.46,\ 7.34]$. The interval is wide because it must contain next month's *shock*, not just estimation error. (B) forgets the 1.96; (C) treats the SER like the standard error of an estimated mean, shrinking it by a sample-size factor and confusing a forecast interval with a CI for the mean; (D) centres the interval on g_T instead of the forecast.

11. In Lab 6 the AR models for the change in inflation are estimated with lag operators, for example `reg d_infl L.d_infl L2.d_infl, robust`. For the `L.` operators to work, the do-file must first declare the quarterly date variable with:

(A) `sort qdate`

(B) `xtset qdate`

(C) `format qdate %tq`

(D) `tsset qdate`

Solution. (D). `tsset qdate` declares the data as a time series indexed by `qdate`, which is what activates `L.`, exactly as in the lab do-file. (A) only orders the rows; (B) makes `qdate` the panel (entity) identifier and declares no time variable, so `L.` still fails; (C) only sets the display format of the date.

12. A distributed-lag regression of inflation on oil-price growth gives an impact multiplier of 0.5 and dynamic multipliers of 0.3 (lag 1) and 0.1 (lag 2). The cumulative dynamic multiplier at horizon $h = 2$ is:

(A) 0.5.

(B) 0.4.

(C) 0.9.

(D) 0.1.

Solution. (C). The cumulative multiplier is the running sum of the dynamic multipliers: $0.5 + 0.3 + 0.1 = 0.9$. (A) is the impact effect alone; (B) sums the lags but omits the impact; (D) is the horizon-2 dynamic multiplier, not the cumulative one.

13. You estimate a distributed-lag regression on $T = 512$ monthly observations. The rule-of-thumb HAC truncation parameter $m = \text{round}\big(0.75\,T^{1/3}\big)$ is:

(A) 6.

(B) 8.

(C) 17.

(D) 384.

Solution. (A). $512^{1/3} = 8$, so $m = 0.75 \times 8 = 6$. (B) forgets the 0.75 factor; (C) applies the 0.75 factor to the *square* root of T instead of the cube root ($0.75 \times \sqrt{512} \approx 17$); (D) computes $0.75\,T$, skipping the cube root.

14. In Lab 7, monthly CPI inflation is regressed on the current value and six lags of oil-price growth. The residuals are strongly serially correlated ($\widehat{\rho}_1 \approx 0.56$), and with $T = 749$ the truncation is $m = 7$. The command that reports the correct (HAC) standard errors is:

(A) `regress infl L(0/6).oilg, robust`

(B) `newey infl L(0/6).oilg, lag(7)`

(C) `regress infl L(0/6).oilg, vce(cluster mdate)`

(D) `arch infl L(0/6).oilg, arch(7)`

Solution. (B). `newey` computes Newey–West (HAC) standard errors with truncation `lag(7)`, as in the lab. (A) is only heteroskedasticity-robust and ignores serial correlation; (C) clusters on the date, which puts each observation in its own cluster and fixes nothing; (D) fits an ARCH volatility model, a different object entirely.

15. You want causal dynamic multipliers from a distributed-lag regression of consumer-price inflation in a small oil-importing economy. Which regressor most plausibly satisfies exogeneity?

(A) the domestic policy interest rate, set in response to inflation.

(B) domestic wage growth, negotiated partly in response to past inflation.

(C) the lagged dependent variable, inflation in the previous month.

(D) the world oil price, set on global markets the economy cannot move.

Solution. (D). A small importer takes the world oil price as given, so feedback from its inflation to the regressor is negligible, which is what exogeneity requires. (A) and (B) respond to inflation, a reverse-causality problem; (C) is correlated with past error terms by construction.

16. An ADF test on a log stock-price index gives $t = -2.10$; the ADF 5% critical value is -2.86. The correct conclusion is:

(A) fail to reject the unit root, but model the series in levels, not differences.

(B) the series is stationary, because -2.10 lies above the critical value -2.86.

(C) fail to reject the unit root; treat the series as I(1) and difference it.

(D) there is no unit root, but heteroskedasticity requires HAC critical values instead.

Solution. (C). -2.10 is not more negative than -2.86, so the unit-root null stands: treat the level as I(1) and work with the differenced series (returns). (A) reaches the right verdict but keeps the I(1) series in levels, inviting a spurious regression; (B) reads the comparison backwards; (D) confuses HAC inference with unit-root testing.

17. Which pair of series is the most plausible candidate for cointegration?

(A) the 10-year and the 2-year government bond yield of the same country.

(B) the GDP levels of two small economies with no economic links.

(C) a stationary interest-rate spread and a random-walk stock price.

(D) daily stock returns and daily bond returns, both roughly white noise.

Solution. (A). Two government bond yields of the same country are pulled along by one common trend, the general level of interest rates, so both are I(1) but their spread is stationary: the definition of cointegration. (B) shares no common trend; (C) mixes I(0) with I(1); (D) is two stationary series, so cointegration does not apply.

18. An error-correction model gives $\Delta Y_t = 0.1 - 0.25\,\widehat{z}_{t-1} + u_t$, where $\widehat{z}_{t-1}$ is the lagged equilibrium error. Last period $\widehat{z}_{t-1} = +2$. The contribution of the error-correction term to this period's change is:

(A) $+0.5$.

(B) -0.5.

(C) -0.25.

(D) -0.125.

Solution. (B). The term contributes $-0.25 \times 2 = -0.5$: Y was above its long-run relationship, so it is pulled back down. (A) has the wrong sign; (C) ignores the size of the gap; (D) divides the coefficient by the gap instead of multiplying.

19. A GARCH(1,1) fit to daily returns gives $\sigma_t^2 = 0.02 + 0.08\,u_{t-1}^2 + 0.90\,\sigma_{t-1}^2$. The unconditional (long-run) daily variance implied by this model is:

(A) 0.02.

(B) 0.20.

(C) 0.98.

(D) 1.0.

Solution. (D). The persistence is $\alpha + \beta = 0.08 + 0.90 = 0.98$, so the unconditional variance is $\omega/(1-\alpha-\beta) = 0.02/(1-0.98) = 0.02/0.02 = 1.0$. (A) reports ω alone, forgetting to divide by one minus the persistence; (B) takes β alone as the persistence, $0.02/(1-0.90) = 0.20$; (C) reports the persistence $\alpha+\beta$ itself, not the variance.

20. Still with $\sigma_t^2 = 0.02 + 0.08\,u_{t-1}^2 + 0.90\,\sigma_{t-1}^2$ (daily returns in %): yesterday's return was -2% and yesterday's conditional variance was 1.5. Today's conditional *volatility* is:

(A) 1.69%.

(B) 1.1%.

(C) 1.3%.

(D) 1.9%.

Solution. (C). $\sigma_t^2 = 0.02 + 0.08 \times 4 + 0.90 \times 1.5 = 0.02 + 0.32 + 1.35 = 1.69$, so $\sigma_t = \sqrt{1.69} = 1.3\%$. (A) reports the variance without taking the square root; (B) plugs in $u_{t-1} = -2$ instead of $u_{t-1}^2 = 4$; (D) swaps α and β in the recursion.

Section B: Analytical Problem (30 marks)

You have joined the rates desk of a London fixed-income fund. You are given monthly data, January 1986 to December 2025 ($T = 480$ months), on the 10-year government bond yield $R10_t$ and the 3-month rate $R3_t$ (both in percentage points). The fund also runs an equity book whose daily percentage returns you model for risk. Your desk head asks for the following analysis.

The spread $S_t = R10_t - R3_t$ has the fitted AR(1)

$$\widehat{S}_t = 0.30 + 0.80\, S_{t-1}, \qquad SER = 0.25,$$

and the latest observed value is $S_T = 2.00$.

(a) **[6 marks]** Compute the one-step-ahead forecast $\widehat{S}_{T+1}$ and the 95% one-step forecast interval. The forecast lies *below* the last observed value; explain why, using the long-run mean of this AR(1).

(b) **[6 marks]** A colleague warns you against regressing the two rates on each other in levels. ADF tests (intercept only) give $t = -1.35$ for $R10$ and $t = -1.62$ for $R3$; the 5% critical value is -2.86. State the conclusion of each test, the implied order of integration of the two rates, and the danger her warning refers to.

(c) **[6 marks]** You run the levels regression anyway and obtain

$$\widehat{R10}_t = 1.50 + 0.90\, R3_t,$$

with residual $\widehat{z}_t$. An ADF test on $\widehat{z}_t$ gives $t = -3.70$; the Engle–Granger 5% critical value is -3.41. State the conclusion, explain why the Engle–Granger critical value is used instead of the standard ADF value, and say what the result means for the validity of the levels regression.

(d) **[6 marks]** The desk's error-correction model for the 10-year rate is

$$\Delta R10_t = 0.02 - 0.20\, \widehat{z}_{t-1} + u_t,$$

and currently the 10-year sits above its long-run relation: $\widehat{z}_{t-1} = +0.60$. Compute the contribution of the error-correction term to next month's change, state what fraction of the gap closes each month, and compute the half-life of the gap.

(e) **[6 marks]** For the equity book, a GARCH(1,1) on daily returns (in %) gives

$$\sigma_t^2 = 0.04 + 0.10\, u_{t-1}^2 + 0.85\, \sigma_{t-1}^2.$$

Yesterday the market fell 4% and yesterday's conditional variance was 2.25. Compute today's conditional variance and volatility, and the 1-day 99% VaR. The long-run (unconditional) daily volatility of this model is about 0.89%; in one sentence, explain why the desk should not freeze its VaR limit at the calm-day level.

Solution. (a) $\widehat{S}_{T+1} = 0.30 + 0.80 \times 2.00 = \mathbf{1.90}$. Interval: $1.90 \pm 1.96 \times 0.25 = 1.90 \pm 0.49 = [\mathbf{1.41},\ \mathbf{2.39}]$. The long-run mean is $\beta_0/(1-\beta_1) = 0.30/0.20 = \mathbf{1.50}$. Because $\beta_1 = 0.80 < 1$, the series is stationary and mean-reverting: from $S_T = 2.00$, above the long-run mean, the forecast moves part of the way back towards 1.50, so it lies below the last value.

(b) $R10$: -1.35 is not more negative than -2.86, so **fail to reject** the unit root. $R3$: -1.62 is not more negative than -2.86, so **fail to reject** as well. Both rates are treated as **I(1)**: their levels behave like random walks and only their changes are stationary. The danger is the **spurious regression**: regressing one I(1) series on another can produce a large t-statistic and a high R^2 even when no true relationship exists, so a levels regression is not to be trusted until cointegration is established.

(c) -3.70 is more negative than -3.41, so **reject** the unit root in the residual: $\widehat{z}_t$ is stationary and the two rates are **cointegrated**, sharing one stochastic trend, with long-run relation $R10 = 1.50 + 0.90\,R3$. The Engle–Granger critical value is used because the residual is *estimated*: OLS chooses the coefficients that make the residual look as stationary as possible, so the test needs a stricter (more negative) cutoff than the standard ADF value. Given cointegration, the levels regression is **valid**: this is the exception to the rule in (b).

(d) Contribution of the error-correction term: $-0.20 \times 0.60 = \mathbf{-0.12}$ percentage points, pulling the 10-year rate back down towards the long-run relation. The coefficient -0.20 means about **20% of the gap closes each month**. Treating the gap itself as decaying at that rate (the short rate holding steady, so $\widehat{z}_t \approx 0.80\,\widehat{z}_{t-1}$), the half-life is $\ln(0.5)/\ln(0.80) = (-0.693)/(-0.223) \approx \mathbf{3.1\ months}$.

(e) $u_{t-1}^2 = (-4)^2 = 16$. $\sigma_t^2 = 0.04 + 0.10 \times 16 + 0.85 \times 2.25 = 0.04 + 1.60 + 1.9125 = \mathbf{3.5525}$ (about 3.55). Today's volatility: $\sigma_t = \sqrt{3.5525} \approx \mathbf{1.88}\%$. VaR: $2.33 \times 1.88 \approx \mathbf{4.39}\%$ of the book, the loss exceeded only about one day in a hundred under normality. A frozen calm-day limit ($2.33 \times 0.89 \approx 2.1\%$) would understate today's risk by half: VaR scales with the *conditional* volatility, which the GARCH recursion updates daily, and with persistence $\alpha + \beta = 0.95$ the elevated volatility will fade only slowly.

Section C: Interpretation (20 marks)

Read the following and answer in **no more than five sentences**. Use the concepts from the course; you do not need any calculation.

A Frankfurt asset manager wants to launch a "Governance Leaders" fund. Its quant team assembles a panel of 400 European firms observed yearly from 2012 to 2021 and regresses return on assets (ROA) on each firm's governance score. Analyst A pools all 4,000 firm-years: the coefficient is $+0.15$ *with* $t = 5.2$ *using ordinary standard errors, and he proposes putting this number in the marketing deck as the payoff to better governance. Analyst B re-runs the regression with firm and year fixed effects and with standard errors clustered by firm: the coefficient falls to* $+0.02$ *with* $t = 0.5$*. Analyst A objects: "your fixed effects threw away most of the variation in the data; my regression uses all of it, and my t-statistic is far bigger."*

Question [20 marks]. Whose estimate is the better guide to what *improving a firm's governance* would do to its profitability, and why? Your answer should say what the fixed effects remove and why the coefficient fell, what is wrong with Analyst A's "more variation" objection, and what the clustered standard errors correct. Answer in at most five sentences.

Solution. Model answer (within is credible, between is contaminated). Analyst B's estimate is the better guide. Analyst A's pooled coefficient compares *different* firms, so it absorbs every time-invariant trait that raises both governance scores and profitability (sector, firm quality, management culture): classic omitted-variable bias that no t-statistic can repair. Firm fixed effects compare each firm *with itself* as its own governance changes, and year fixed effects absorb shocks common to all firms in a year, so the fall from 0.15 to 0.02 reveals that the raw association was almost entirely a between-firm difference, not the effect of improving governance. Analyst A's objection has it backwards: the between-firm variation his regression "uses" is precisely the contaminated part, so discarding it is the point, not a flaw. Finally, a firm's errors are correlated across years, so ordinary standard errors treat 4,000 firm-years as independent and come out too small; the clustered ones are the honest ones, and they say the within-firm effect is statistically indistinguishable from zero.

Marking guide (20 marks).

- **8 marks** – diagnoses the pooled estimate: time-invariant firm traits confound the between-firm comparison (omitted-variable bias); fixed effects compare a firm with itself over time, and the collapse from 0.15 to 0.02 shows the raw association was between firms.
- **6 marks** – answers the "more variation" objection: between variation is the contaminated part, so removing it is the design, not a defect; year effects absorb common

shocks.

- **4 marks** – explains clustering: errors correlated within a firm over time make ordinary SEs too small, so the verdict must use the clustered SE, under which the effect is not significant.
- **2 marks** – clear writing within the five-sentence limit.

Answers that side with Analyst A because of the larger t-statistic earn very little: that is precisely the fallacy the question tests.