Financial Econometrics 2026
Shanghai Jiao Tong University
Summer School 2026
Quiz 2
With Solutions

Reading time: 10 minutes. **Writing time:** 60 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. All standard errors are heteroskedasticity-robust.

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 $\pm\, 1.96\, SE$ (90%: 1.64; 99%: 2.58); small-sample two-sided 5% t: 2.23 ($df = 10$), 2.09 ($df = 20$), 2.00 ($df = 60$), $\to$ 1.96 as df grows; sign of omitted-variable bias = $\text{sign}\big[\text{corr}(X, Z) \times \text{effect of } Z \text{ on } Y\big]$; $\bar{R}^2 = 1 - \frac{n-1}{n-k-1}\frac{SSR}{TSS}$; 5% critical values of F: 3.84 ($q = 1$), 3.00 ($q = 2$), 2.60 ($q = 3$); 1%, $q = 2$: 4.61; $F = t^2$ for one restriction.

Section A: Multiple Choice (20 questions × 2.5 marks = 50 marks)

1. The standard error of an estimated regression slope measures:

(A) how much the estimate would vary across repeated samples.

(B) the spread of the outcome Y around its sample mean.

(C) the typical size of the residuals around the fitted line.

(D) the share of the variation in Y the line explains.

Solution. (A). The standard error is the standard deviation of the estimator's sampling distribution: roughly how far the slope computed from one sample typically lands from the true slope. The average miss of the line is the SER (C), and the explained share is R^2 (D); the three numbers answer different questions.

2. A regression gives a slope of 1.5 with standard error 0.6. Testing $H_0 : \beta_1 = 0$ at the 5% level:

(A) $t = 0.4$; do not reject.

(B) $t = 2.5$; reject.

(C) $t = 2.5$; do not reject.

(D) $t = 0.9$; reject.

Solution. (B). $t = (1.5 - 0)/0.6 = 2.5$, and $2.5 > 1.96$, so the estimate is too far from zero to be sampling luck at the 5% level: reject. The t-statistic counts standard errors between the estimate and the claimed value.

3. A slope is estimated as -2.0 with standard error 0.5. The 95% confidence interval is:

(A) $[-2.50, -1.50]$.

(B) $[-2.98, -1.02]$.

(C) $[-2.82, -1.18]$.

(D) $[-3.29, -0.71]$.

Solution. (B). $-2.0 \pm 1.96 \times 0.5 = -2.0 \pm 0.98 = [-2.98, -1.02]$. Option (A) uses one standard error, (C) uses the 90% multiplier 1.64, and (D) the 99% multiplier 2.58. Since 0 is outside the interval, a 5% test would also reject $\beta_1 = 0$.

4. In the regression $Y_i = \beta_0 + \beta_1 D_i + u_i$, where D_i is 1 for firms with a female CEO and 0 otherwise, β_0 is:

(A) the average of Y among firms with a female CEO.

(B) the gap in average Y between the two groups.

(C) the average of Y among firms without a female CEO.

(D) the sample share of firms with a female CEO.

Solution. (C). With a binary regressor there is no line, just two group averages: $\beta_0 = \mathbb{E}[Y \mid D = 0]$, the average outcome in the $D = 0$ group, and $\beta_1 = \mathbb{E}[Y \mid D = 1] - \mathbb{E}[Y \mid D = 0]$, the gap between the groups (B). The usual standard errors, t-statistics and intervals apply unchanged to both.

5. *Heteroskedasticity* means that:

(A) the errors average to zero at every X.

(B) the values of X are all very similar.

(C) the residuals sum to zero by construction.

(D) the spread of the errors changes with X.

Solution. (D). Homoskedasticity is a constant error spread at every value of X; heteroskedasticity is the fan shape, a spread that widens or narrows with X. Returns data are a standard example: calm years and crisis years have very different spreads. The simple SE formula assumes constant spread and gives the wrong number when that fails.

6. The rule taught in this course about robust (heteroskedasticity-robust) standard errors is that they:

(A) are always larger than the homoskedasticity-only ones.

(B) are valid whether or not the error spread is constant.

(C) are needed only when the sample size is very small.

(D) are the default output printed by most software.

Solution. (B). Robust standard errors are correct with or without a constant error spread, and they agree with the simple formula when the spread happens to be constant, so nothing is lost by using them: hence the rule, always use robust. They are not always larger (in the growth-on-trade example the robust SE was smaller), and most software prints the simple ones by default, so you must ask for them.

7. With only $n = 12$ observations, the correct two-sided 5% cutoff for a t-statistic is closest to:

(A) 1.64.

(B) 1.96.

(C) 2.00.

(D) 2.23.

Solution. (D). With few observations the Student t cutoff replaces 1.96: at $n-2=10$ degrees of freedom the 5% two-sided value is 2.23. The cutoff shrinks toward 1.96 as n grows (2.09 near $n=22$, 2.00 near $n=62$). In finance samples are usually large, but a two-year monthly track record is exactly the case to watch.

8. Leaving a variable Z out of a regression of Y on X biases the slope on X when:

(A) Z affects Y, whether or not it is related to X.

(B) Z is correlated with X, whether or not it affects Y.

(C) either one of the two conditions holds.

(D) Z affects Y and Z is correlated with X.

Solution. (D). Omitted-variable bias needs both conditions at once. If Z affects Y but is unrelated to X, it adds noise but steals no credit; if Z moves with X but does not affect Y, there is nothing to steal. Only when the omitted influence travels together with X does the slope on X absorb credit that belongs to Z.

9. An omitted variable Z *raises* Y, and Z is *positively* correlated with X. The estimated slope on X in the regression without Z is biased:

(A) downward.

(B) not at all.

(C) upward.

(D) towards exactly zero.

Solution. (C). Sign of the bias $=$ sign[corr$(X,Z)\times$ effect of Z on Y] $=(+)\times(+)=+$: the slope is pushed up. Intuitively, whenever X is high, the helpful Z tends to be high too, and the lone regressor X collects credit for Z's work.

10. If a regression suffers from omitted-variable bias, doubling the sample size:

(A) removes the bias slowly.

(B) removes it once n is large enough.

(C) cuts the bias roughly in half.

(D) does not reduce the bias at all.

Solution. (D). The bias formula, $\hat{\beta}_1 \to \beta_1 + \rho_{Xu}\sigma_u/\sigma_X$, contains no n anywhere: omitted-variable bias is a systematic error of aim, not sampling noise. More data makes the estimate settle ever more precisely on the *wrong* value. The cures are an experiment, within-group comparison, or putting Z into the regression.

11. In the multiple regression $Y_i = \beta_0 + \beta_1 X_{1i} + \beta_2 X_{2i} + u_i$, the coefficient β_1 measures the effect on Y of a one-unit change in X_1:

(A) holding X_2 fixed.

(B) and the accompanying change in X_2.

(C) averaged over all values of X_2 and u.

(D) ignoring the presence of X_2.

Solution. (A). This is the partial effect, the one idea of the chapter: with X_2 in the equation, β_1 is the effect of moving X_1 while X_2 is held fixed, so each coefficient earns only its own credit. Option (D) describes the single-regressor slope, which is exactly what omitted-variable bias contaminates.

12. Adding a regressor that is pure noise to a regression will:

(A) lower R^2, while possibly raising $\bar{R}^2$.

(B) lower both R^2 and $\bar{R}^2$ at the same time.

(C) never lower R^2, and possibly lower $\bar{R}^2$.

(D) raise both R^2 and $\bar{R}^2$ in every case.

Solution. (C). R^2 can only rise or stay equal when a variable is added, even a useless one, because OLS can always set its coefficient to zero or better in-sample. Adjusted R^2 charges a penalty per variable, so a noise regressor typically pulls it down. A widening gap between the two is the sign that a model carries variables it cannot afford.

13. A regression includes an intercept and dummies for *every* sector (technology, financials, other). The standard fix is to:

(A) add one further dummy variable for large firms.

(B) drop one sector dummy, which becomes the baseline.

(C) drop the dependent variable from the equation.

(D) estimate the equation without any sector dummies.

Solution. (B). With all sectors dummied, the dummies sum to one for every firm, an exact copy of the intercept's constant: perfect multicollinearity, and the software cannot run (or silently drops something). Drop one category and it becomes the baseline; every remaining dummy coefficient is then a comparison against it. Keeping all dummies but dropping the intercept works too.

14. Two regressors are highly, but not perfectly, correlated (say 0.95). The consequence is:

(A) unbiased estimates with large standard errors.

(B) biased estimates with small standard errors.

(C) that the software refuses to run the regression.

(D) that the residuals no longer sum to zero.

Solution. (A). Imperfect multicollinearity is not a mistake; it is a feature of the data. OLS runs and remains unbiased, but the sample contains few moments where one twin moved and the other did not, so their separate effects are hard to tell apart: honest but blurry, meaning large standard errors. Refusal to run (C) is the symptom of *perfect* multicollinearity.

15. A null hypothesis restricts two coefficients to zero ($q = 2$). If you instead run two separate t-tests at 5% and reject when either is significant, the true false-alarm rate is about:

(A) 5%.

(B) 2.5%.

(C) 10%.

(D) 25%.

Solution. (C). If the null is true, each t stays inside ± 1.96 with probability 0.95, so both stay inside with probability $0.95^2 = 0.9025$: the chance of at least one false alarm is 9.75%, nearly double the promised 5%. A test whose real error rate differs from its promised one has the wrong size, which is why joint hypotheses need the F-test.

16. A robust F-statistic for a joint null with $q = 2$ restrictions comes out at 4.2. The conclusion is:

(A) reject at 1% and at 5%.

(B) reject at neither level.

(C) reject at 1% but not at 5%.

(D) reject at 5% but not at 1%.

Solution. (D). The large-sample critical values for $q = 2$ are 3.00 at 5% and 4.61 at 1%. Since $3.00 < 4.2 < 4.61$, the joint null is rejected at the 5% level but survives at 1%. Rejection at 1% always implies rejection at 5%, so option (C) is impossible.

17. For a single restriction ($q = 1$), the F-statistic and the t-statistic for the same null are related by:

(A) $F = t$.

(B) $F = 2t$.

(C) $F = t^2$.

(D) $F = t/2$.

Solution. (C). With one restriction the two tests are the same test: $F = t^2$, and the critical values agree exactly, $3.84 = 1.96^2$. This is a useful consistency check whenever software reports both.

18. In a regression with two correlated regressors, each t-statistic is small, yet the joint F-test rejects strongly. The explanation is:

(A) the regressors are correlated and share the credit.

(B) the sample is far too small for reliable t-tests.

(C) a coding error somewhere in the estimation software.

(D) heteroskedasticity in the regression's error terms.

Solution. (A). Correlated regressors are twins: the data cannot say which one did the work, so each coefficient alone is imprecise (small t), while together their joint contribution is unmistakable (large F). This is not a paradox but the standard signature of shared credit; the F-statistic's correlation term is built to handle exactly this case.

19. A very high R^2 in a regression is, by itself, evidence that:

(A) none of the following.

(B) the included variables cause Y.

(C) there is no omitted-variable bias.

(D) the right variables have been chosen.

Solution. (A). R^2 measures in-sample fit and nothing else. It does not establish significance, causation, the absence of omitted-variable bias, or that the variable list is right; a beautiful fit around a biased coefficient is convincing and wrong. Variables are chosen by hunting confounders of the variable of interest, not by chasing fit.

20. In a carefully designed regression, the coefficient on a *control* variable (included only to guard the variable of interest):

(A) is the causal effect of the control.

(B) is not given a causal reading.

(C) must exceed 1.96 to keep the control.

(D) equals the bias it removes.

Solution. (B). Controls stand guard: they hold a would-be credit thief fixed so the variable of interest is measured fairly. Their own coefficients typically mix several influences (the lunch-programme coefficient bundles everything that travels with poverty) and are not interpreted causally. A control also stays in the regression on design grounds, significant or not.

Section B: Analytical Problem (30 marks)

A fund's monthly excess return R (in % per month) is regressed on the market factor MKT and the size factor SMB, using $n = 120$ months of data (robust standard errors in parentheses):

$$\widehat{R} = \underset{(0.10)}{0.20} + \underset{(0.05)}{1.10}\, MKT + \underset{(0.20)}{0.50}\, SMB, \qquad R^2 = 0.62, \quad \bar{R}^2 = 0.61.$$

(a) [**6 marks**] Interpret the coefficient 1.10 on MKT in one sentence, and test whether it differs from zero at the 5% level.

(b) [**6 marks**] The intercept (0.20) is the fund's alpha. Test $H_0 : \alpha = 0$ at the 5% level and give the 95% confidence interval for alpha. What does your conclusion say about the manager?

(c) [**6 marks**] Test whether the fund's market exposure differs from exactly 1, that is $H_0 : \beta_{MKT} = 1$, at the 5% level.

(d) [**6 marks**] A one-factor regression of the same fund on MKT alone gives an alpha of 0.35. Using the omitted-variable-bias sign rule, explain why the one-factor alpha is larger than the three lines above suggest it should be. (The fund's returns load positively on SMB, and assume the fund's SMB exposure is positively related to its market exposure across months.)

(e) [**6 marks**] A sceptic claims that neither factor matters: $H_0 : \beta_{MKT} = 0$ **and** $\beta_{SMB} = 0$. The software reports a robust F-statistic of 8.2 for this joint null. State the conclusion at the 5% and 1% levels, and explain in one sentence why this joint test, rather than the two t-tests, is the right tool.

Solution. (a) Holding the fund's SMB exposure fixed, a month in which the market factor is one percentage point higher is associated with a fund return 1.10 points higher: the fund moves slightly more than one-for-one with the market. $t = 1.10/0.05 = \mathbf{22}$, far beyond 1.96: the market loading is overwhelmingly significant.

(b) $t = 0.20/0.10 = \mathbf{2.0} > 1.96$: reject $H_0 : \alpha = 0$ at the 5% level, though only just. The 95% CI is $0.20 \pm 1.96 \times 0.10 = [\mathbf{0.004}, \mathbf{0.396}]$, which excludes zero by a whisker. Statistically, there is evidence of performance beyond the two factor exposures; economically, the plausible range runs from almost nothing to about 0.4% per month, so the evidence for skill is real but fragile.

(c) $t = \dfrac{1.10 - 1}{0.05} = \dfrac{0.10}{0.05} = \mathbf{2.0} > 1.96$: reject $H_0 : \beta_{MKT} = 1$ at the 5% level. The fund is statistically distinguishable from a pure index exposure, though the difference

(0.10) is economically modest. Note the t-statistic measures distance from the *claimed* value, which need not be zero.

(d) Omitting SMB biases what remains. The omitted factor *raises* the fund's returns (positive loading, +0.50) and is *positively* related to the included exposure, so the bias sign is $(+) \times (+) = +$: the one-factor regression pushes credit for the size tilt into the intercept and slope. The inflated alpha of 0.35 is mostly the reward for holding small stocks, not manager skill; with SMB included, alpha falls to 0.20. Never judge an alpha before the relevant factor exposures are held fixed.

(e) For $q = 2$ the critical values are 3.00 (5%) and 4.61 (1%); since $8.2 > 4.61$, **reject the joint null at both levels**. The joint claim must be judged jointly: two separate 5% t-tests have a real false-alarm rate near 10% (wrong size), and with correlated factors each t alone can understate what the pair explains together.

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.

An analyst regresses the annual stock returns of 300 firms on each firm's ESG score (a measure of environmental and social practices) and finds a positive, statistically significant slope. She concludes that good ESG practices raise returns and recommends investing accordingly. A colleague repeats the exercise with one change: he adds firm size to the regression, since large firms tend to have both high ESG scores and returns that behave differently. In his regression the ESG coefficient falls by more than half and is no longer statistically significant.

Question [20 marks]. Explain what happened to the ESG coefficient and why, name the problem in the analyst's original regression, and state what her one-variable slope was actually measuring. Answer in at most five sentences.

Solution. Model answer (omitted-variable bias, exposed by adding the control). The analyst's regression omitted firm size, a variable that both affects returns and is correlated with the ESG score, so the two conditions for **omitted-variable bias** were met and her slope was biased. Her one-variable coefficient was not the effect of ESG; it was the ESG effect *plus stolen credit* for whatever large firms do differently, since high-ESG firms tend to be large. When the colleague held size fixed, each variable could earn only its own credit, and the ESG coefficient collapsed, showing most of the original slope belonged to size. Statistical significance in the first regression was never evidence against this: bias is systematic, so a large sample simply measures the wrong number precisely. The honest conclusion is that the data provide little evidence that ESG practices themselves raise returns once size is held fixed, and any remaining claim should be tested with further sensible controls in place.

Marking guide (20 marks).

- **8 marks** — names **omitted-variable bias** and verifies its two conditions here: size affects returns **and** size is correlated with the ESG score.
- **6 marks** — explains the mechanism: the one-variable slope bundled ESG's effect with size's effect (stolen credit); holding size fixed reassigns the credit, so the coefficient falls.
- **4 marks** — notes that significance does not protect against bias (no n in the bias formula), and states what the original slope actually measured: a mixture, not the ESG effect.
- **2 marks** — clear writing within the five-sentence limit.

Answers that say only "correlation is not causation" without naming the omitted variable and its two conditions earn partial credit at most: the question asks for the specific mechanism.