Problem Set 4

Econ 6376 — Time Series Econometrics

Parameters

- Feel free to work in groups, however your code and write ups must be your own.
- Submissions should be made through Blackboard and include two files:
  - Your write up as a single PDF, **no more than 10 pages**, ending with the **AI Use Disclosure** paragraph.
  - Your **organized, commented** R script that reproduces every result in your write up, top to bottom, from the supplied `ps4_series.csv` and your original midterm `data.csv`.
- Your submission is due by 11:59pm on the due date listed in Blackboard.
- Late submissions **WILL NOT** be accepted unless previously arranged.
- Make sure your answer is sufficient for a reader that is not familiar with the problem and/or time series to understand your solution.
- Treat your write up like a professional report. If you wouldn't turn it in to your boss, don't turn it in to me.

Use of AI

AI assistance is permitted on this problem set, as it is on every graded deliverable in this course. What gets graded is understanding you can defend.

End your write up with a short, free-form paragraph headed **AI Use Disclosure** describing how you used it — which tools, for what, and where you checked its work. It does not count toward your 10 pages and it is not graded on its content. A submission without it is incomplete, in the same way a submission without your R script is incomplete.

Grading

Component	Points
Question 1 — Cointegration and predictive direction	60
Question 2 — Cointegration in your midterm data	40
Total	100

Part	Question 1	Question 2
(a)	6	10
(b)	10	14
(c)	14	8
(d)	12	8
(e)	10	—
(f)	8	—

The Data and the Design

Question 1 uses the supplied `ps4_series.csv`. It contains 500 equally spaced observations of three simulated series, `X`, `Y`, and `Z`, and an observation index, `t`. The index is not a calendar date or another variable to include automatically in your regressions. Treat the generating process as unknown. Use all 500 observations, allowing for the rows needed to construct lags.

Question 2 uses **the exact `data.csv` you submitted for the midterm**. Use the original **training block only** throughout that question. Do not use the `weights` or `holdout` blocks, change the original split, or download a revised vintage. Work with the original levels, or a justified transformation such as logs, before assessing integration; do not begin a cointegration analysis by replacing every series with its first difference.

```
sim <- read.csv("ps4_series.csv")
midterm <- read.csv("data.csv")
train <- midterm[midterm$block == "train", , drop = FALSE]
```

Keep observations in their original chronological order. For each pair, align the two series by their training dates and verify an equally spaced, complete common sample. Do not silently delete interior missing dates and treat the remaining rows as consecutive periods. Report how you handle missing observations or an unsuitable sample; do not invent values to make a test run.

Engle–Granger Regression Framework

Use this framework in both questions. It concerns **two plausibly I(1) series**, an intercept but **no deterministic time trend** in the candidate levels relationship, and a 5% test. A high levels correlation or R^2 does not establish cointegration. First determine whether the integration evidence makes this procedure appropriate.

For an eligible pair, call the chosen left-hand-side series A_t and the other series B_t. State your normalization before inspecting the cointegration test result. Estimate

$$A_t = c + \kappa B_t + v_t, \qquad e_t = A_t - \hat{c} - \hat{\kappa} B_t.$$

The fitted residual e_t estimates a candidate equilibrium deviation; it is not an observed innovation. Estimate the residual ADF regression

$$\Delta e_t = \gamma e_{t-1} + \sum_{j=1}^{p_e} \psi_j \Delta e_{t-j} + \epsilon_t^{ADF}.$$

Include **no additional intercept or trend** in this second regression. The test statistic is the t-ratio on the lagged residual,

$$\tau_{EG} = \frac{\hat{\gamma}}{\widehat{\mathrm{se}}(\hat{\gamma})}.$$

Test H_0: **no cointegration**, against the stationary-deviation alternative under the maintained I(1) setup. Use the following approximate 5% cointegration critical value from Module 10:

$$\mathrm{CV}_{0.05} = -3.33613 - \frac{6.1101}{n_c} - \frac{6.823}{n_c^2}, \qquad n_c = n - 1,$$

where n is the number of aligned **levels observations for that pair**. It is not the length of the full midterm file, the residual degrees of freedom, or the number of usable residual-ADF rows. The latter is $n - 1 - p_e$. Reject at 5% when $\tau_{EG} < \mathrm{CV}_{0.05}$.

This is the MacKinnon (2010) bivariate intercept/no-trend response-surface approximation used in the notes, not a guarantee of exact finite-sample size. **Do not use an ordinary regression t-test p-value or an ordinary ADF critical value for the estimated levels residuals.** The statistic from `urca::ur.df(type = "none")` may be used to check your residual regression; its printed critical values are not the Engle–Granger reference values. The Module 10 `eg_test()` helper may also be used, but explain the two regressions and report their specifications.

Inspect the **remaining residuals from the residual-ADF regression**, not just the fitted levels deviation, for serial dependence. For Question 1, use $p_e = 1$ as the baseline and $p_e = 4$ as the sensitivity check. For Question 2, choose and state a baseline lag rule and one alternative before comparing the cointegration decisions; justify them using the sample frequency, sample size, and diagnostics. Report both results, including disagreement. The supplied critical value applies only to the stated two-variable intercept/no-trend framework. If that framework is not defensible for a proposed economic pair, explain the limitation rather than applying its critical value to a different specification.

Question 1 — Cointegration and Predictive Direction *(Modules 9–10)*

(a) Define cointegration for two I(1) series. Explain how it differs from high correlation in levels, and why an I(0) series needs different treatment when deciding whether to use the Engle–Granger framework above.

(b) Assess the order of integration of `X`, `Y`, and `Z`. Support your conclusions with plots, appropriately specified ADF and KPSS tests, and evidence from first differences where needed. State and justify the deterministic terms and lag choices in the individual-series tests. Identify which pairs are eligible for the supplied Engle–Granger procedure. Discuss conflicting evidence rather than treating a single p-value as proof.

(c) Apply the supplied Engle–Granger framework to each eligible pair. Report the estimated levels relationship, normalization, n, p_e, usable residual-ADF rows, τ_{EG}, matching critical value, and decision. Include the prescribed lag sensitivity and examine the test regression's remaining serial dependence. Which pairs show evidence of cointegration? Explain why the individual-series classification in (b) matters to that conclusion.

(d) Does the past of `Z` add information for predicting `Y`, **conditional on the specified histories of both `Y` and `X`**? Begin with the one-lag ARDL:

$$Y_t = \alpha + \phi_1 Y_{t-1} + \beta_{X,1} X_{t-1} + \beta_{Z,1} Z_{t-1} + \epsilon_t.$$

State the null hypothesis and estimate the unrestricted and restricted models. The restricted model drops only the lagged-`Z` term: retain the intercept and the same `Y` and `X` lags. Use exactly the same dependent observations in both models. Calculate and report the RSS-based F statistic, its degrees of freedom, p-value, and 5% decision; do not report only regression significance stars.

Then repeat the comparison using **two lags of each of `Y`, `X`, and `Z`**. Test the two `Z` coefficients jointly. Inspect the dynamic regressions' remaining residual dependence and compare the one- and two-lag findings. Keep current `X_t` and current `Z_t` out of these past-only prediction comparisons.

(e) Repeat the one- and two-lag comparisons with **`X` as the outcome**, retaining the same histories of `X` and `Y` and testing only the `Z`-lag block. Report the same statistics and diagnostics. Do the data provide evidence that `Z` Granger-causes `Y`, `X`, both, or neither under these specifications?

(f) Explain what your cointegration and Granger findings say about long-run comovement and predictive direction. Does either finding establish an intervention effect? Why can a well-supported nonrejection differ from a claim that no relationship exists? Your conclusions should follow the observed evidence, including any ambiguity, rather than an assumed answer about simulated data.

F-test convention. Use

$$F = \frac{(RSS_R - RSS_U)/r}{RSS_U/(n_G - k_U)},$$

where r is the number of excluded `Z`-lag coefficients, n_G is the common usable ARDL sample size, and k_U counts all unrestricted regression coefficients, including the intercept. Compare with $F(r, n_G - k_U)$. For this supplied exercise, use this as an **approximate reference for the stated lagged-`Z` exclusion test**. This is not a general instruction to apply ordinary F tests to arbitrary restrictions on integrated level series. Assess the regression's remaining serial dependence and interpret a nonrejection as insufficient evidence, not proof of no predictive contribution.

Question 2 — Cointegration in Your Midterm Data *(Module 10)*

Use **only the original training block** of your frozen midterm dataset. The objective is to identify and assess candidate pairwise relationships, not to produce a rejection.

(a) Identify the series, transformations, training dates, and sample sizes used. Reassess their integration orders using evidence appropriate to the training block. Do not rely on a midterm test that included observations outside that block. Identify every unordered pair eligible for the supplied two-I(1)-series Engle–Granger framework. For each eligible pair, briefly explain what a stable relationship would mean economically. Distinguish economic plausibility from statistical evidence.

(b) Apply the supplied Engle–Granger framework to **all eligible pairs** using one declared normalization per pair. Follow the alignment requirements and use each pair's actual n when calculating its critical value. State your baseline lag rule and one alternative before comparing the decisions, and report both. Use one compact table containing the pair, normalization, training span, n, estimated $\hat{c}$ and $\hat{\kappa}$, p_e, τ_{EG}, critical value, and decision. You may split this into two readable tables rather than shrinking it beyond legibility.

(c) For one eligible pair, reverse the normalization and repeat the baseline test on **the same observations**. Choose that pair before looking for a favorable reverse-normalization result. Examine the baseline test regression's remaining serial dependence, and explain how normalization, lag choice, sample length, or a plausible break affects your confidence in the finding. If the evidence is sensitive, retain that qualification in your conclusion; do not choose whichever specification rejects.

(d) How many eligible pairs show evidence of cointegration under your declared baseline, and which findings remain supported by the checks you reported? Explain the limits of interpreting a collection of pairwise tests, including the chance of isolated rejections when several pairs are examined. The number of rejecting pairs is **not** the number of independent cointegrating relationships in the full system. A reversed normalization is not an additional relationship, and pairwise tests need not detect a relationship that requires three or more series. No system-rank test is required.

A null result is a valid result. You can receive full credit for a careful finding of no evidence of cointegration. If fewer than two series are plausibly I(1), or no pair supports the stated framework and usable sample, explain that conclusion and why the tests in (b)–(c) are not warranted. Do not substitute new series or a new vintage merely to obtain an eligible pair or a rejection. In that case, use (b)–(d) to explain the proposed testing procedure, its unmet requirements, and what your data do and do not permit you to conclude; numerical tests are not required for an ineligible pair.