Seminar 4 Solutions

Giulio Rossetti*

giuliorossetti94.github.io

February 6, 2025

* email: giulio.rossetti.1@wbs.ac.uk

Disclaimer

Full solutions are available on my.wbs. All exercises are examinable material, not just the ones we covered in the seminars.

TS vs CS

Roadmap

Exercise 1

Exercise 2

Exercise 3

Exercise 4

Exercise 1 (Part 1)

Q: As for cross sections, can we assume that time-series observations are independent of each other?

Exercise 1 (Part 1)

Q: As for cross sections, can we assume that time-series observations are independent of each other?

- NO. In a time series setting, the temporal ordering of observations matters.
- Cannot safely assume they are independent, because a typical feature of time series is serial correlation/dependence.
- In a time-series context, the randomness does not come from sampling from a population (as in cross sections), but rather from observing one realization of a stochastic process through time.

Exercise 1 (Part 2)

Q: How would you estimate a multiple linear regression model in a time-series setting?

Exercise 1 (Part 2)

Q: How would you estimate a multiple linear regression model in a time-series setting?

model

$$y_t = \beta_0 + \beta_1 x_{t1} + \beta_2 x_{t2} + \cdots + \beta_k x_{tk} + u_t$$

- In matrix form: $y = X\beta + u$.
- The Ordinary Least Squares (OLS) estimator $\hat{\beta} = (\hat{\beta}_0, \ldots, \hat{\beta}_k)'$ minimizes the sum of squared residuals:

$$\hat{\beta} = \arg\min_{\beta} (y - X\beta)'(y - X\beta) = \arg\min_{\beta} \mathbf{u}'\mathbf{u}.$$

- Equivalently, we look for $\hat{\beta}$ that minimizes $S(\beta)$ (the sum of squared errors).

Exercise 1 (Part 3)

Q: What assumptions do you need to obtain unbiasedness of the OLS estimator in a time-series setting?

Exercise 1 (Part 3)

Q: What assumptions do you need to obtain unbiasedness of the OLS estimator in a time-series setting?

- Finite-sample properties of OLS under classical assumptions:
  - TS-1: Linear in parameters.
  - TS-2: No perfect collinearity among regressors.
  - TS-3: Zero conditional mean, $E[u_t \mid X] = 0$.
- Under these assumptions, $\hat{\beta}$ is an unbiased estimator of β.

Exercise 1 (Part 4)

Q: Is the zero conditional mean assumption more restrictive in a time-series setting than in a cross-sectional setting?

Exercise 1 (Part 4)

Q: Is the zero conditional mean assumption more restrictive in a time-series setting than in a cross-sectional setting?

- YES. **Strict exogeneity** (TS.3) is often questionable because it rules out any feedback from the dependent variable on future values of the explanatory variables.
- Exogeneity: $E[u_t \mid x_t] = 0$, i.e., the error is uncorrelated with regressors at the same period.

$$r_t = \beta_0 + \beta_1 \, \mathrm{MKT}_t + u_t$$

- TS.3 implies $E[u_t \mid \mathrm{MKT}_{t-j}] = 0$, but this may be violated (e.g., MKT_{t-1} could be correlated with u_t).
- In reality, MKT might be pro-cyclical or correlated with consumption, leading to endogeneity.

Exercise 1 (Part 5)

Q: What assumptions are needed for the OLS estimator to be BLUE, and what do we need for valid F- and t-tests?

Exercise 1 (Part 5)

Q: What assumptions are needed for the OLS estimator to be BLUE, and what do we need for valid F- and t-tests?

- For efficiency (BLUE), in addition to TS.1–TS.3, we also need:
  - TS.4 Homoskedasticity: $\mathrm{Var}(u_t \mid X) = \sigma^2$.
  - TS.5 No serial correlation: $\mathrm{Corr}(u_t, u_s) = 0$ for $t \neq s$.
- Under TS.1–TS.5, OLS is BLUE (Best Linear Unbiased Estimator).
- For valid F- and t-tests, we also assume:
  - TS.6 Normality: $u_t \sim N(0, \sigma^2)$, independent of X.
- Then $\hat{\beta}$ has a normal sampling distribution, and the usual F- and t-tests are valid.

Roadmap

Exercise 1

Exercise 2

Exercise 3

Exercise 4

Exercise 2 (Part 1)

model

$$\pi_t = \beta_0 + \beta_1 \,\mathrm{Unemp}_t + \beta_2 \,\mathrm{Unemp}_{t-1} + \beta_3 \,\mathrm{Unemp}_{t-2} + \beta_4 \,\mathrm{Unemp}_{t-3} + u_t$$

Exercise 2 (Part 1)

model

$$\pi_t = \beta_0 + \beta_1 \, \text{Unemp}_t + \beta_2 \, \text{Unemp}_{t-1} + \beta_3 \, \text{Unemp}_{t-2} + \beta_4 \, \text{Unemp}_{t-3} + u_t$$

- The transitory effect from one year ago (i.e., 4 quarters ago) is measured by β_4.
- The transitory effect of a current change in unemployment is given by β_1.
- The persistent effect is measured by the sum of the lag coefficients:

$$\beta_1 + \beta_2 + \beta_3 + \beta_4.$$

Interpretation of the coefficients in an FDL model

Transitory increase in z_t

$$y_t = \alpha_0 + \delta_0 z_t + \delta_1 z_{t-1} + \delta_2 z_{t-2} + u_t \quad \text{(FDL of order two).}$$

- Scenario: For $t < 0$, assume $z_t = c$. At time $t = 0$, z_0 increases to $c + 1$ *just for that period*, and then at $t = 1$, it reverts to c.
- Key equations (setting $u_t = 0$ for simplicity):

$$\begin{aligned}
y_{-1} &= \alpha_0 + \delta_0 c + \delta_1 c + \delta_2 c, \\
y_0 &= \alpha_0 + \delta_0 (c+1) + \delta_1 c + \delta_2 c, \\
y_1 &= \alpha_0 + \delta_0 c + \delta_1 (c+1) + \delta_2 c, \\
y_2 &= \alpha_0 + \delta_0 c + \delta_1 c + \delta_2 (c+1), \\
y_3 &= \alpha_0 + \delta_0 c + \delta_1 c + \delta_2 c.
\end{aligned}$$

Interpretation of the coefficients in an FDL model

- Interpretation:
  - The *immediate* effect on y_0 (from y_{-1}) is δ_0.
  - After one period, $y_1 - y_{-1} = \delta_1$, etc.
  - By $t = 3$, y_3 has returned to its initial level, so the effect of the increase in z_0 is transitory.

Interpretation of the coefficients in an FDL model

Permanent increase in z_t

$$y_t = \alpha_0 + \delta_0 z_t + \delta_1 z_{t-1} + \delta_2 z_{t-2} + u_t.$$

- Scenario: Suppose now that at $t = 0$, z_0 increases from c to $c + 1$ and **stays** at $c + 1$ for all subsequent periods.
- Key equations (still setting $u_t = 0$):

$$\begin{aligned}
y_{-1} &= \alpha_0 + \delta_0 c + \delta_1 c + \delta_2 c, \\
y_0 &= \alpha_0 + \delta_0 (c+1) + \delta_1 c + \delta_2 c, \\
y_1 &= \alpha_0 + \delta_0 (c+1) + \delta_1 (c+1) + \delta_2 c, \\
y_2 &= \alpha_0 + \delta_0 (c+1) + \delta_1 (c+1) + \delta_2 (c+1), \\
y_3 &= \alpha_0 + \delta_0 (c+1) + \delta_1 (c+1) + \delta_2 (c+1), \quad \ldots
\end{aligned}$$

Interpretation of the coefficients in an FDL model

- Long-run effect:
  - For large t, $z_t = c + 1$. Thus y_t stabilizes at $\alpha_0 + (\delta_0 + \delta_1 + \delta_2)(c + 1)$.
  - The cumulative impact of a permanent $+1$ in z is $\delta_0 + \delta_1 + \delta_2$.

Roadmap

Exercise 1

Exercise 2

Exercise 3

Exercise 4

Visual Representation of the Problem

- **Option 1:** Buy a 3-month T-bill at time $t-1$, hold it to t.
  - Its yield, $hy3_{t-1}$, is known at $t-1$.
- **Option 2:** Buy a 6-month T-bill at time $t-1$, sell after 3 months (at t).
  - Its 3-month holding-period yield, $hy6_t$, is *unknown* at $t-1$.
- The Expectations Hypothesis suggests $hy3_{t-1}$ and $hy6_t$ should be the same on average.
- We test this by estimating:

$$hy6_t \;=\; \beta_0 \;+\; \beta_1\, hy3_{t-1} \;+\; u_t$$

and checking if $\beta_1 = 1$.

Visual Representation

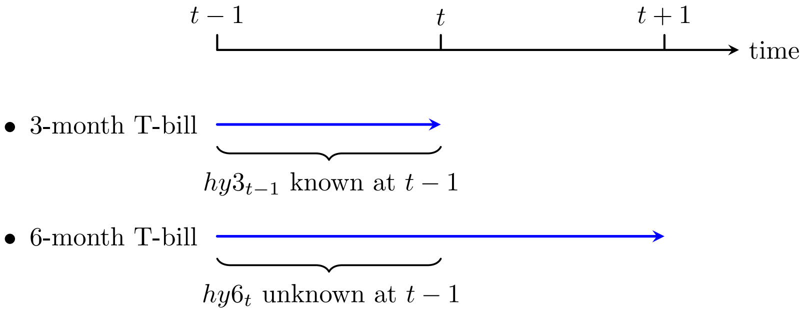

Figure: Visual representation of the problem

Estimation Results

Source	SS	df	MS
Model	84.9875172	1	84.9875172
Residual	13.1294786	121	.108508087
Total	98.1169958	122	.804237671

Number of obs =	123
F(1, 121) =	783.24
Prob > F =	0.0000
R-squared =	0.8662
Adj R-squared =	0.8651
Root MSE =	.32941

hy6	Coef.	Std. Err.	t	P>|t|	[95% Conf. Interval]
hy3_1	1.104309	.0394568	27.99	0.000	
_cons	-.0579439	.0699576	-0.83	0.409	

Figure: Estimation results for the Expectations Hypothesis

- We test the null hypothesis $H_0 : \beta_1 = 1$.

Exercise 3

Q: How do we compute the t-statistic for hypothesis testing on a single parameter $\hat{\beta}_1$?

Exercise 3

Q: How do we compute the t-statistic for hypothesis testing on a single parameter $\hat{\beta}_1$?

- We use the ratio of the estimated parameter minus its hypothesized value over the standard error:

$$t_{\hat{\beta}_1} = \frac{\hat{\beta}_1 - 1}{\mathrm{se}\big(\hat{\beta}_1\big)}.$$

- In this example:

$$\hat{\beta}_1 = 1.1043, \quad \mathrm{se}\big(\hat{\beta}_1\big) = 0.039 \quad \Longrightarrow \quad t_{\hat{\beta}_1} = \frac{1.1043 - 1}{0.039} = 2.67.$$

- $Interpretation:$ The larger $|t_{\hat{\beta}_1}|$ is, the more evidence we have that β_1 differs from 1.

Exercise 3

Q: What is the two-sided rejection rule, and how do we apply it?

Exercise 3

Q: What is the two-sided rejection rule, and how do we apply it?

- For a two-sided null hypothesis $H_0\colon \beta_1 = 1$, we reject H_0 in favor of $H_a\colon \beta_1 \neq 1$ if

$$\left|t_{\hat{\beta}_1}\right| > c,$$

  where c is the critical value from a t-distribution with $T - k - 1$ degrees of freedom.

- At the 1% significance level, $c = 2.62$. Because our computed statistic $t_{\hat{\beta}_1} = 2.67$ is greater than 2.62, we reject H_0 and conclude $\beta_1 \neq 1$ at the 1% level.

Roadmap

Exercise 1

Exercise 2

Exercise 3

Exercise 4

Exercise 4

model

$$\text{Return}_t = \beta_0 + \beta_1 \,\text{Return}_{t-1} + \beta_2 \,\text{Return}^2_{t-1} + u_t, \quad u_t \sim N(0, \sigma^2).$$

Source	SS	df	MS
Model	19.2169743	2	9.60848717
Residual	3051.20782	686	4.4478248
Total	3070.42479	688	4.46282673

Number of obs	= 689
F(2, 686)	= 2.16
Prob > F	= 0.1161
R-squared	= 0.0063
Adj R-squared	= 0.0034
Root MSE	= 2.109

return	Coef.	Std. Err.	t	P>|t|	[95% Conf.	Interval]
return_1	.0485723	.0387224	1.25	0.210	-.0274563	.1246009
ret2	-.009735	.0070296	-1.38	0.167	-.023537	.004067
_cons	.2255486	.087234	2.59	0.010	.0542708	.3968263

Figure: Predictive Model for Stock Returns

Exercise 4

$$E[\text{Return}_t \mid \text{Return}_{t-1}] = E[\text{Return}_t].$$

Exercise 4

$E[\mathrm{Return}_t \mid \mathrm{Return}_{t-1}] = E[\mathrm{Return}_t]$.

- Intuitively, if both β_1 and β_2 are zero, then $E[\mathrm{Return}_t \mid \mathrm{Return}_{t-1}]$ does not depend on Return_{t-1}.
- So we set up the null hypothesis as $H_0\colon \beta_1 = \beta_2 = 0$.
- The F-statistic is about 2.16 with a p-value ≈ 0.116.
- Conclusion: Since the p-value exceeds 0.10, we cannot reject H_0 at the 10% level.
- This suggests that Return_t does not significantly depend on past returns.

Exercise 4

Q: Are weekly stock returns predictable?

- Predicting Return_t based on Return_{t-1} (and Return_{t-1}^2) does not appear promising:
  - The F-statistic is borderline significant at the 10% level.
  - The model explains less than 1% of the variation in Return_t.
- Hence, there is little evidence that weekly stock returns are predictable using only past returns.