

Seminar 6 Solutions

Time Series Regression

Giulio Rossetti*

giuliorossetti94.github.io

February 13, 2026

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

Roadmap

Recap

Part 1: Theory

Exercise 1: Time Series Fundamentals

Part 2: Finite Distributed Lag Model

Exercise 2: FDL Model

Part 3: Expectations Hypothesis

Exercise 3: Expectations Hypothesis

Part 4: Stock Return Predictability

Exercise 4: Stock Return Predictability

Roadmap

Recap

Part 1: Theory

Exercise 1: Time Series Fundamentals

Part 2: Finite Distributed Lag Model

Exercise 2: FDL Model

Part 3: Expectations Hypothesis

Exercise 3: Expectations Hypothesis

Part 4: Stock Return Predictability

Exercise 4: Stock Return Predictability

Disclaimer

Disclaimer

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

Recap

Data Structures in Economics

- *Cross-sectional data* (weeks 1-4): one time period, many units
 - Example: grades of students in a given exam, wages of individuals in a given year;
- *Time series data* (weeks 5-8): one unit, many time periods
 - Example: inflation over time for a given country, price of a stock over time;
- *Panel data* (weeks 9-10): many units, many time periods for the same units
 - Example: grades of several students in several exams, holdings of several mutual funds over several years;
- *Alternative data*
 - Example: text, images, spatial data,..., etc



Roadmap

Recap

Part 1: Theory

Exercise 1: Time Series Fundamentals

Part 2: Finite Distributed Lag Model

Exercise 2: FDL Model

Part 3: Expectations Hypothesis

Exercise 3: Expectations Hypothesis

Part 4: Stock Return Predictability

Exercise 4: Stock Return Predictability

Exercise 1

Cross-Sectional vs. Time-Series

- So far: cross-sectional econometrics: everything about Y_i 
- Now: time series data: everything about Y_t
- Next: panel data: everything about $Y_{i,t}$
- Can we assume that time-series observations are independent of each other?

Exercise 1

Independence in Time Series

NO. In a time series setting the temporal ordering of the observations matters.

- A *time series* is a sequence of observations of a variable (e.g., inflation) over time
- We cannot assume that observations are independent of each other in a time-series setting: time series data is ordered
- A typical feature of time series is **serial correlation/dependence**.

Exercise 1

Serial Correlation/Dependence

- In cross-sectional data we assume that observations are i.i.d (*random sampling*)
- *Intuition*: if you relabel individual A as B and B as A, the joint distribution of the data remains the same
- This is not true in time series data: independence is rare
- If we relabel observations, we change the meaning of the data

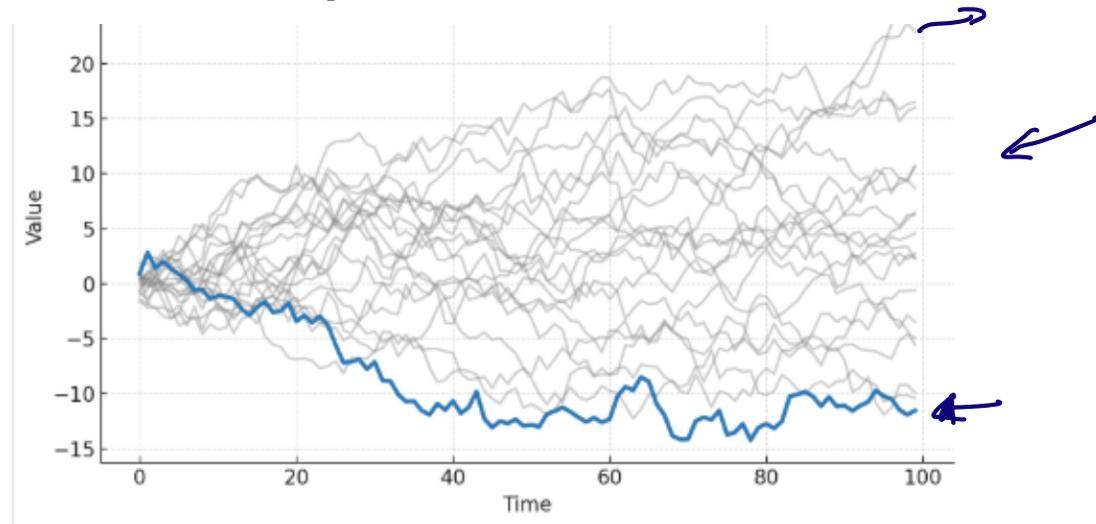
Examples:

- High inflation this month likely means high inflation next month;
- A student doing well on exams will likely do well on the next one;
- Unlike cross-sectional data, the future can depend on the past;

Exercise 1

Stochastic Processes

We model time series as a sequence of **random variables**. A collection of random variables indexed by time is called a **stochastic process**



- We observe only one realization of the process
- But all the grey paths were equally likely to have happened
- We want to make inference about the **whole process**, and we see only one path

Exercise 1

$$y_t = x_t \beta + u$$

OLS in a Time-Series Setting

Let us consider the static model with k explanatory variables:

$$y_t = \beta_0 + \beta_1 x_{1t} + \beta_2 x_{2t} + \cdots + \beta_k x_{kt} + u_t.$$

We can introduce the following matrix notation. Let

$$\mathbf{y}_{(T \times 1)} = \begin{matrix} \text{dep} \\ \left[\begin{array}{c} y_1 \\ y_2 \\ \vdots \\ y_T \end{array} \right] \end{matrix}, \quad \mathbf{X}_{(T \times (k+1))} = \begin{matrix} \text{incl} \\ \left[\begin{array}{cccc} 1 & x_{11} & \cdots & x_{k1} \\ \vdots & \cdots & \cdots & \vdots \\ 1 & x_{1T} & \cdots & x_{kT} \end{array} \right] \end{matrix}$$

$$\boldsymbol{\beta}_{((k+1) \times 1)} = [\beta_0, \beta_1, \dots, \beta_k]', \quad \mathbf{u}_{(T \times 1)} = [u_1, u_2, \dots, u_T]'$$

Exercise 1

OLS Estimator

Then, we can rewrite the MLR model as

$$\mathbf{y} = \mathbf{X}\boldsymbol{\beta} + \mathbf{u}.$$

- The Ordinary Least Squares (OLS) estimator $\hat{\boldsymbol{\beta}} = (\hat{\beta}_0, \dots, \hat{\beta}_k)'$ minimizes the sum of squared residuals:

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

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

Exercise 1

$$E[\hat{\beta}] = \beta$$

Unbiasedness Assumptions

Finite sample properties of OLS under classical assumptions

- TS.1 Linear in the parameters. ←
- TS.2 No perfect collinearity (or no multi-collinearity). ←
- TS.3 Zero conditional mean ←

STRICT EXOGENEITY

$$E[u_t | \mathbf{X}] = 0, \quad E[u_t | x_t, x_{t-1}, x_{t-2}, x_{t+1}, x_{t+2}, \dots] = 0$$

- Under these assumptions, $\hat{\beta}$ is an unbiased estimator of β .

Large samples

$$E[u_t | x_t] = 0$$

Exercise 1

Strict Exogeneity

YES. Strict exogeneity: rules out any possible feedback from the dependent variable on future values of the explanatory variables.

Notice the difference between strict exogeneity and exogeneity!

- Exogeneity: $E[u_t|x_t] = 0$, error term uncorrelated with the regressors at the same period

CAPM: $y_t = \beta_0 + \beta_1 MKT_t + u_t$. $\rightarrow c_t$

- TS.3 implies that $E[u_t|MKT_{1:T}] = 0$. \leftarrow
- error terms contain factors such as consumption
- $MKT_{t-1} \downarrow \rightarrow u_t \downarrow$ as MKT is pro-cyclical and might be correlated with lower consumption.

Exercise 1

Strict Exogeneity – CAPM Example

- The error terms contain factors such as investors' expectations of the business cycle, unexpected inflation, oil shocks.
- A linear relationship might be restrictive, but it should be a good approximation.

TS.3 implies that $E[u_t | MKT_{1:T}] = 0$. This could be easily violated. For example:

- $MKT_{t-1} \downarrow \rightarrow u_t \downarrow$ as MKT is pro-cyclical and might be correlated with lower consumption expenditures and flight-to-quality.
- $u_{t-1} \downarrow \rightarrow MKT_t \downarrow$ as Utils might represents a relevant fraction of MKT.

Exercise 1

BLUE

Efficiency is obtained if two additional assumptions are satisfied

- TS.4 Homoskedasticity

$$\text{Var}(u_t|\mathbf{X}) = \text{Var}(u_t) = \sigma^2.$$

- TS.5 No serial correlation

$$\text{Corr}(u_t, u_s) = 0, \quad t \neq s.$$

Under assumptions TS.1 through TS.5, the OLS estimator is BLUE

Exercise 1

Valid Inference

$$\begin{array}{l} \text{TS6} \Rightarrow \text{TS5} \\ \Rightarrow \text{TS4} \end{array}$$

Assumption TS.6 Normality



$$u_t \sim N(0, \sigma^2) \quad \text{independent of } \mathbf{X}.$$

Under assumptions TS.1 through TS.6, the OLS estimator has the usual normal distribution (conditional on \mathbf{X}). The usual F - and t -tests are valid.

Roadmap

Recap

Part 1: Theory

Exercise 1: Time Series Fundamentals

Part 2: Finite Distributed Lag Model

Exercise 2: FDL Model

Part 3: Expectations Hypothesis

Exercise 3: Expectations Hypothesis

Part 4: Stock Return Predictability

Exercise 4: Stock Return Predictability

Exercise 2

Phillips Curve with Distributed Lags

Model:

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

Exercise 2

Phillips Curve with Distributed Lags

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:

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

Exercise 2

FDL – Transitory Increase

Model:

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

- 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):

$$y_{-1} = \alpha_0 + \delta_0 c + \delta_1 c + \delta_2 c,$$

$$y_0 = \alpha_0 + \delta_0(c + 1) + \delta_1 c + \delta_2 c,$$

$$y_1 = \alpha_0 + \delta_0 c + \delta_1(c + 1) + \delta_2 c,$$

$$y_2 = \alpha_0 + \delta_0 c + \delta_1 c + \delta_2(c + 1),$$

$$y_3 = \alpha_0 + \delta_0 c + \delta_1 c + \delta_2 c.$$

Exercise 2

FDL – Transitory Interpretation

- 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.

Exercise 2

FDL – Permanent Increase

Model:

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

- 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 $u_t = 0$):

$$y_{-1} = \alpha_0 + \delta_0 c + \delta_1 c + \delta_2 c,$$

$$y_0 = \alpha_0 + \delta_0(c + 1) + \delta_1 c + \delta_2 c,$$

$$y_1 = \alpha_0 + \delta_0(c + 1) + \delta_1(c + 1) + \delta_2 c,$$

$$y_2 = \alpha_0 + \delta_0(c + 1) + \delta_1(c + 1) + \delta_2(c + 1),$$

$$y_3 = \alpha_0 + \delta_0(c + 1) + \delta_1(c + 1) + \delta_2(c + 1), \quad \dots$$

Exercise 2

FDL – Long-Run Effect

- 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

Recap

Part 1: Theory

Exercise 1: Time Series Fundamentals

Part 2: Finite Distributed Lag Model

Exercise 2: FDL Model

Part 3: Expectations Hypothesis

Exercise 3: Expectations Hypothesis

Part 4: Stock Return Predictability

Exercise 4: Stock Return Predictability

Exercise 3

Expectations Hypothesis – Setup

Let $hy6_t$ denote the three-month holding-period yield (in percent) from buying a six-month T-bill at time $(t - 1)$ and selling it at time t (three months later) as a three-month T-bill. Let $hy3_{t-1}$ be the three-month holding-period yield from buying a three-month T-bill at time $(t - 1)$.

At time $(t - 1)$, $hy3_{t-1}$ is known, whereas $hy6_t$ is unknown because $p3_t$ (the price of the three-month T-bill) is unknown at time $(t - 1)$.

Exercise 3

Expectations Hypothesis – 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 $\underline{hy3_{t-1}}$ and $\underline{hy6_t}$ should be the same on average.
- We test this by estimating:

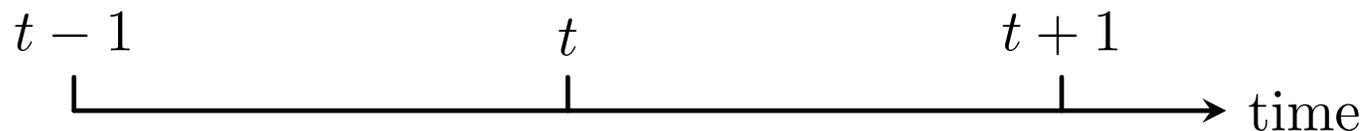
$$hy6_t = \beta_0 + \beta_1 hy3_{t-1} + u_t$$

and checking if $\beta_1 = 1$.

$$E[hy6_t | x] = hy3_{t-1}$$

Exercise 3

Visual Representation



• 3-month T-bill  


 $hy3_{t+1}$ known at $t - 1$

• 6-month T-bill 


 $hy6_t$ unknown at $t - 1$

Exercise 3

Estimation Results

$$H_0: \beta_1 = 1 \quad \leftarrow$$
$$H_1: \beta_1 \neq 1$$

Source	SS	df	MS			
Model	84.9875173	1	84.9875173	Number of obs =	123	
Residual	13.1294786	121	.108508087	F(1, 121) =	783.24	
Total	98.1169958	122	.804227671	Prob > F =	0.0000	
				R-squared =	0.8662	
				Adj R-squared =	0.8651	
				Root MSE =	.22941	

hy6	Coef.	Std. Err.	t	P> t	[95% Conf. Interval]
hy2_1	1.104209	.0294588	27.99	0.000	
_cons	-.0579439	.0699576	-0.83	0.409	

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

Exercise 3

Computing the t -Statistic

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

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

↙ H₀

- In this example:

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

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

Exercise 3

Two-Sided Rejection Rule

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

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

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

Recap

Part 1: Theory

Exercise 1: Time Series Fundamentals

Part 2: Finite Distributed Lag Model

Exercise 2: FDL Model

Part 3: Expectations Hypothesis

Exercise 3: Expectations Hypothesis

Part 4: Stock Return Predictability

Exercise 4: Stock Return Predictability

$$E[r_t | r_{t-1}] = E[r_t]$$

Exercise 4

Predictive Model – Regression Output

Model:

H_0 : no predictability $\Rightarrow \beta_1 = \beta_2 = 0$
 H_1 : predictability $\Rightarrow H_0$ is false

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

Source	SS	df	MS
Model	19.2169749	2	9.60848717
Residual	3051.20762	686	4.4478240
Total	3070.42470	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	Coeff.	Std. Err.	t	P> t	[95% Conf. Interval]
return_1	.0485723	.0287224	1.25	0.210	-.0274563 .1246009
ret2	-.009735	.0070296	-1.38	0.167	-.023527 .004067
_cons	.2255486	.087234	2.59	0.010	.0542708 .3968269

$$Y = \beta_0 + \beta_1 x_1 + \beta_2 x_2 + \beta_3 x_3 + u_t$$

Exercise 4

Hypothesis Testing

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

$$H_0: \beta_1 = \beta_2 = \beta_3 = 0$$

$$H_0: \beta_1 = \beta_2$$

Exercise 4

Hypothesis Testing

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

- Intuitively, if both β_1 and β_2 are zero, then $E[\text{Return}_t \mid \text{Return}_{t-1}]$ does not depend on Return_{t-1} .
- So we set up the null hypothesis as $H_0: \beta_1 = \beta_2 = 0$.
- 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

Are Weekly 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 . $\leftarrow R^2$
- Hence, there is little evidence that weekly stock returns are predictable using only past returns.