



Time Series Analysis and Forecasting

Seminar 2: ARMA Models



Daniel Traian PELE

Bucharest University of Economic Studies

IDA Institute Digital Assets

Blockchain Research Center

AI4EFin Artificial Intelligence for Energy Finance

Romanian Academy, Institute for Economic Forecasting

MSCA Digital Finance

Seminar Outline

- ▣ **Multiple Choice Quiz** – Knowledge check
- ▣ **True/False** – Conceptual checks
- ▣ **Calculation Exercises** – Applied practice
- ▣ **Worked Examples** – Fitting and diagnostics
- ▣ **Discussion Topics** – Practical applications
- ▣ **AI-Assisted Exercise** – Critical thinking

Quiz 1: Lag Operator

Question

What is the result of applying $(1 - L)^2$ to X_t ?

Answer choices

- (A) $X_t - X_{t-1}$
- (B) $X_t - 2X_{t-1} + X_{t-2}$
- (C) $X_t + X_{t-1} + X_{t-2}$
- (D) $X_t - X_{t-2}$

Answer on next slide...



Quiz 1: Answer

Answer: $B - X_t - 2X_{t-1} + X_{t-2}$

▣ **Question:** What is the result of applying $(1 - L)^2$ to X_t ?

Answer choices

- (A) $X_t - X_{t-1}$ ✗
- (B) $X_t - 2X_{t-1} + X_{t-2}$ ✓
- (C) $X_t + X_{t-1} + X_{t-2}$ ✗
- (D) $X_t - X_{t-2}$ ✗

▣ $(1 - L)^2 X_t = (1 - 2L + L^2)X_t = X_t - 2X_{t-1} + X_{t-2}$ — the **second difference** of X_t

Quiz 2: AR(1) Stationarity

Question

For which value of ϕ is the AR(1) process $X_t = 0.5 + \phi X_{t-1} + \varepsilon_t$ stationary?

Answer choices

- (A) $\phi = 1.2$
- (B) $\phi = 1.0$
- (C) $\phi = -0.8$
- (D) $\phi = -1.5$

Answer on next slide...

Quiz 2: Answer

Answer: C – $\phi = -0.8$ (Stationary)

▣ **Question:** For which value of ϕ is the AR(1) process stationary?

Answer choices

- (A) $\phi = 1.2$ ✗
- (B) $\phi = 1.0$ ✗
- (C) $\phi = -0.8$ ✓
- (D) $\phi = -1.5$ ✗

▣ **AR(1) stationarity condition:** $|\phi| < 1$

▣ A: $|1.2| = 1.2 > 1$ B: $|1.0| = 1$ (unit root) C: $|-0.8| = 0.8 < 1$ D: $|-1.5| = 1.5 > 1$

Quiz 3: ACF Pattern

Question

You observe the following ACF pattern: significant spike at lag 1, then all lags within confidence bands. PACF shows gradual decay. What model is suggested?

Answer choices

- (A) AR(1)
- (B) MA(1)
- (C) ARMA(1,1)
- (D) White noise

Answer on next slide...

Quiz 3: Answer

Answer: B – MA(1)

▣ **Question:** ACF cuts off at lag 1, PACF decays. What model?

Answer choices

- (A) AR(1) ✗
- (B) **MA(1)** ✓
- (C) ARMA(1,1) ✗
- (D) White noise ✗

- ▣ ACF cuts off after lag $q \Rightarrow \text{MA}(q)$; PACF cuts off after lag $p \Rightarrow \text{AR}(p)$
- ▣ Here: ACF cuts off at lag 1, PACF decays $\Rightarrow$ **MA(1)**

Quiz 4: MA Invertibility

Question

For the MA(1) process $X_t = \varepsilon_t + 1.5\varepsilon_{t-1}$, is the process invertible?

Answer choices

- (A) Yes, because MA processes are always invertible
- (B) Yes, because $1.5 > 0$
- (C) No, because $|\theta| = 1.5 > 1$
- (D) No, because MA processes are never invertible

Answer on next slide...

Quiz 4: Answer

Answer: C – Not invertible ($|\theta| = 1.5 > 1$)

▣ **Question:** Is the MA(1) process $X_t = \varepsilon_t + 1.5\varepsilon_{t-1}$ invertible?

Answer choices

- (A) Yes, because MA processes are always invertible ✗
- (B) Yes, because $1.5 > 0$ ✗
- (C) **No, because $|\theta| = 1.5 > 1$ ✓**
- (D) No, because MA processes are never invertible ✗

▣ **MA(1) invertibility:** Requires $|\theta| < 1$

▣ The root of $\theta(z) = 1 + \theta z = 0$ must be outside the unit circle. Here: $z = -1/1.5 = -0.67$ is **inside** $\Rightarrow$ **Not invertible**

Quiz 5: ARMA Representation

Question

The compact form $\phi(L)X_t = \theta(L)\varepsilon_t$ represents which model?

Answer choices

- (A) Pure AR model
- (B) Pure MA model
- (C) ARMA model
- (D) None of the above

Answer on next slide...

Quiz 5: Answer

Answer: C – ARMA model

▣ **Question:** The compact form $\phi(L)X_t = \theta(L)\varepsilon_t$ represents which model?

Answer choices

(A) Pure AR model ✗

(B) Pure MA model ✗

(C) **ARMA model** ✓

(D) None of the above ✗

▣ $\phi(L) = 1 - \phi_1 L - \dots - \phi_p L^p; \quad \theta(L) = 1 + \theta_1 L + \dots + \theta_q L^q$

▣ Special cases: $\theta(L) = 1 \Rightarrow$ Pure AR; $\phi(L) = 1 \Rightarrow$ Pure MA

Quiz 6: Information Criteria

Question

When comparing ARMA(1,1) and ARMA(2,1) using BIC, which statement is correct?

Answer choices

- (A) Lower BIC always means better forecasts
- (B) BIC penalizes complexity less than AIC
- (C) The model with lower BIC is preferred
- (D) BIC can only compare models with the same number of parameters

Answer on next slide...

Quiz 6: Answer

Answer: C – The model with lower BIC is preferred

▣ **Question:** When comparing ARMA(1,1) and ARMA(2,1) using BIC, which statement is correct?

Answer choices

- (A) Lower BIC always means better forecasts ✗
- (B) BIC penalizes complexity less than AIC ✗
- (C) **The model with lower BIC is preferred** ✓
- (D) BIC can only compare models with the same number of parameters ✗

▣ $AIC = -2 \ln(\hat{L}) + 2k; \quad BIC = -2 \ln(\hat{L}) + k \ln(n)$

▣ BIC penalizes complexity **more** than AIC (for $n > 7$) $\Rightarrow$ BIC favors simpler models

Quiz 7: Ljung-Box Test

Question

After fitting an ARMA(2,1) model, you run the Ljung-Box test on residuals and get $p\text{-value} = 0.02$. What do you conclude?

Answer choices

- (A) The model is adequate
- (B) Residuals are white noise
- (C) There is significant autocorrelation in residuals
- (D) The model has too many parameters

Answer on next slide...

Quiz 7: Answer

Answer: C – Significant autocorrelation in residuals

▣ **Question:** Ljung-Box test on ARMA(2,1) residuals gives $p\text{-value} = 0.02$. What do you conclude?

Answer choices

- (A) The model is adequate ✗
- (B) Residuals are white noise ✗
- (C) **There is significant autocorrelation in residuals** ✓
- (D) The model has too many parameters ✗

▣ **Ljung-Box test:** H_0 : Residuals are white noise; H_1 : Autocorrelation present

▣ $p\text{-value} = 0.02 < 0.05 \Rightarrow$ **Reject H_0** . The model is **inadequate** — try other orders

Quiz 8: Forecasting

Question

For an AR(1) model with $\phi = 0.6$ and mean $\mu = 10$, what happens to forecasts as horizon $h \rightarrow \infty$?

Answer choices

- (A) Forecasts grow without bound
- (B) Forecasts converge to 0
- (C) Forecasts converge to $\mu = 10$
- (D) Forecasts oscillate forever

Answer on next slide...

Quiz 8: Answer

Answer: C – Forecasts converge to $\mu = 10$

▣ **Question:** AR(1) with $\phi = 0.6$, $\mu = 10$. What happens as $h \rightarrow \infty$?

Answer choices

- (A) Forecasts grow without bound ✗
- (B) Forecasts converge to 0 ✗
- (C) **Forecasts converge to $\mu = 10$** ✓
- (D) Forecasts oscillate forever ✗

▣ **AR(1) forecast formula:** $\hat{X}_{n+h|n} = \mu + \phi^h(X_n - \mu)$

▣ Since $|\phi| = 0.6 < 1$: $\lim_{h \rightarrow \infty} \phi^h = 0 \Rightarrow$ Forecasts converge to μ . **Mean reversion!**

Quiz 9: AR(2) Roots

Question

An AR(2) process has characteristic roots (of $\phi(z) = 1 - \phi_1 z - \phi_2 z^2 = 0$) $z_1 = 0.8$ and $z_2 = -0.5$. Is it stationary?

Answer choices

- (A) Yes, because both roots are inside the unit circle
- (B) No, because one root is negative
- (C) No, because roots must be outside the unit circle
- (D) Cannot determine without more information

Answer on next slide...

Quiz 9: Answer

Answer: C – Roots must be outside the unit circle

▣ **Question:** AR(2) has roots $z_1 = 0.8$ and $z_2 = -0.5$. Is it stationary?

Answer choices

- (A) Yes, because both roots are inside the unit circle ✗
- (B) No, because one root is negative ✗
- (C) **No, because roots must be outside the unit circle ✓**
- (D) Cannot determine without more information ✗

▣ **Stationarity condition:** The roots of $\phi(z) = 0$ must be **outside** the unit circle ($|z| > 1$)

▣ Here: $|z_1| = 0.8 < 1$ ✗; $|z_2| = 0.5 < 1$ ✗. Both inside $\Rightarrow$ **Non-stationary**

Quiz 10: MA(q) Properties

Question

For an MA(2) process, the ACF:

Answer choices

- (A) Decays exponentially
- (B) Cuts off after lag 2
- (C) Cuts off after lag 1
- (D) Never cuts off

Answer on next slide...

Quiz 10: Answer

Answer: B – Cuts off after lag 2

▣ **Question:** For an MA(2) process, the ACF:

Answer choices

- (A) Decays exponentially ✗
- (B) **Cuts off after lag 2** ✓
- (C) Cuts off after lag 1 ✗
- (D) Never cuts off ✗

▣ **ACF property for MA(q):** $\rho(h) = 0$ for $h > q$

▣ MA(1): cuts off after lag 1; MA(2): cuts off after lag 2; MA(q): cuts off after lag q — key identification feature!

True or False? — Questions

1. An AR(2) process is stationary.
 2. MA processes are always stationary.
 3. The PACF of an AR(2) process is zero for all lags greater than 2.
 4. If AIC selects ARMA(2,1) and BIC selects ARMA(1,1), then the true model is ARMA(1,1).
 5. Confidence intervals for the parameters of a fitted model are always symmetric.
 6. The Yule-Walker equations can be used to estimate the parameters of an AR process.
-

True or False? — Answers

-
-
1. An AR(2) process can exhibit pseudocyclical behavior. **True**
 2. MA processes require a seasonal differencing. **False**
 3. The PACF of an AR(p) process has a significant peak at lag p. **True**
 4. If AIC selects ARMA(2,1) and BIC selects ARMA(1,1), they can be both correct. **True**
 5. Confidence intervals narrow as the sample size increases. **True**
 6. The Yule-Walker equations can be used to estimate the parameters of an AR process. **True**

 TSA_ch2_ar2

Exercise 1: AR(1) Properties

Problem

- ▣ **Data:** $X_t = 2 + 0.7X_{t-1} + \varepsilon_t$, $\varepsilon_t \sim WN(0, 9)$
- ▣ **Calculate:** a) Mean μ ; b) Variance $\gamma(0)$; c) Autocovariance $\gamma(1)$ and $\gamma(2)$; d) Autocorrelation $\rho(1)$ and $\rho(2)$

Solution

- ▣ Given: $c = 2$, $\phi = 0.7$, $\sigma^2 = 9$
- ▣ **a) Mean:** $\mu = \frac{c}{1-\phi} = \frac{2}{0.3} = \mathbf{6.67}$
- ▣ **b) Variance:** $\gamma(0) = \frac{\sigma^2}{1-\phi^2} = \frac{9}{0.51} = \mathbf{17.65}$
- ▣ **c) Autocovariance:** $\gamma(1) = \phi \cdot \gamma(0) = 0.7 \times 17.65 = \mathbf{12.35}$;
 $\gamma(2) = \phi^2 \cdot \gamma(0) = 0.49 \times 17.65 = \mathbf{8.65}$
- ▣ **d) Autocorrelation:** $\rho(1) = \phi = \mathbf{0.7}$, $\rho(2) = \phi^2 = \mathbf{0.49}$

Exercise 2: MA(1) Properties

Problem

- ▣ **Data:** $X_t = 5 + \varepsilon_t - 0.4\varepsilon_{t-1}$, $\varepsilon_t \sim WN(0, 4)$
- ▣ **Calculate:** a) Mean μ ; b) Variance $\gamma(0)$; c) Autocovariance $\gamma(1)$; d) Autocorrelation $\rho(1)$; e) Is this process invertible?

Solution

- ▣ Given: $\mu = 5$, $\theta = -0.4$, $\sigma^2 = 4$
- ▣ **a) Mean:** $\mathbb{E}[X_t] = \mu = 5$
- ▣ **b) Variance:** $\gamma(0) = \sigma^2(1 + \theta^2) = 4(1.16) = 4.64$
- ▣ **c) Autocovariance:** $\gamma(1) = \theta\sigma^2 = -0.4 \times 4 = -1.6$
- ▣ **d) Autocorrelation:** $\rho(1) = \frac{-1.6}{4.64} = -0.345$
- ▣ **e) Invertibility:** $|\theta| = 0.4 < 1 \Rightarrow \text{Yes}$

Exercise 3: Characteristic Roots

Problem

- ▣ **Data:** $X_t = 0.5X_{t-1} + 0.3X_{t-2} + \varepsilon_t$
- ▣ **Calculate:** a) Write the characteristic equation; b) Find the characteristic roots; c) Is this process stationary?

Solution

- ▣ **a) Characteristic equation:** $\phi(z) = 1 - 0.5z - 0.3z^2 = 0$, i.e., $0.3z^2 + 0.5z - 1 = 0$
- ▣ **b) Roots (quadratic formula):** $z = \frac{-0.5 \pm \sqrt{0.25 + 1.2}}{0.6} \Rightarrow z_1 = 1.17, z_2 = -2.84$
- ▣ **c) Stationarity check:** $|z_1| = 1.17 > 1$ ✓; $|z_2| = 2.84 > 1$ ✓. Both outside the unit circle $\Rightarrow$ **Stationary**

Exercise 4: Forecasting

Problem

- ▣ **Data:** $X_t = 3 + 0.8X_{t-1} + \varepsilon_t$, $\sigma^2 = 4$, $X_{100} = 20$
- ▣ **Calculate:** a) 1-step ahead forecast $\hat{X}_{101|100}$; b) 2-step ahead forecast $\hat{X}_{102|100}$; c) Long-run forecast as $h \rightarrow \infty$; d) 95% CI for $\hat{X}_{101|100}$

Solution

- ▣ Given: $c = 3$, $\phi = 0.8$, $\sigma^2 = 4$, $X_{100} = 20$. **Mean:** $\mu = \frac{3}{1-0.8} = 15$
- ▣ **a) One-step:** $\hat{X}_{101|100} = 3 + 0.8 \times 20 = 19$
- ▣ **b) Two-step:** $\hat{X}_{102|100} = 3 + 0.8 \times 19 = 18.2$
- ▣ **c) Long-run:** $\lim_{h \rightarrow \infty} \hat{X}_{100+h|100} = \mu = 15$
- ▣ **d) 95% CI:** $19 \pm 1.96 \times 2 = [15.08, 22.92]$

Python Exercise 1: AR(1) Simulation and Fitting

Task

- ▣ Simulate 300 observations from an AR(1) with $\phi = 0.6$
- ▣ Plot the series and ACF/PACF
- ▣ Fit an AR(1) model and compare $\hat{\phi}$ vs actual ϕ
- ▣ Examine residual diagnostics

Key code

```
from statsmodels.tsa.arima.model import ARIMA
model = ARIMA(x, order=(1, 0, 0)).fit()
print(model.summary())
```

Python Exercise 2: Model Selection

Task

- ▣ Load a time series and check stationarity (ADF test)
- ▣ Compare AIC/BIC for AR(1), MA(1), ARMA(1,1), ARMA(2,1)
- ▣ Select the best model
- ▣ Generate forecasts with confidence intervals

Key functions

- ▣ `adfuller(x)` for stationarity test
- ▣ `model.aic`, `model.bic` for criteria
- ▣ `model.get_forecast(h)` for predictions

Python Exercise 3: Diagnostic Checking

Task

After fitting a model, perform complete diagnostics:

- ▣ Plot residuals over time
- ▣ Plot the ACF of residuals
- ▣ Create a Q-Q plot for normality
- ▣ Run the Ljung-Box test

Key functions

- ▣ `model.resid` for residuals
- ▣ `plot_acf(resid)` for ACF plot
- ▣ `stats.probplot(resid)` for Q-Q plot
- ▣ `acorr_ljungbox(resid, lags=[10])` for test

Discussion 1: Model Selection

Scenario

You are modeling monthly inflation rates. After checking stationarity:

- ▣ ACF: significant at lags 1, 2, 3, then decays
- ▣ PACF: significant at lags 1, 2, then cuts off
- ▣ AIC selects ARMA(2,3)
- ▣ BIC selects AR(2)

Questions

1. What does the ACF/PACF pattern suggest?
2. Why do AIC and BIC disagree?
3. Which model would you choose and why?
4. What additional checks would you perform?

Discussion 2: Forecast Evaluation

Scenario

You fit an ARMA(1,1) model to daily stock returns. The in-sample fit looks good (Ljung-Box $p = 0.45$), but out-of-sample RMSE is worse than a random walk.

Questions

1. Is this surprising? Why or why not?
2. What does this tell us about return predictability?
3. Should you conclude that the ARMA model is useless?
4. What alternatives might you consider?

Hint: Think about the Efficient Market Hypothesis and what ARMA captures vs. volatility clustering.

AI Exercise: Critical Thinking

Context

AI tools can fit ARMA models and generate diagnostics automatically. The critical skill is **evaluating whether the methodology is correct**.

Key questions to ask about any AI-generated ARMA analysis:

1. Did it check stationarity **before** fitting?
2. Is the model order justified by ACF/PACF?
3. Are residuals white noise (Ljung-Box test)?
4. Are the roots inside the unit circle?
5. Is the forecast horizon reasonable for the model?

Warning: AI-generated code may run without errors and look professional. *That does not mean it is correct.*

AI Exercise 1: Critique an AI Analysis

Scenario

You asked an AI: “Fit the best model to this sunspot data.” It returned:

- ▣ Fitted ARMA(4,3) with AIC = 2415.3
- ▣ No stationarity test performed
- ▣ Ljung-Box p-value = 0.03 (reported as “acceptable”)
- ▣ 50-year forecast with tight confidence intervals

Your critique

1. Is ARMA(4,3) over-parameterized? What would BIC suggest?
2. Why is Ljung-Box $p = 0.03$ **not** acceptable at 5% level?
3. Are 50-year forecasts reliable for ARMA models? Why/why not?
4. What is the correct Box-Jenkins methodology that was skipped?

AI Exercise 2: Prompt Refinement for ARMA

Task

Iteratively improve prompts for fitting an AR model to sunspot data.

Round 1 (vague): *"Fit a time series model to sunspots"*

- What did the AI produce? What's missing?

Round 2 (better): *"Test stationarity with ADF, examine ACF/PACF, fit $AR(p)$ using BIC, check residuals with Ljung-Box"*

- Did the AI follow the Box-Jenkins methodology?

Round 3 (expert): *"Follow Box-Jenkins: (1) plot & test stationarity, (2) identify order from ACF/PACF, (3) estimate $AR(2)$, (4) Ljung-Box on residuals, (5) forecast 20 steps with 95% CI"*

- Compare results across all three rounds

AI Exercise 3: Model Selection Competition

Task

Download monthly unemployment data from `statsmodels.datasets`.

Your approach (manual):

- ACF/PACF analysis → candidate models
- Compare AIC/BIC across AR(1), AR(2), MA(1), ARMA(1,1)
- Residual diagnostics for selected model
- Rolling 1-step forecast on last 20 observations

AI approach:

- Ask AI to “find the best ARMA model and forecast”

Compare:

- Which model did each select? Do they agree?
- Compare out-of-sample RMSE
- Did the AI use proper rolling forecasts or just multi-step?
- **Submit:** 1-page reflection on AI strengths and weaknesses

Key Formulas Summary

Concept	Formula
AR(1) mean	$\mu = c/(1 - \phi)$
AR(1) variance	$\gamma(0) = \sigma^2/(1 - \phi^2)$
AR(1) ACF	$\rho(h) = \phi^h$
AR(1) stationarity	$ \phi < 1$
MA(1) variance	$\gamma(0) = \sigma^2(1 + \theta^2)$
MA(1) ACF	$\rho(1) = \theta/(1 + \theta^2), \rho(h) = 0 \text{ for } h > 1$
MA(1) invertibility	$ \theta < 1$
AR(1) forecast	$\hat{X}_{n+h n} = \mu + \phi^h(X_n - \mu)$
Forecast CI	$\hat{X} \pm z_{\alpha/2} \times \sqrt{\text{MSFE}(h)}$
AIC	$-2 \ln(\hat{L}) + 2k$
BIC	$-2 \ln(\hat{L}) + k \ln(n)$

Notation: $\hat{L}$ = maximum of the likelihood function, k = no. of parameters, n = sample size, c = constant, σ^2 = white noise variance

Bibliography I

Fundamental textbooks

- ▣ Hyndman, R.J., & Athanasopoulos, G. (2021). *Forecasting: Principles and Practice*, 3rd ed., OTexts.
- ▣ Shumway, R.H., & Stoffer, D.S. (2017). *Time Series Analysis and Its Applications*, 4th ed., Springer.
- ▣ Brockwell, P.J., & Davis, R.A. (2016). *Introduction to Time Series and Forecasting*, 3rd ed., Springer.

Financial time series

- ▣ Tsay, R.S. (2010). *Analysis of Financial Time Series*, 3rd ed., Wiley.
- ▣ Franke, J., Härdle, W.K., & Hafner, C.M. (2019). *Statistics of Financial Markets*, 4th ed., Springer.

Bibliography II

Modern approaches and Machine Learning

- ▣ Nielsen, A. (2019). *Practical Time Series Analysis*, O'Reilly Media.
- ▣ Petropoulos, F., et al. (2022). *Forecasting: Theory and Practice*, International Journal of Forecasting.
- ▣ Makridakis, S., Spiliotis, E., & Assimakopoulos, V. (2020). The M4 Competition, International Journal of Forecasting.

Online resources and code

- ▣ **Quantlet**: <https://quantlet.com> — Code repository for statistics
- ▣ **Quantinar**: <https://quantinar.com> — Learning platform for quantitative methods
- ▣ **GitHub TSA**: https://github.com/QuantLet/TSA/tree/main/TSA_ch2 — Python code for this seminar

Thank You!

Questions?

Seminar materials are available at: <https://danpele.github.io/Time-Series-Analysis/>



Quantlet



Quantinar