



Time Series Analysis and Forecasting

Seminar 3: ARIMA 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 on ARIMA concepts
- ▣ **True/False** – Conceptual checks
- ▣ **Calculation Exercises** – Applied practice with ARIMA
- ▣ **Worked Examples** – Real-world applications
- ▣ **Real Data Analysis** – GDP case study
- ▣ **AI-Assisted Exercise** – Critical thinking
- ▣ **Summary** – Key takeaways

Quiz 1: Integration Order

Question

A time series Y_t requires two differences to become stationary. What is its order of integration?

Answer choices

- (A) $I(0)$
- (B) $I(1)$
- (C) $I(2)$
- (D) Cannot be determined

Answer on next slide...

Quiz 1: Answer




Answer: $C - I(2)$

Answer choices

- (A) $I(0)$ ✗
- (B) $I(1)$ ✗
- (C) $I(2)$ ✓
- (D) Cannot be determined ✗

- ▣ **Question:** What is the order of integration if two differences are needed?
- ▣ **Definition:** $Y_t \sim I(d)$ if $\Delta^d Y_t$ is stationary but $\Delta^{d-1} Y_t$ is not
- ▣ If $\Delta^2 Y_t = \varepsilon_t$, then ΔY_t still has a unit root
- ▣ **Real-world:** Price levels may be $I(2)$ when inflation itself is non-stationary

Visual: Integrated Processes



$I(0)$: stationary. $I(1)$: one difference needed. $I(2)$: two differences needed to become stationary.

 TSA_ch3_def_integrated


Quiz 2: Random Walk Properties

Question

For a random walk $Y_t = Y_{t-1} + \varepsilon_t$ with $\text{Var}(\varepsilon_t) = \sigma^2$, what is $\text{Var}(Y_t)$?

Answer choices

- (A) $\text{Var}(Y_t) = \sigma^2$ (constant)
- (B) $\text{Var}(Y_t) = t\sigma^2$ (grows linearly with time)
- (C) $\text{Var}(Y_t) = t^2\sigma^2$ (grows quadratically)
- (D) $\text{Var}(Y_t) = \sigma^2/t$ (decreases with time)



Answer on next slide...

 TSA_ch3_rw_variance

Quiz 3: ADF Test Specification

Question

When applying the ADF test to GDP data (which shows a clear upward trend), what specification should be used?

Answer choices

- (A) No constant, no trend
- (B) With constant, no trend
- (C) With constant and trend
- (D) The specification does not matter

Answer on next slide...

Quiz 3: Answer

Answer: C – With constant and trend

Answer choices

- (A) No constant, no trend ✗
- (B) With constant, no trend ✗
- (C) **With constant and trend** ✓
- (D) The specification does not matter ✗

- ▣ **ADF regression with trend:** $\Delta Y_t = \alpha + \beta t + \gamma Y_{t-1} + \sum \delta_j \Delta Y_{t-j} + \varepsilon_t$
- ▣ **No constant:** series with zero mean (rarely used)
- ▣ **With constant:** series with non-zero mean but no visible trend
- ▣ **With constant + trend:** series with visible deterministic trend (GDP, prices)
- ▣ **Warning:** Wrong specification reduces the power of the test!

Quiz 4: ARIMA Notation

Question

What does ARIMA(2,1,1) mean?

Answer choices

- (A) AR(2) on differenced data with MA(1) errors
- (B) AR(1) with 2 differences and MA(1)
- (C) MA(2) with 1 difference and AR(1)
- (D) 2 lags, 1 trend, 1 seasonal component

Answer on next slide...

Quiz 4: Answer


Answer: A – AR(2) on differenced data with MA(1) errors

Answer choices

- (A) AR(2) on differenced data with MA(1) errors ✓
- (B) AR(1) with 2 differences and MA(1) ✗
- (C) MA(2) with 1 difference and AR(1) ✗
- (D) 2 lags, 1 trend, 1 seasonal component ✗

- Question: What does ARIMA(2,1,1) mean?
- ARIMA(p, d, q): $\phi(L)(1-L)^d Y_t = \theta(L)\varepsilon_t$
- ARIMA(2,1,1): $(1 - \phi_1 L - \phi_2 L^2)(1-L)Y_t = (1 + \theta_1 L)\varepsilon_t$
- Interpretation: First difference the series, then fit ARMA(2,1) to ΔY_t

Visual: ARIMA Process



Top: original ARIMA series. Bottom: after differencing, use ACF/PACF to identify AR and MA orders.

Quiz 5: ARIMA Equivalence

Question

The ARIMA(0,1,1) model without a constant, $(1 - L)Y_t = (1 + \theta L)\varepsilon_t$, is equivalent to:

Answer choices

- (A) Simple Exponential Smoothing (SES)
- (B) A stationary AR(1) model
- (C) A pure random walk
- (D) A stationary MA(1) model

Answer on next slide...

Quiz 5: Answer

Answer: A – Simple Exponential Smoothing (SES)

Answer choices

(A) **Simple Exponential Smoothing (SES)** ✓

(B) A stationary AR(1) model ✗

(C) A pure random walk ✗

(D) A stationary MA(1) model ✗

▣ **ARIMA(0,1,1):** $Y_t = Y_{t-1} + \varepsilon_t + \theta\varepsilon_{t-1}$

▣ **SES:** $\hat{Y}_{t+1} = \alpha Y_t + (1 - \alpha)\hat{Y}_t$ with $\alpha = 1 + \theta$

▣ When $\theta = 0$: pure random walk (naive); when $-1 < \theta < 0$: smoothing ($0 < \alpha < 1$)

▣ **Conclusion:** SES is the optimal case of an ARIMA(0,1,1)!

Quiz 6: ADF+KPSS Decision Matrix

Question

ADF fails to reject H_0 ($p = 0.15$) and KPSS fails to reject H_0 ($p = 0.08$). What is the conclusion?

Answer choices

- (A) The series is stationary
- (B) The series has a unit root
- (C) Results are inconclusive — insufficient statistical power
- (D) Both tests are wrong

Answer on next slide...

Quiz 6: Answer

Answer: C – Inconclusive results

Question: ADF fails to reject, KPSS fails to reject. Conclusion?

Answer choices

- (A) The series is stationary ✗
- (B) The series has a unit root ✗
- (C) **Results are inconclusive — insufficient statistical power** ✓
- (D) Both tests are wrong ✗

	ADF fails to rej.	ADF rejects
KPSS fails to rej.	Inconclusive	Stationary
KPSS rejects	Unit root	Inconclusive

☐ **Solutions:** Larger sample, PP or ERS tests, or sequential procedure — difference and re-test.

Quiz 7: Overdifferencing

Question

If $Y_t \sim I(1)$ and we compute $\Delta^2 Y_t$, what happens?




Answer choices

- (A) We get a better stationary series
- (B) We introduce artificial negative autocorrelation
- (C) The variance decreases
- (D) Nothing changes

Answer on next slide...

Quiz 7: Answer

Answer: B – Artificial negative autocorrelation



☐ **Diagnostic:** ACF at lag 1 ≈ -0.5 signals overdifferencing. Reduce d by 1!

TSA_ch3_overdifferencing

Quiz 8: Forecast Variance

Question

For an ARIMA(0,1,0) model (random walk), how does forecast variance behave as horizon h increases?

Answer choices

- (A) Stays constant
- (B) Decreases to zero
- (C) Grows linearly with h
- (D) Converges to a finite limit

Answer on next slide...

Quiz 8: Answer

Answer: C – Grows linearly with h

Answer choices

- (A) Stays constant ✗
- (B) Decreases to zero ✗
- (C) **Grows linearly with h** ✓
- (D) Converges to a finite limit ✗

- ▣ **Random walk forecast:** $\hat{Y}_{T+h|T} = Y_T$ (best forecast is current value)
- ▣ **Forecast error:** $Y_{T+h} - \hat{Y}_{T+h|T} = \sum_{i=1}^h \varepsilon_{T+i}$
- ▣ **Variance:** $\text{Var}(Y_{T+h} - \hat{Y}_{T+h|T}) = h\sigma^2$
- ▣ **95% CI:** $Y_T \pm 1.96\sqrt{h}\sigma$ (widens with $\sqrt{h}$)

Quiz 9: Unit Root Test Power

Question

The ADF test has low power when:

Answer choices

- (A) Sample size is very large
- (B) The true root is close to but not equal to 1
- (C) The series has no trend
- (D) The series is clearly stationary

Answer on next slide...

Quiz 9: Answer

Answer: B – Root close to but not equal to 1

Answer choices

- (A) Sample size is very large ✗
 - (B) **The true root is close to but not equal to 1** ✓
 - (C) The series has no trend ✗
 - (D) The series is clearly stationary ✗
-
- ▣ **Example:** AR(1) with $\phi = 0.95$ vs random walk ($\phi = 1$)
 - ▣ Both have similar ACF patterns (slow decay), but one is stationary!
 - ▣ **Low power:** High probability of Type II error (failing to reject false H_0)
 - ▣ **Solutions:** Larger sample sizes, Phillips-Perron test, panel unit root tests

Quiz 10: ARIMA Model Selection

Question

After differencing once, the ACF shows a spike at lag 1 only, and PACF decays. The appropriate model is:


Answer choices

- (A) ARIMA(1,1,0)
- (B) ARIMA(0,1,1)
- (C) ARIMA(1,1,1)
- (D) ARIMA(0,2,1)

Answer on next slide...

Quiz 10: Answer

Answer: B – ARIMA(0,1,1)



- **Pattern:** ACF cuts off at lag 1, PACF decays $\Rightarrow$ MA(1) for differenced series. Full model: ARIMA(0,1,1) = IMA(1,1)

Quiz 11: Trend Stationarity vs Difference Stationarity

Question

A trend-stationary process is made stationary by:


Answer choices

- (A) Taking first differences
- (B) Removing the deterministic trend via regression
- (C) Taking second differences
- (D) Applying seasonal adjustment

Answer on next slide...

Quiz 11: Answer

Answer: B – Removing deterministic trend via regression



- **Trend-stationary:** Detrend (shocks are temporary). **Difference-stationary:** Difference (shocks are permanent). Wrong treatment affects the model!

Quiz 12: ARIMA Invertibility

Question

ARIMA(0,1,1) with $\theta_1 = 1.2$ is:

Answer choices

- (A) Stationary and invertible
- (B) Non-stationary but invertible
- (C) Non-stationary and non-invertible
- (D) Stationary but non-invertible

Answer on next slide...

Quiz 12: Answer

Answer: C – Non-stationary and non-invertible

Answer choices

- (A) Stationary and invertible ✗
- (B) Non-stationary but invertible ✗
- (C) **Non-stationary and non-invertible** ✓
- (D) Stationary but non-invertible ✗

- ▣ **Stationarity:** $d = 1$ means a unit root $\Rightarrow$ **Non-stationary**
- ▣ **Invertibility:** MA root $z = -1/1.2 = -0.833$ (inside the unit circle)
- ▣ $|\theta_1| = 1.2 > 1 \Rightarrow$ **Non-invertible**
- ▣ **Fix:** Rewrite with $\theta^* = 1/1.2 = 0.833$ and adjust variance

Quiz 13: Spurious Regression

Question

Regressing one random walk on another independent random walk typically shows:

Answer choices

- (A) No significant relationship
- (B) High R^2 and significant t-statistics (spuriously)
- (C) Negative correlation
- (D) Perfect multicollinearity

Answer on next slide...

Quiz 13: Answer

Answer: B – High R^2 and significant t -statistics (spuriously)

Answer choices

- (A) No significant relationship ✗
- (B) High R^2 and significant t -statistics (spuriously) ✓
- (C) Negative correlation ✗
- (D) Perfect multicollinearity ✗

- **Question:** Regressing one random walk on another independent random walk shows?
- **Granger & Newbold (1974):** Spurious regression phenomenon
- High R^2 (often > 0.9), significant t -statistics, very low Durbin-Watson ($\ll 2$)
- **Solutions:** (1) Difference both series, or (2) Test for cointegration

Quiz 14: Long-Run Forecast

Question

The long-run forecast from ARIMA(1,1,0) with $\phi_1 = 0.7$ converges to:

Answer choices

- (A) Zero
- (B) The unconditional mean
- (C) A linear trend extrapolation
- (D) The last observed value

Answer on next slide...

Quiz 14: Answer

Answer: C – A linear trend extrapolation

Answer choices

- (A) Zero ✗
- (B) The unconditional mean ✗
- (C) **A linear trend extrapolation ✓**
- (D) The last observed value ✗

- ▣ **Model:** $(1 - \phi_1 L)(1 - L)Y_t = c + \varepsilon_t$
- ▣ **Long-run forecast:** $\hat{Y}_{T+h} \approx Y_T + h \cdot \frac{c}{1 - \phi_1}$
- ▣ **Stationary ARMA:** forecasts $\rightarrow$ unconditional mean
- ▣ **I(1) without drift:** forecasts $\rightarrow$ last value (flat)
- ▣ **I(1) with drift:** forecasts $\rightarrow$ linear extrapolation

True or False? — Questions

1. An $I(2)$ process requires two differences to be stationary.
 2. The ADF test always includes a constant term.
 3. $ARIMA(0,1,0)$ is another name for a random walk.
 4. Differencing a stationary series makes it “non-stationary”.
 5. The KPSS test has stationarity as the null hypothesis.
 6. $ARIMA$ models can only capture linear dependencies.
-

True or False? — Answers

Statement
1. An $I(2)$ process requires two differences to become stationary.
2. The ADF test always includes a constant term.
3. ARIMA(0,1,0) is another name for a random walk.
4. Differencing a stationary series makes it “more stationary.”
5. The KPSS test has stationarity as the null hypothesis.
6. ARIMA models can only capture linear patterns.

Exercise 1: Unit Root Testing

Problem

- ▣ **Data:** Quarterly GDP data for 80 quarters
- ▣ ADF test (with constant and trend) gives test statistic = -2.85 ; 5% critical value = -3.41
- ▣ **Calculate:** a) Conclusion about stationarity? b) What would you do next?

Solution

- ▣ **a)** Since $-2.85 > -3.41$, we **fail to reject** H_0 . The data appears to have a unit root (non-stationary).
- ▣ **b)** Take the first difference ΔY_t and repeat the ADF test on the differenced series to confirm it is now stationary.

Exercise 2: Model Identification

Problem

- After differencing a time series once, the ACF shows: significant spike at lag 1 ($\rho_1 = 0.4$), all other lags insignificant
- The PACF shows gradual decay
- **Identify:** What ARIMA model is suggested?

Solution

- ACF cuts off after lag 1 $\Rightarrow$ MA(1) component
- PACF decays $\Rightarrow$ Confirms MA structure
- Since we differenced once: $d = 1$
- **Suggested model: ARIMA(0,1,1) or IMA(1,1)**

Exercise 3: ARIMA Equation

Problem

- Write out the full equation for ARIMA(1,1,1):

$$(1 - \phi_1 L)(1 - L)Y_t = c + (1 + \theta_1 L)\varepsilon_t$$

- **Expand:** completely in terms of Y_t , Y_{t-1} , Y_{t-2} , etc.

Solution

- Expanding $(1 - \phi_1 L)(1 - L) = 1 - L - \phi_1 L + \phi_1 L^2 = 1 - (1 + \phi_1)L + \phi_1 L^2$:

$$Y_t - (1 + \phi_1)Y_{t-1} + \phi_1 Y_{t-2} = c + \varepsilon_t + \theta_1 \varepsilon_{t-1}$$

- Or equivalently:

$$Y_t = c + (1 + \phi_1)Y_{t-1} - \phi_1 Y_{t-2} + \varepsilon_t + \theta_1 \varepsilon_{t-1}$$

Exercise 4: Forecast Calculation

Problem

- ▣ **Model:** ARIMA(0,1,1): $\Delta Y_t = \varepsilon_t + 0.3\varepsilon_{t-1}$
- ▣ At time T : $Y_T = 100$, $\hat{\varepsilon}_T = 2$, $\sigma^2 = 4$
- ▣ **Calculate:** a) $\hat{Y}_{T+1|T}$ (one-step forecast); b) $\hat{Y}_{T+2|T}$ (two-step forecast)

Solution

- ▣ a) $\hat{Y}_{T+1|T} = Y_T + 0.3\hat{\varepsilon}_T = 100 + 0.3(2) = 100.6$
- ▣ b) $\hat{Y}_{T+2|T} = \hat{Y}_{T+1|T} + 0.3 \cdot 0 = 100.6 + 0 = 100.6$
(Future shocks $\varepsilon_{T+1}, \varepsilon_{T+2}$ are forecast as 0)

Exercise 5: Confidence Intervals

Problem

- Continuing from Exercise 4, calculate the 95% forecast intervals for $\hat{Y}_{T+1|T}$ and $\hat{Y}_{T+2|T}$
- Recall: $\sigma^2 = 4$, $\theta_1 = 0.3$

Solution

- For IMA(1,1), the $MA(\infty)$ weights are $\psi_0 = 1$, $\psi_j = 1 + \theta_1$ for $j \geq 1$.
- 1-step:** $\text{Var}(e_{T+1}) = \sigma^2 \psi_0^2 = 4$, so $SE = 2$
 $100.6 \pm 1.96(2) = [96.68, 104.52]$
- 2-step:** $\text{Var}(e_{T+2}) = \sigma^2(\psi_0^2 + \psi_1^2) = 4(1 + 1.3^2) = 10.76$, $SE = 3.28$
 $100.6 \pm 1.96(3.28) = [94.17, 107.03]$

Example: Testing for Unit Root in Stock Prices

Scenario

You have daily closing prices for a stock over 500 days. You want to determine if prices follow a random walk.

Step-by-step Approach

1. **Visual inspection:** Plot prices – likely shows trend
2. **ADF test on prices:** Expect to fail to reject H_0 (unit root)
3. **Take log returns:** $r_t = \ln(P_t/P_{t-1}) = \Delta \ln(P_t)$
4. **ADF test on returns:** Should reject H_0 (stationary)
5. **Conclusion:** Log prices are $I(1)$, returns are $I(0)$

Example: Box-Jenkins for Inflation Data

Scenario

Monthly inflation rates for 10 years. Build an ARIMA model.

Workflow

1. **Plot & test:** ADF suggests borderline – try both $d = 0$ and $d = 1$
2. **If $d = 0$:** Fit ARMA models, compare AIC
3. **If $d = 1$:** Examine ACF/PACF of ΔY_t
 - ▶ ACF: spike at lag 1, then cuts off
 - ▶ PACF: decays
 - ▶ $\Rightarrow$ Try ARIMA(0,1,1)
4. **Estimate:** Fit ARIMA(0,1,1), check coefficients
5. **Diagnose:** Ljung-Box on residuals (want $p > 0.05$)
6. **Compare:** AIC of ARIMA(0,1,1) vs ARMA(1,1) on levels

Example: Interpreting Python Output

statsmodels ARIMA Output

```


                        ARIMA Model Results
=====
Dep. Variable:          D.y    No. Observations:   99
Model:                ARIMA(1,1,1)    AIC          285.32
                                   BIC          295.63
=====
               coef    std err          z      P>|z|
-----
const         0.0521    0.048      1.085    0.278
ar.L1         0.4532    0.102      4.443    0.000
ma.L1        -0.2891    0.118     -2.450    0.014
sigma2        1.2340    0.176      7.011    0.000

```

Interpretation

- ▣ AR (0.45) significant, MA (-0.29) significant
- ▣ Constant (0.052) not significant – could set $c = 0$
- ▣ Check: $|\phi_1| < 1$ (stationary), $|\theta_1| < 1$ (invertible) – OK!

Case Study: US Real GDP (1990–2024)




Observations

US Real GDP in billions of 2017 dollars (quarterly). Clear **upward trend**. Drops during recessions (2008–09, 2020). Non-stationary: needs differencing.

 TSA_ch3_gdp_levels


Stationarity Through Differencing



Observations

- **Left:** GDP in levels — clear upward trend (non-stationary)
- **Right:** GDP growth rate = $\Delta \log(Y_t) \times 100$ — stationary, fluctuates around mean ($\approx 0.6\%/quarter$)

ACF/PACF: Levels vs Differenced



Observations

- ▣ **Top row:** ACF/PACF of GDP levels — slow decay $\Rightarrow$ non-stationarity
- ▣ **Bottom row:** ACF/PACF of GDP growth — values within confidence bands
- ▣ A low-order ARIMA model is appropriate

ARIMA Estimation Results: US GDP Growth


Model: ARIMA(1,1,1) on log(GDP)

Parameter	Estimate	Std. Error	z-stat	p-value
ϕ_1 (AR.L1)	0.312	0.185	1.69	0.091
θ_1 (MA.L1)	-0.087	0.203	-0.43	0.668
σ^2	0.00012	-	-	-

Interpretation

- ▣ Low-order ARIMA captures GDP dynamics reasonably well
- ▣ AR(1) coefficient positive – GDP growth shows persistence
- ▣ Alternative: simpler random walk (ARIMA(0,1,0)) often competitive


Forecast: ARIMA vs Actual



Observations

- Blue: historical training data; Green: actual test data
- Red: ARIMA forecasts with 95% CI — CI widens with forecast horizon

Model Diagnostics: Residual Analysis



Observations

- Residuals without systematic patterns over time; approximately normal distribution (histogram, Q-Q)
- ACF of residuals within bounds — no autocorrelation; model adequately captures the data generating process

Discussion: Deterministic vs Stochastic Trends

Key Question

Why is it important to distinguish between deterministic and stochastic trends?

Discussion Points

- ▣ **Wrong treatment consequences:**
 - ▶ Detrending a unit root $\Rightarrow$ spurious stationarity
 - ▶ Differencing a trend-stationary $\Rightarrow$ over differencing
- ▣ **Economic interpretation:**
 - ▶ Deterministic trend: shocks are temporary
 - ▶ Stochastic trend: shocks have permanent effects
- ▣ **Policy implications:**
 - ▶ Does a recession permanently lower GDP, or does the economy return to trend?

Discussion: Model Selection Criteria

Key Question

When should you use AIC vs BIC for ARIMA model selection?

Considerations

- ▣ **AIC:** Minimizes prediction error, may overfit
 - ▶ Better for forecasting
 - ▶ Tends to select larger models
- ▣ **BIC:** Consistent model selection, more parsimonious
 - ▶ Better for identifying “true” model
 - ▶ Penalizes complexity more heavily
- ▣ **Practical advice:** Report both, prefer BIC if they disagree substantially

Discussion: Limitations of ARIMA

Key Question

What are the main limitations of ARIMA models?

Discussion Points

- ▣ **Linearity:** Cannot capture nonlinear dynamics
- ▣ **Constant variance:** Assumes homoskedasticity (no GARCH)
- ▣ **No structural breaks:** Parameters assumed constant
- ▣ **Univariate:** Ignores relationships with other variables
- ▣ **Symmetric:** Treats positive and negative shocks equally
- ▣ **Long-horizon forecasts:** Uncertainty grows rapidly

Extensions

These limitations motivate GARCH (volatility), VAR (multivariate), regime-switching models, etc.

AI Exercise: Critique an AI ARIMA Analysis

Prompt to test in ChatGPT / Claude / Copilot

“Fit the best ARIMA model to this GDP data.”

The AI returned:

- ▣ Fitted ARIMA(3,2,3) with AIC = 1542.7
- ▣ No ADF test performed
- ▣ Ljung-Box p-value = 0.02 (reported as “acceptable”)
- ▣ 30-year forecast with narrow confidence intervals

Exercise:

1. Is ARIMA(3,2,3) over-parameterized? What would BIC suggest?
2. Why is Ljung-Box $p = 0.02$ **not** acceptable at the 5% level?
3. Are 30-year forecasts reliable for ARIMA models? Why?
4. What steps from the Box-Jenkins methodology were skipped?

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

AI Exercise: Prompt Refinement for ARIMA

Task

Iteratively improve prompts for fitting an ARIMA model to GDP data.

Round 1 (vague): *“Fit a time series model to GDP”*

- What did the AI produce? What is missing?

Round 2 (better): *“Test stationarity with ADF and KPSS, difference if needed, examine ACF/PACF, fit ARIMA(p,d,q) 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 ADF+KPSS, (2) differencing, (3) identify orders from ACF/PACF, (4) estimate ARIMA(1,1,1), (5) Ljung-Box on residuals, (6) forecast 8 quarters with 95% CI”*

- Compare results across all three rounds

AI Exercise: Model Selection Competition

Task

Download quarterly US Real GDP data from FRED (series GDPC1).

Your approach (manual):

- ▣ ADF + KPSS tests → differencing
- ▣ ACF/PACF → candidate models
- ▣ AIC/BIC: ARIMA(0,1,0), (1,1,0), (0,1,1), (1,1,1)
- ▣ Residual diagnostics + rolling 1-step forecast

AI approach:

- ▣ Ask the AI: “find the best ARIMA and make forecasts”

Compare:

- ▣ What model did each select? Compare RMSE
- ▣ Rolling vs multi-step forecasts?
- ▣ **Submit:** 1-page reflection on AI

Summary: Chapter 3

Concept	Formula
Random walk	$Y_t = Y_{t-1} + \varepsilon_t$
Random walk variance	$\text{Var}(Y_t) = t\sigma^2$
ARIMA(p, d, q)	$\phi(L)(1-L)^d Y_t = \theta(L)\varepsilon_t$
First difference	$\Delta Y_t = Y_t - Y_{t-1} = (1-L)Y_t$
Second difference	$\Delta^2 Y_t = Y_t - 2Y_{t-1} + Y_{t-2}$
ADF regression	$\Delta Y_t = \alpha + \gamma Y_{t-1} + \sum \delta_j \Delta Y_{t-j} + \varepsilon_t$
ADF null	$H_0 : \gamma = 0 \text{ (unit root)}$
RW forecast	$\hat{Y}_{T+h T} = Y_T$
RW forecast CI	$Y_T \pm z_{\alpha/2} \sqrt{h} \sigma$
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, σ^2 = white noise variance

Questions?

Bibliography I

Time Series Fundamentals

- ▣ 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> — Quantitative methods learning platform
- ▣ **GitHub TSA**: https://github.com/QuantLet/TSA/tree/main/TSA_ch3 — Python code for this seminar

Thank You!

Questions?

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



Quantlet



Quantinar