Time-series,
Spring, 2026

Stationarity – Unit Root Tests – Differencing – AR / MA / ARMA models

Faculty of DS & AI
Spring semester, 2026

Trong-Nghia Nguyen

Business AI Lab

Content

- Stationarity
- Unit Root Test
- Differencing
- AR / MA / ARMA Models

Content

- Stationarity
- Unit Root Test
- Differencing
- AR / MA / ARMA Models

Stationarity

Definition

- Stationarity refers to the **stability** of a sequence **over time**.
- A series is called stationary if its **statistical characteristics** do not change over time.
- Most **classic time-series models** (AR, MA, ARMA, ARIMA) assume the data is **stationary**.

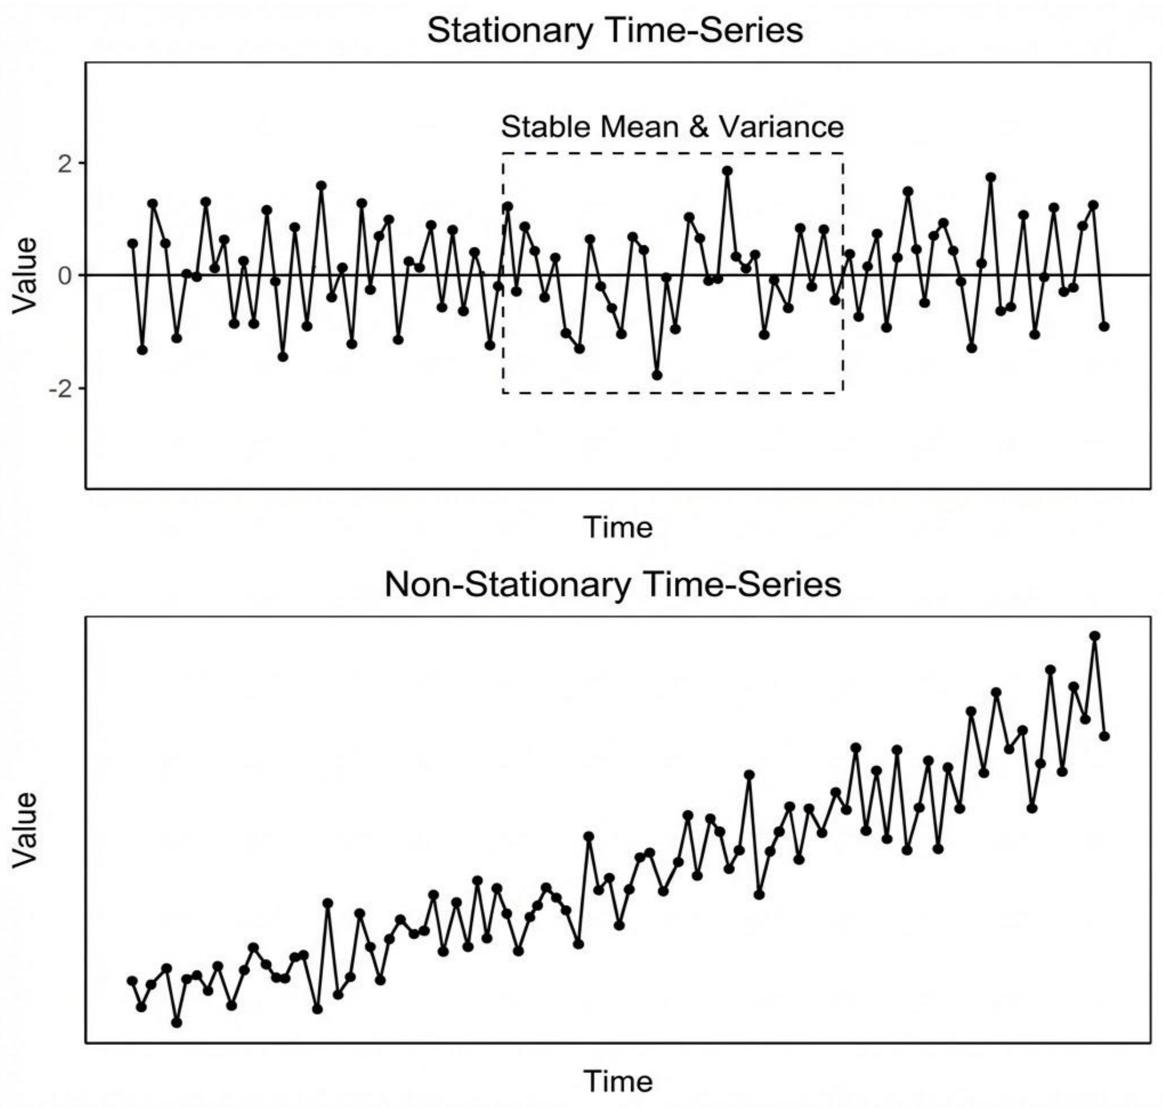

Stationarity

Two common notions of stationarity:

- Strict Stationarity:
  - The entire distribution of the series is invariant over time
- Weak (Covariance) Stationarity
  - Mean is constant over time
  - Variance is constant over time
  - Autocovariance depends only on lag, not on time
- Most time-series models assume **weak stationarity**

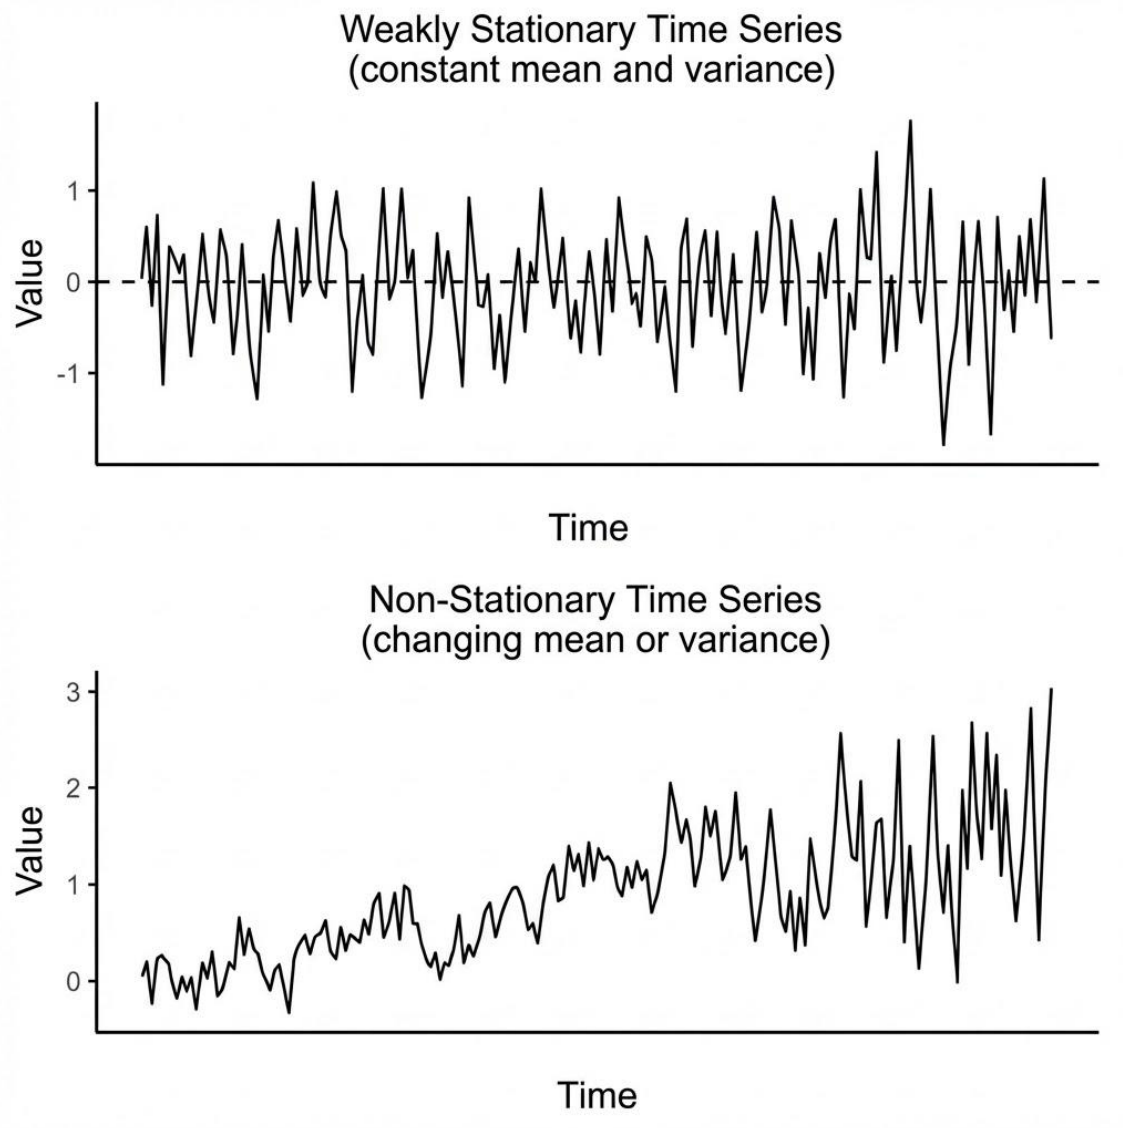

Stationarity

Two common notions of stationarity:

- In weak Stationarity:
  - Mean is constant over time
  - Variance is constant over time
  - Autocovariance depends **only on lag**, **not on time**

$$\mathbb{E}[X_t] = \mu, \quad \forall t$$

$$\mathrm{Var}(X_t) = \sigma^2, \quad \forall t$$

$$\mathrm{Cov}(X_t, X_{t+h}) = \gamma(h)$$

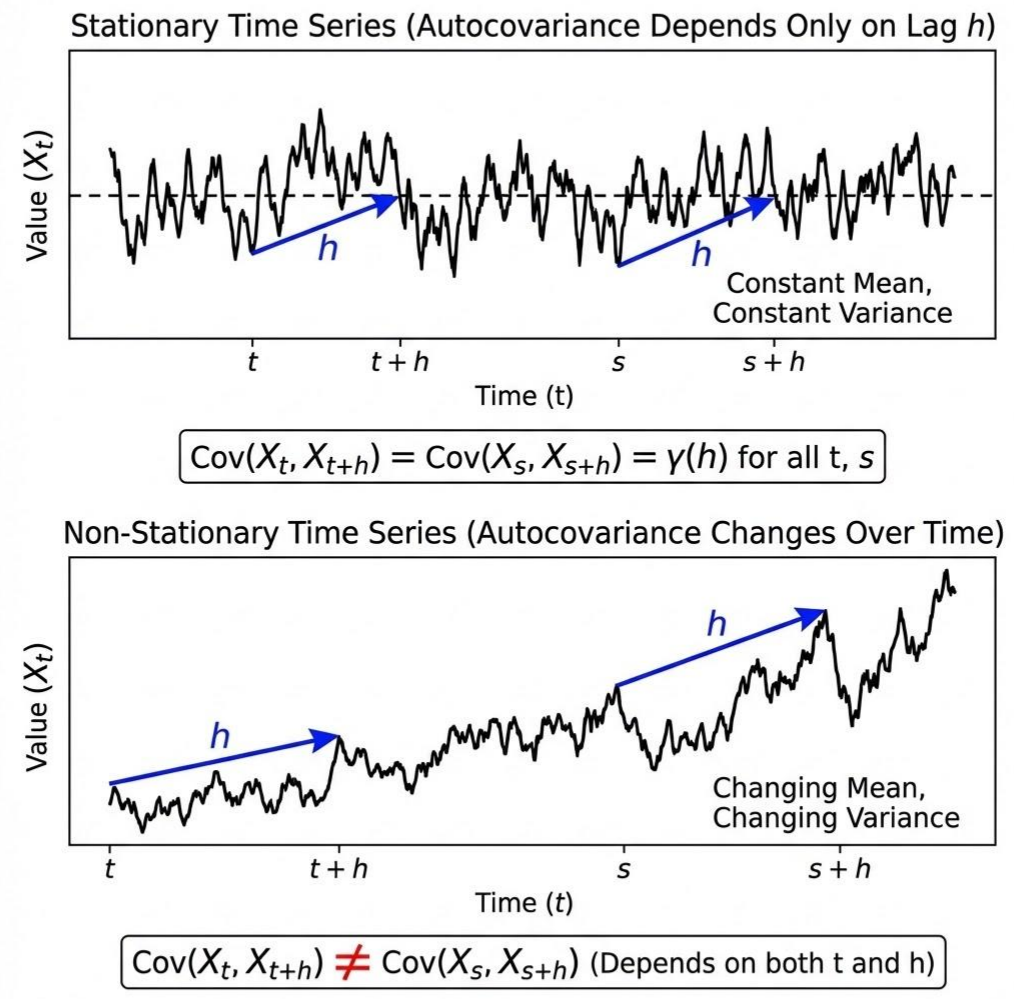

Stationarity

Two common notions of stationarity:

Stationary Time Series

- Model: $x_t = \varepsilon_t$
- Constant mean: $\mu = 0$
- Constant variance: $\sigma^2 = 1$
- No trend over time

Non-Stationary Time Series

- Model: $x_t = 0.05t + \varepsilon_t$
- Mean changes over time
- Presence of deterministic trend
- Violates stationarity assumption

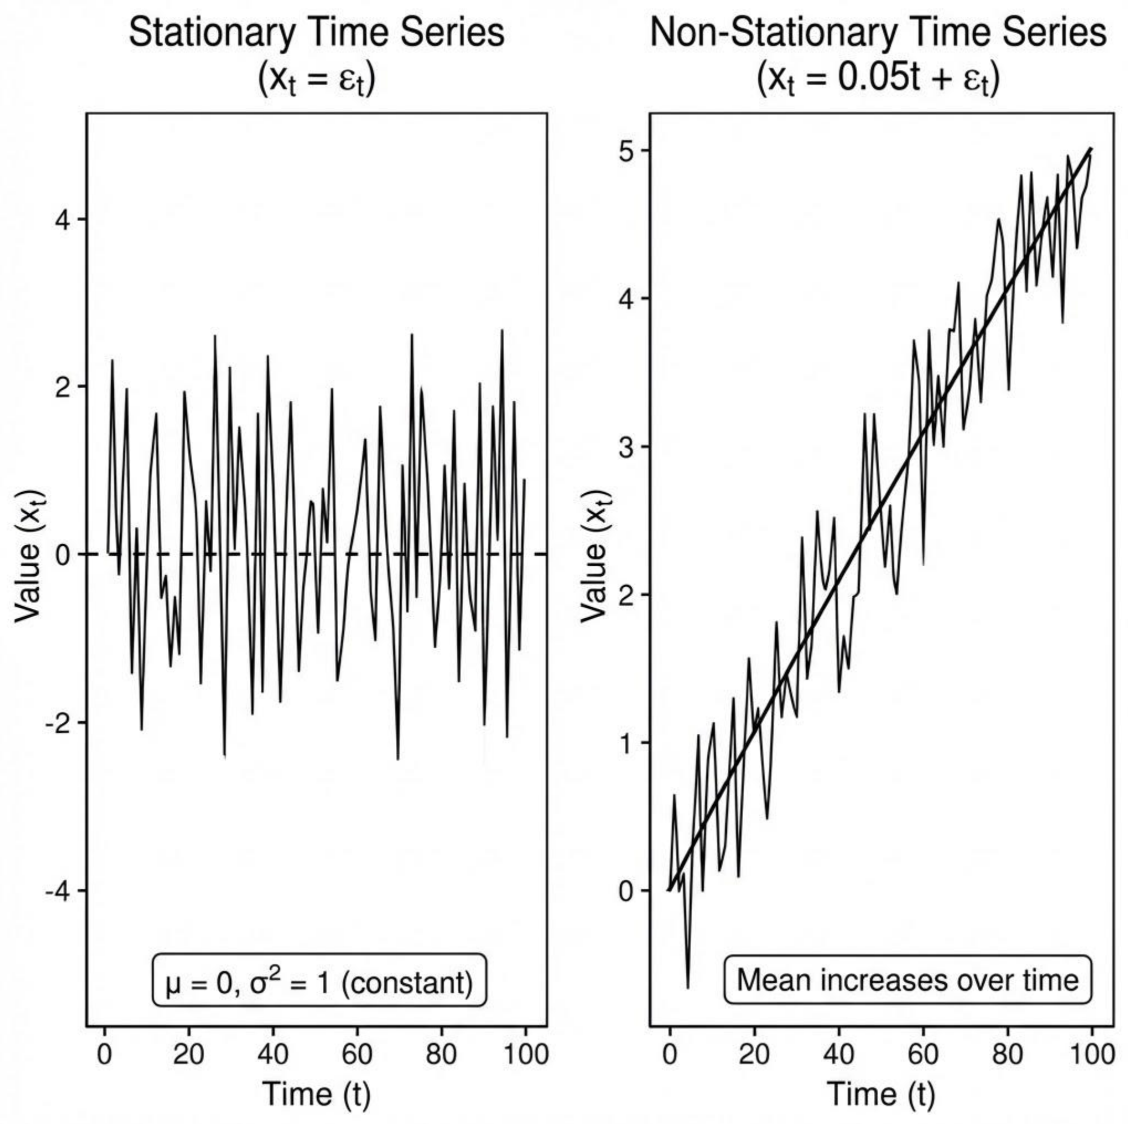

Stationarity

Two common notions of stationarity:

- Based on the figure, which series is stationary and which is non-stationary?
- For each series, comment on:
    - Mean behavior over time
    - Variance behavior over time
- Propose **possible mathematical model** for each series.
- Which series violates the stationarity assumption required by AR / MA models? Explain why.

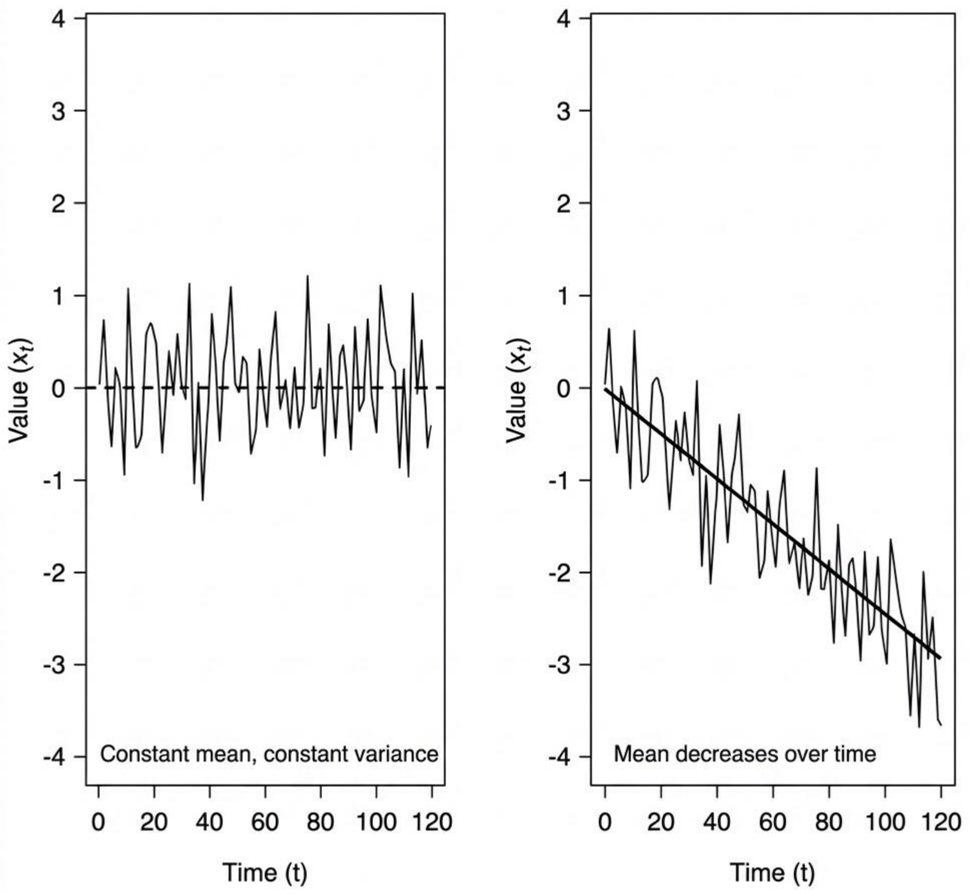

Stationarity

Two common notions of stationarity:

Question 1

Which series is stationary and which is non-stationary?

- **Series A**: Stationary
- **Series B**: Non-stationary

Question 2

Mean and variance behavior

Series A

- Mean: constant over time (≈ 0)
- Variance: constant
- Fluctuations are symmetric and stable

Series B

- Mean: changes over time (downward trend)
- Variance: approximately constant
- Presence of deterministic trend

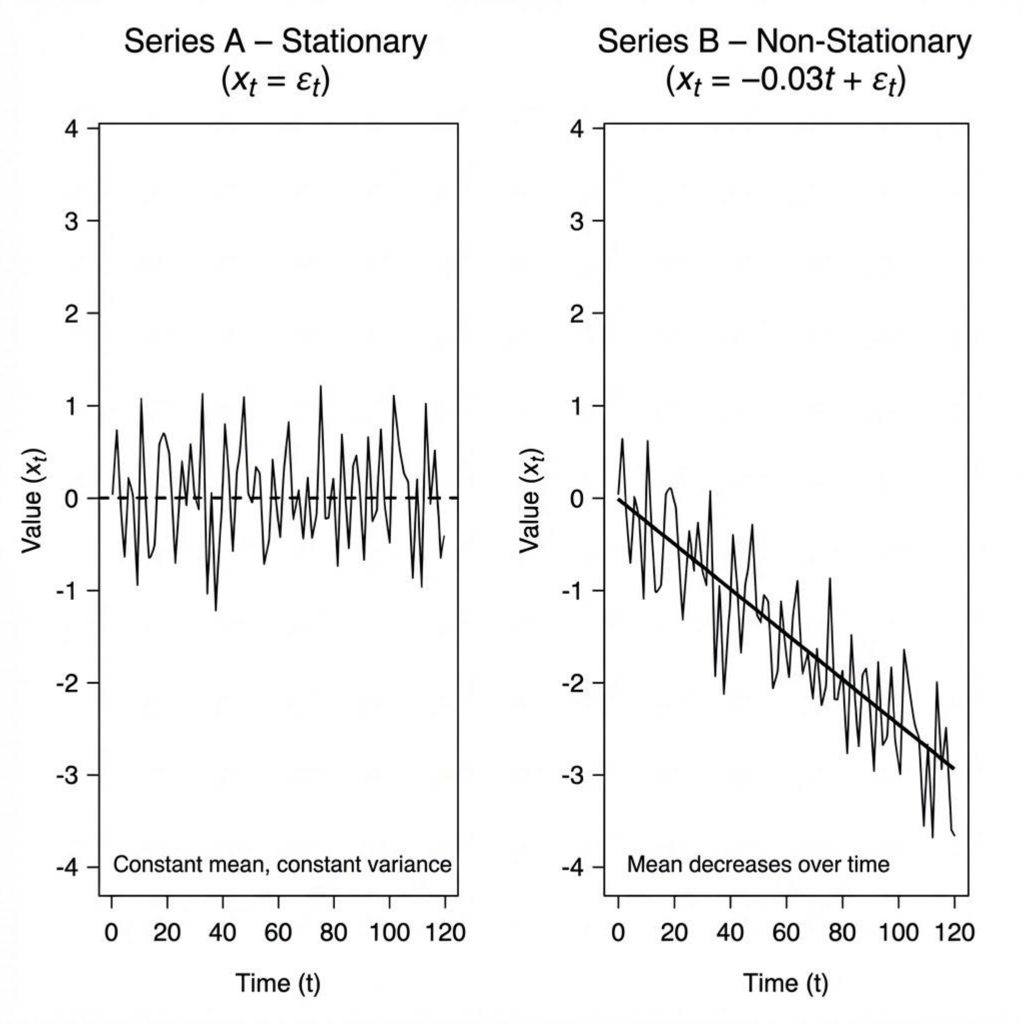

Stationarity

Two common notions of stationarity:

Question 3

Possible mathematical models

- Series A:

$$x_t = \varepsilon_t, \quad \varepsilon_t \sim \mathcal{N}(0,1)$$

- Series B:

$$x_t = -0.03t + \varepsilon_t, \quad \varepsilon_t \sim \mathcal{N}(0,1)$$

Question 4

Which series violates stationarity and why?

- Series B violates stationarity
- Because:

$$\mathbb{E}[x_t] = -0.03t$$

depends explicitly on time

➡️ Mean is not constant → non-stationary

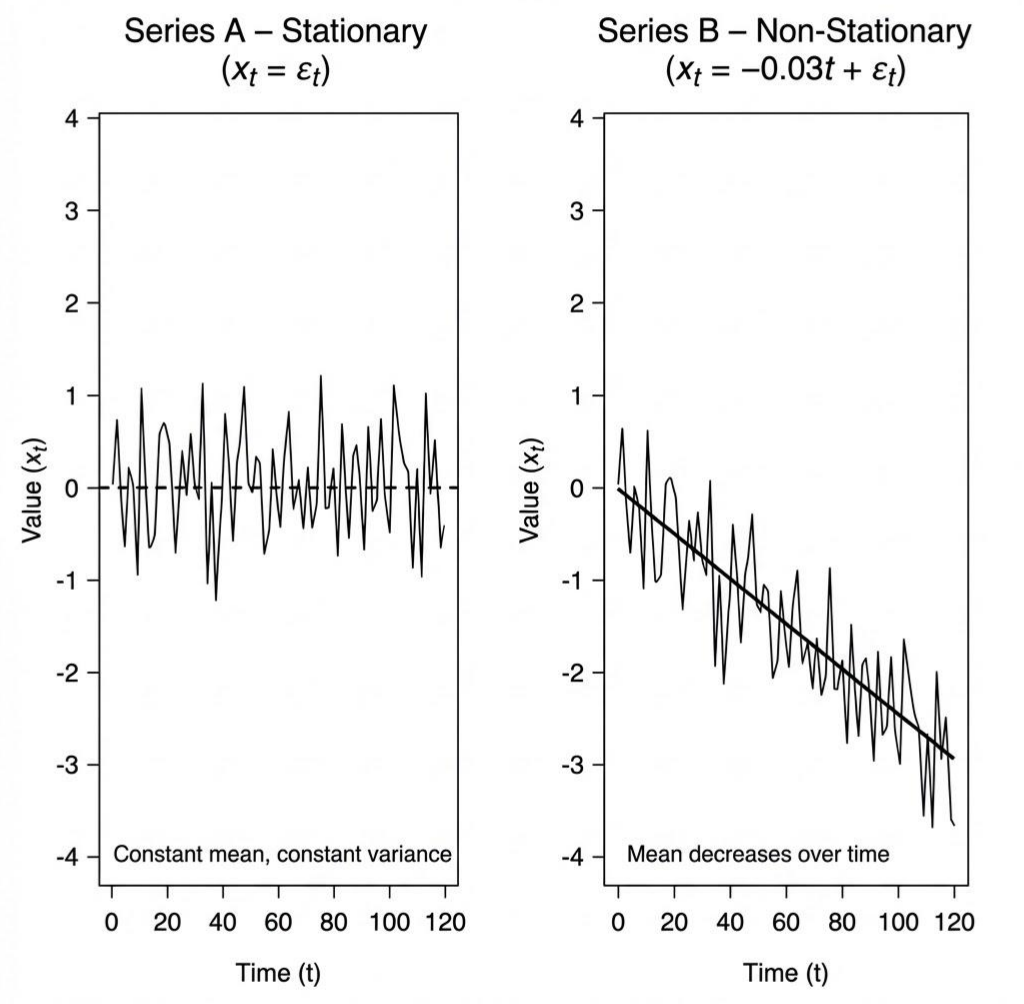

Content

- Stationarity
- Unit Root Test
- Differencing
- AR / MA / ARMA Models

Unit Root Tests

Why Do We Need Unit Root Tests?

- **Visual inspection** is not sufficient
- Some **non-stationary** series look stationary
- Unit root causes **persistent shocks**
- Formal statistical testing is required

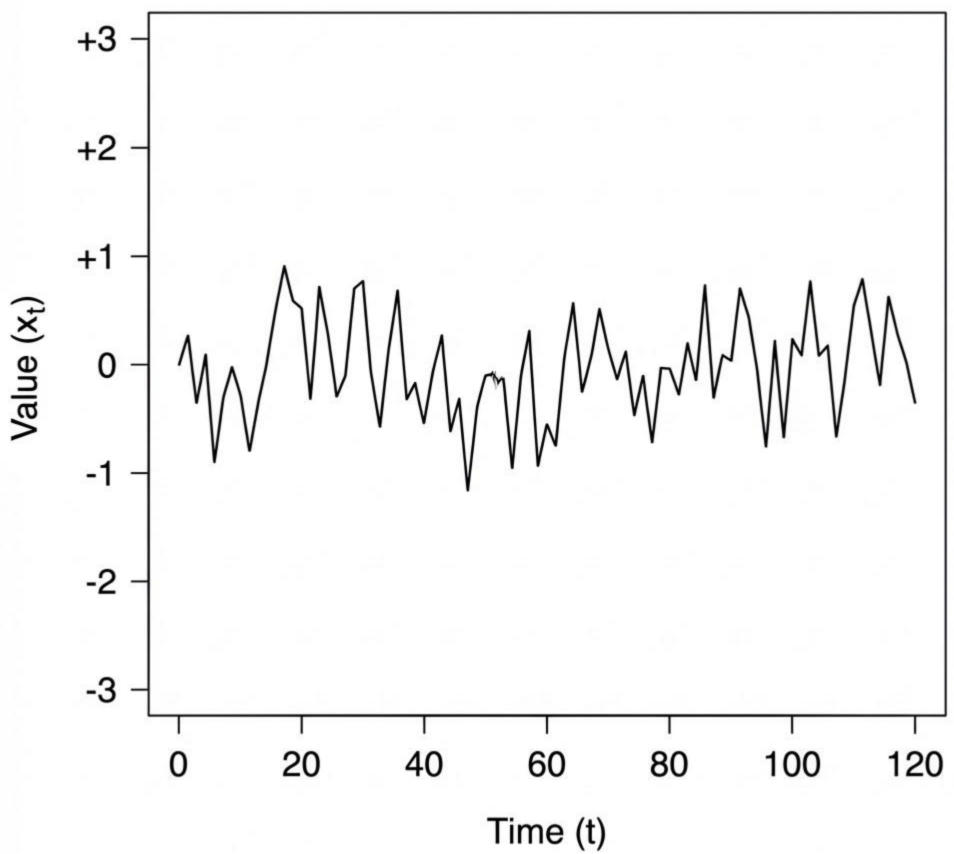

Unit Root Tests

Why Do We Need Unit Root Tests?

- Review this figure:
    - Small amplitude of oscillation (low noise variance)
    - There is no clear trend in the short term.
    - Fluctuating around a "seemingly stable" level.
- However: This is a non-stationary series
    - Mathematical model

$$x_t = x_{t-1} + \varepsilon_t$$

- This is a **random walk**, with:
    - Unit root
    - Without string pull, the mean returns to zero.

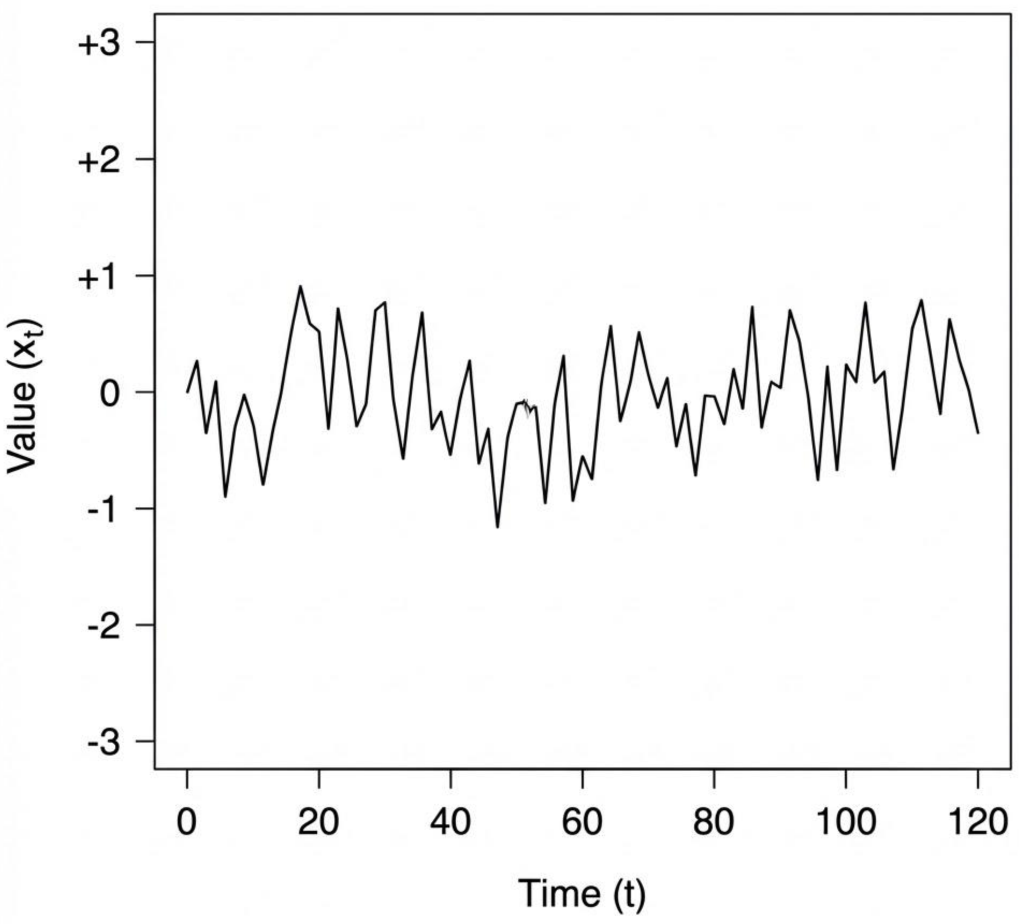

Unit Root Tests

What Is a Unit Root?

- Consider an AR(1) process: $x_t = \phi x_{t-1} + \varepsilon_t$
- The process has a **unit root** if: $\phi = 1$
- In this case:

$$x_t = x_{t-1} + \varepsilon_t$$

  - The series becomes a **random walk**
  - Shocks have **permanent effects**
  - Mean and variance are **not time-invariant**

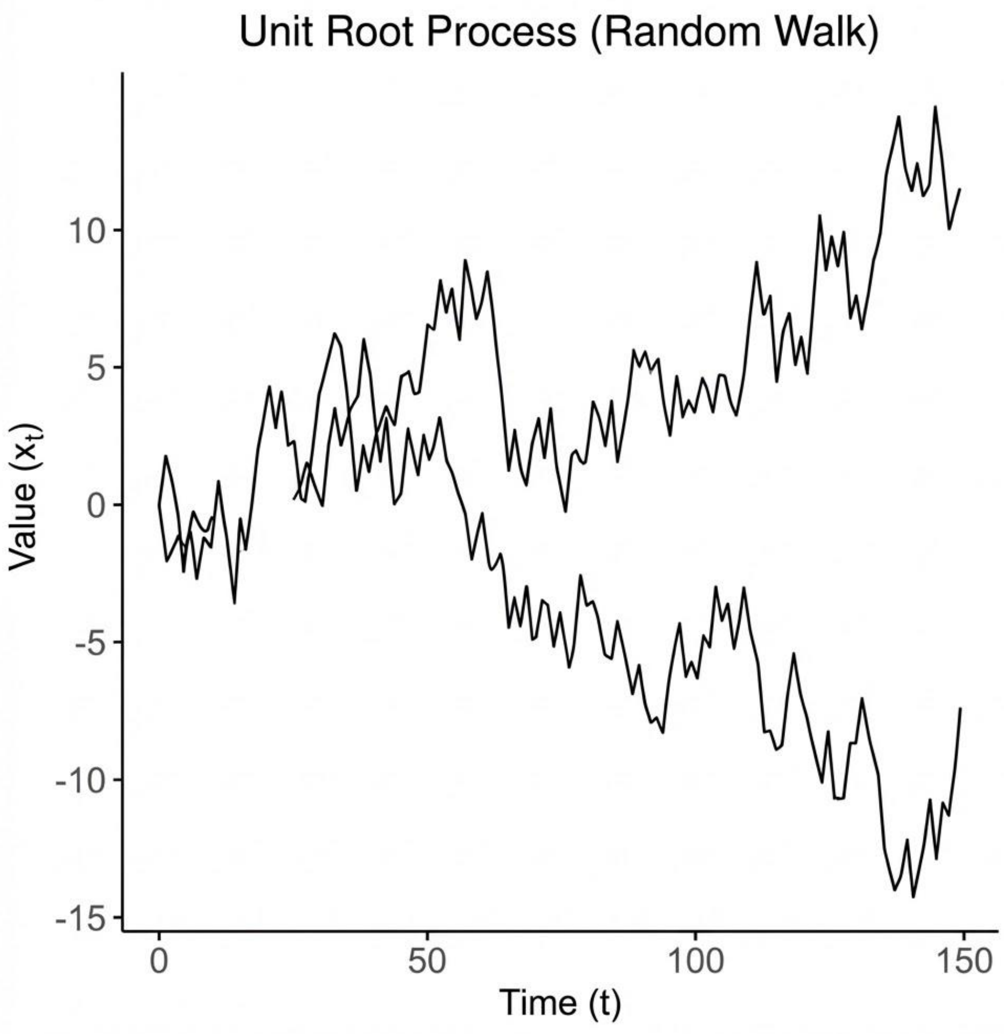

Unit Root Tests

What is being tested?

- Whether a time series is stationary or contains a unit root
- Focus on the stochastic trend in the data
- Hypotheses
  - Null hypothesis (H_0): The series has a unit root (non-stationary)
  - Alternative hypothesis (H_1): The series is stationary
- Common unit root tests
  - ADF (Augmented Dickey–Fuller)
  - KPSS (reverse hypothesis)

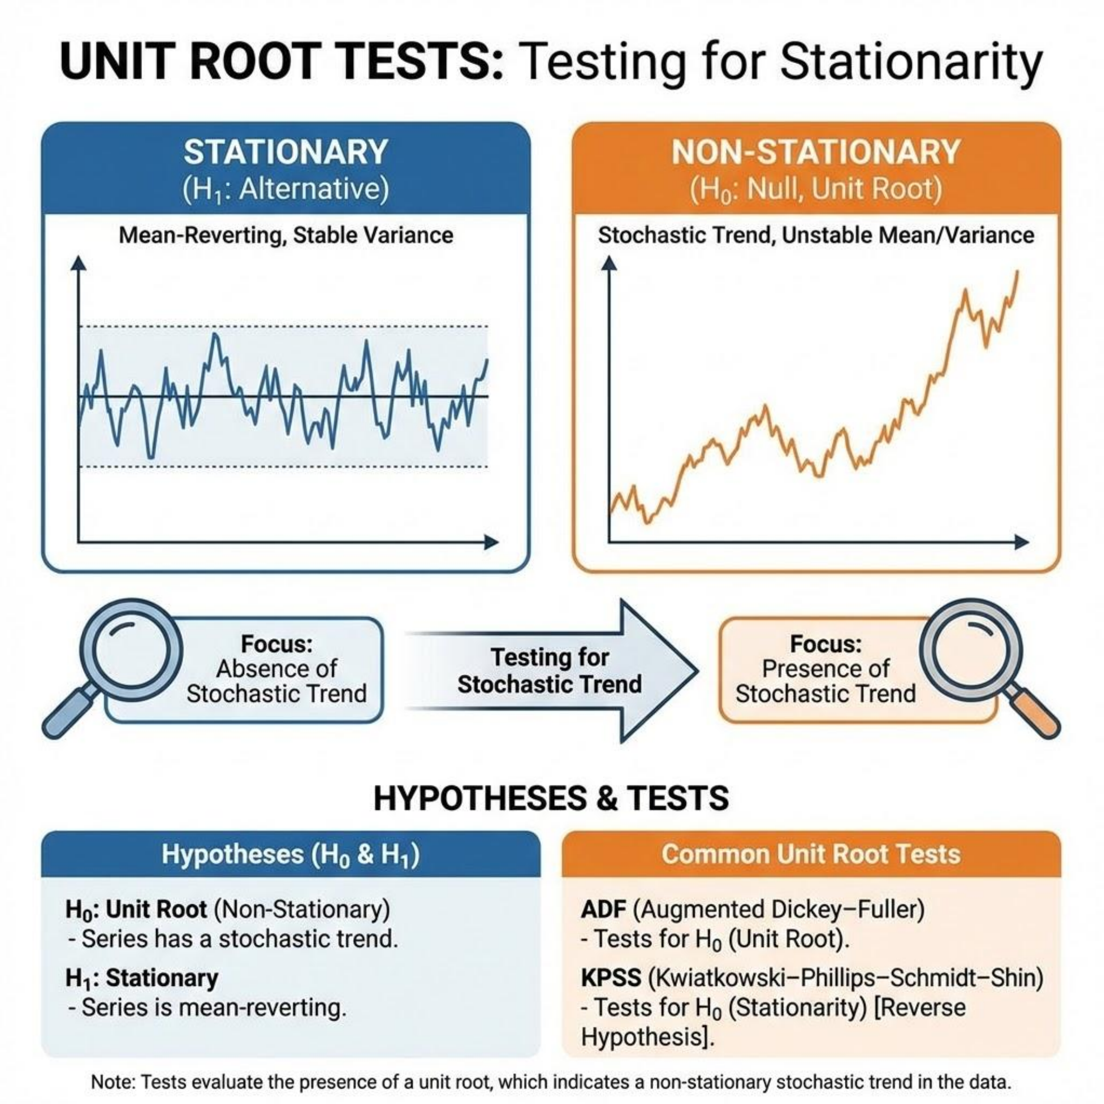

Unit Root Tests

Augmented Dickey–Fuller (ADF) Test

Purpose

- Test whether a time series has a **unit root**
- Extension of the Dickey–Fuller test to handle **autocorrelation**

Test equation

$$\Delta y_t = \gamma y_{t-1} + \sum_{i=1}^{p} \alpha_i \Delta y_{t-i} + \varepsilon_t$$

Hypotheses

- **H₀**: $\gamma = 0$ → unit root (non-stationary)
- **H₁**: $\gamma < 0$ → stationary

Key idea

- Test the **significance of** γ
- Uses **non-standard critical values**

Where:

- Δy_t: first difference of the series
- y_{t-1}: lagged level (unit root term)
- γ: unit root coefficient (key parameter)
- Δy_{t-i}: lagged differences (augmentation terms)
- p: number of lags
- ε_t: white noise error

ADF intuition

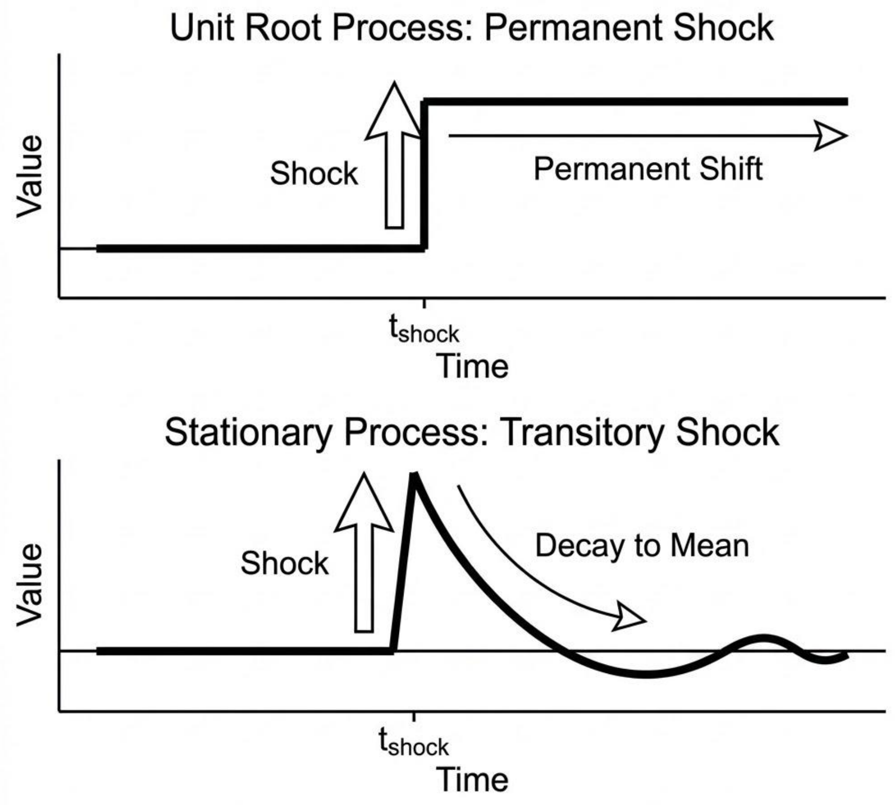

Unit Root Tests

Augmented Dickey–Fuller (ADF) Test

Current state of the system
The level the series was at previously

$$\Delta y_t = \gamma y_{t-1} + \sum_{i=1}^{p} \alpha_i \Delta y_{t-i} + \varepsilon_t$$

"Augmented" part

$\Delta y_t = y_t - y_{t-1}$ *: First-order difference of a series*

Unit Root Tests

Augmented Dickey–Fuller (ADF) Test

Parameters to be tested
Answer the question: "When the trend is at a high, does it tend to reverse or continue further?"

$$\Delta y_t = \gamma y_{t-1} + \sum_{i=1}^{p} \alpha_i \Delta y_{t-i} + \varepsilon_t$$

- $\gamma = 0$: No pull-back → Random walk
- $\gamma < 0$: With pull-back → Mean-reverting
- $\gamma > 0$: Explosive sequence

- The entire ADF test only tests a single hypothesis: $H_0 : \gamma = 0$

Unit Root Tests

Augmented Dickey–Fuller (ADF) Test

```
from statsmodels.tsa.stattools import adfuller


def adf_test(timeseries):
    print("Results of Dickey-Fuller Test:")
    dftest = adfuller(timeseries, autolag="AIC")
    dfoutput = pd.Series(
        dftest[0:4],
        index=[
            "Test Statistic",
            "p-value",
            "#Lags Used",
            "Number of Observations Used",
        ],
    )
    for key, value in dftest[4].items():
        dfoutput["Critical Value (%s)" % key] = value
    print(dfoutput)
```

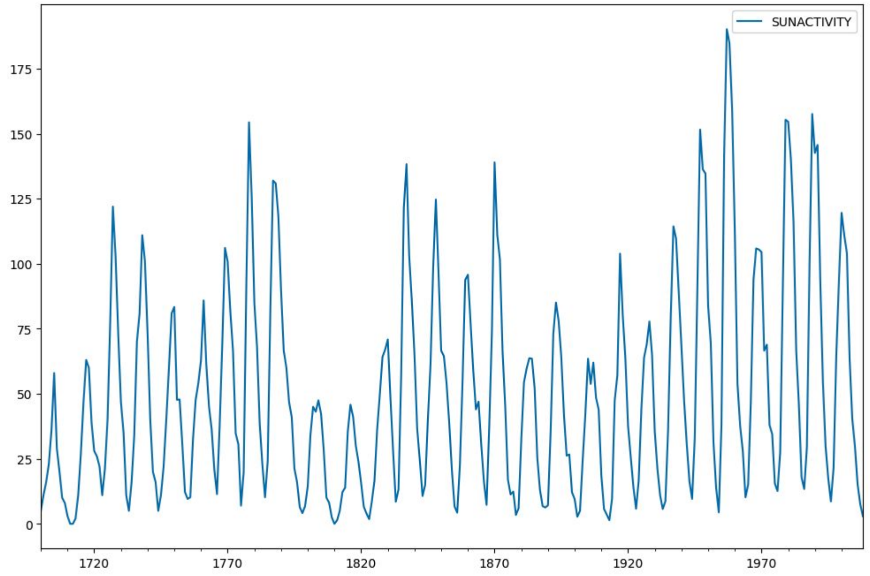


```
Results of Dickey-Fuller Test:
Test Statistic                  -2.837781
p-value                          0.053076
#Lags Used                       8.000000
Number of Observations Used    300.000000
Critical Value (1%)             -3.452337
Critical Value (5%)             -2.871223
Critical Value (10%)            -2.571929
dtype: float64
```

Unit Root Tests

Augmented Dickey–Fuller (ADF) Test

If ADF statistic < Critical Value → Reject H_0 (unit root)

γ

```
Results of Dickey-Fuller Test:
Test Statistic                  -2.837781
p-value                          0.053076
#Lags Used                       8.000000
Number of Observations Used    300.000000
Critical Value (1%)             -3.452337
Critical Value (5%)             -2.871223
Critical Value (10%)            -2.571929
dtype: float64
```

Threshold provided by Dickey-Fuller

Unit Root Tests

KPSS Test (Kwiatkowski–Phillips–Schmidt–Shin)

- KPSS tests stationarity directly, not unit root.
- Whether a time series is stationary around a level or a trend.
- Hypotheses:
  - H_0: The series is **stationary**
  - H_1: The series is **non-stationary** (has unit root)

- Refer code example

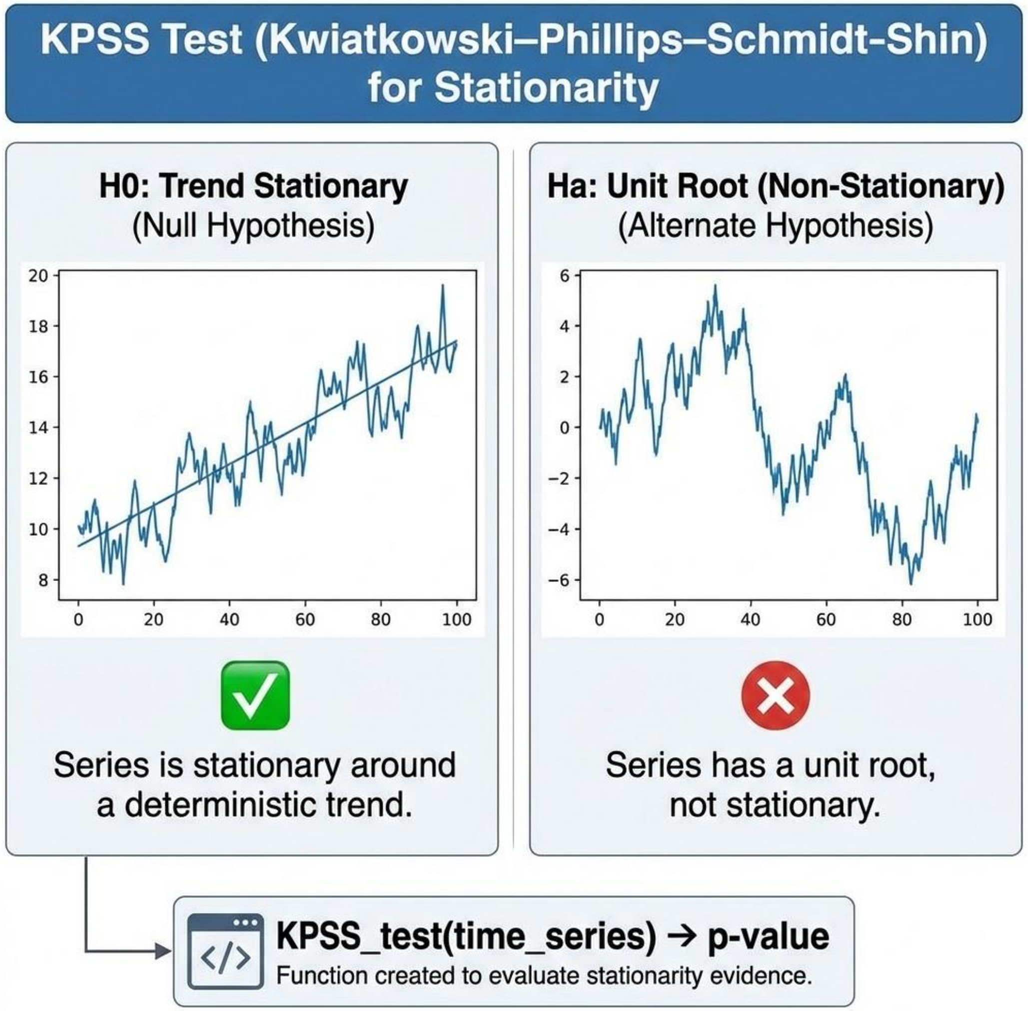

Content

- Stationarity
- Unit Root Test
- Differencing
- AR / MA / ARMA Models

Differencing

Differencing

Goal:

> Remove unit root and make the series stationary.

First-order differencing:

$$\Delta y_t = y_t - y_{t-1}$$

Higher-order differencing:

$$\Delta^d y_t = (1 - L)^d y_t$$

Key idea:

- Differencing removes **stochastic trend**
- Transforms **non-stationary → stationary**

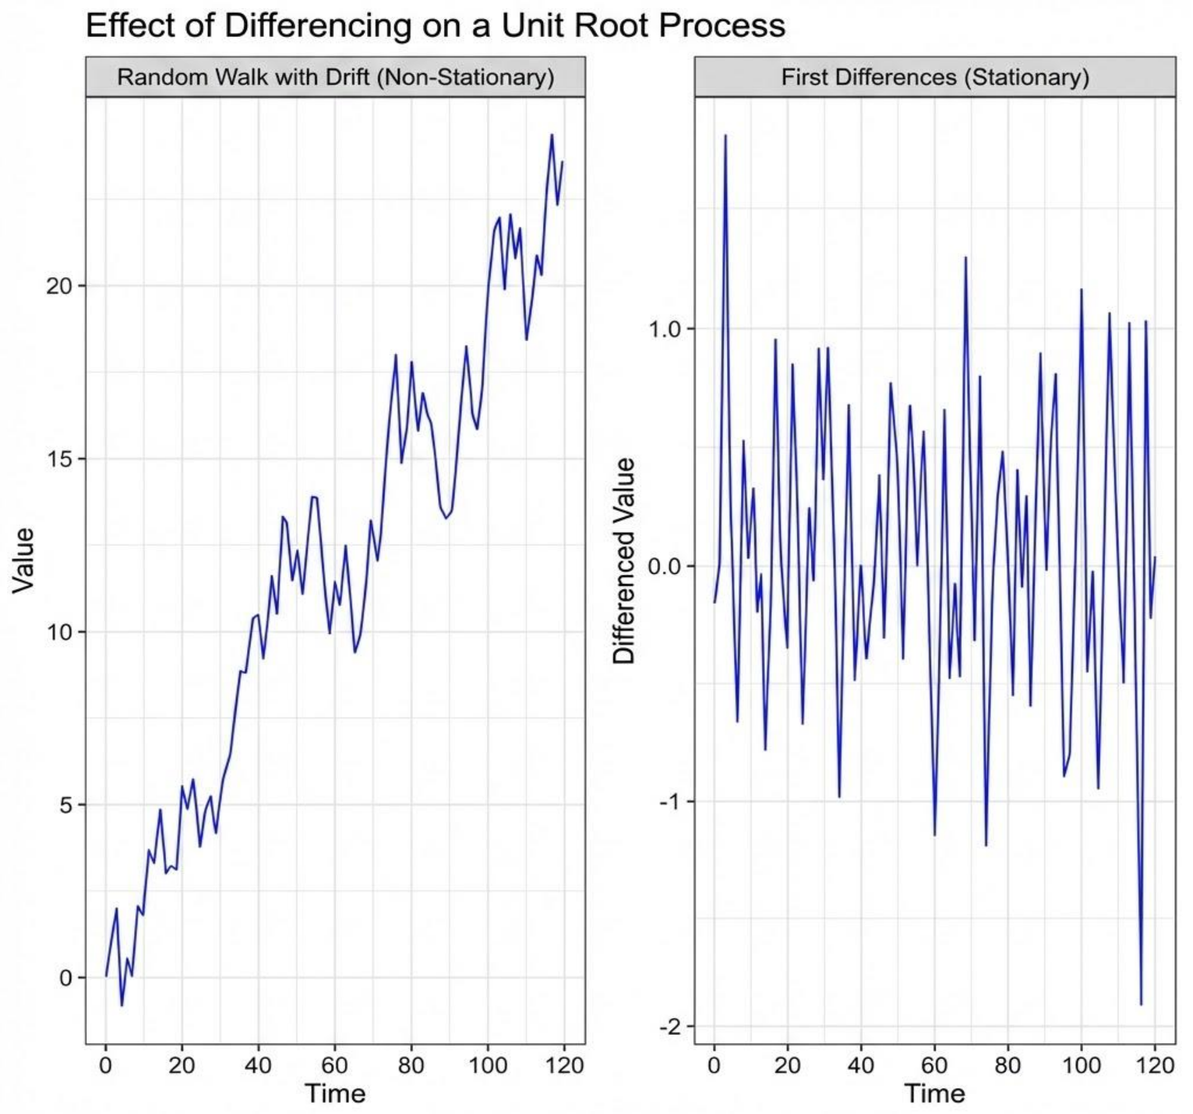

- Refer code example

Differencing

Differencing

```
Results of Dickey-Fuller Test:
Test Statistic                    -1.486166e+01
p-value                            1.715552e-27
#Lags Used                         7.000000e+00
Number of Observations Used        3.000000e+02
Critical Value (1%)               -3.452337e+00
Critical Value (5%)               -2.871223e+00
Critical Value (10%)              -2.571929e+00
dtype: float64
```

- Refer code example

Before:

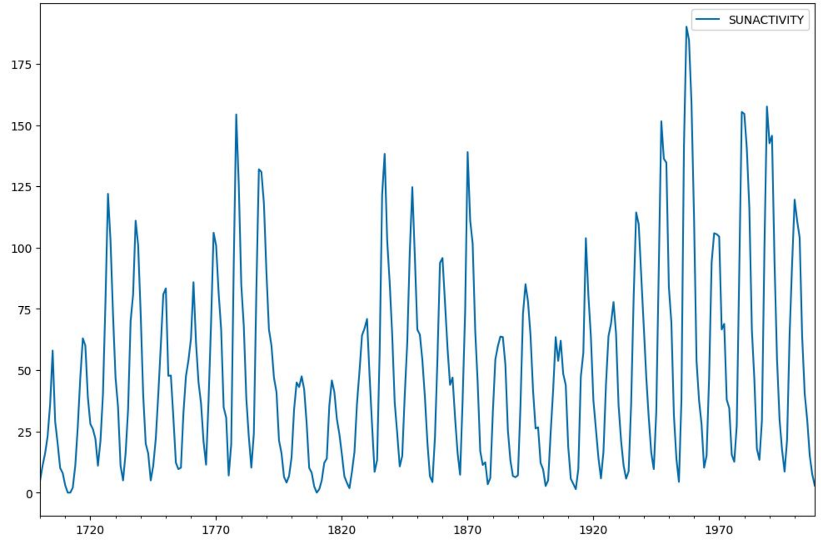


```python
sunspots["SUNACTIVITY_diff"] = sunspots["SUNACTIVITY"] - sunspots["SUNACTIVITY"].shift(
    1
)
sunspots["SUNACTIVITY_diff"].dropna().plot(figsize=(12, 8))
```

After:

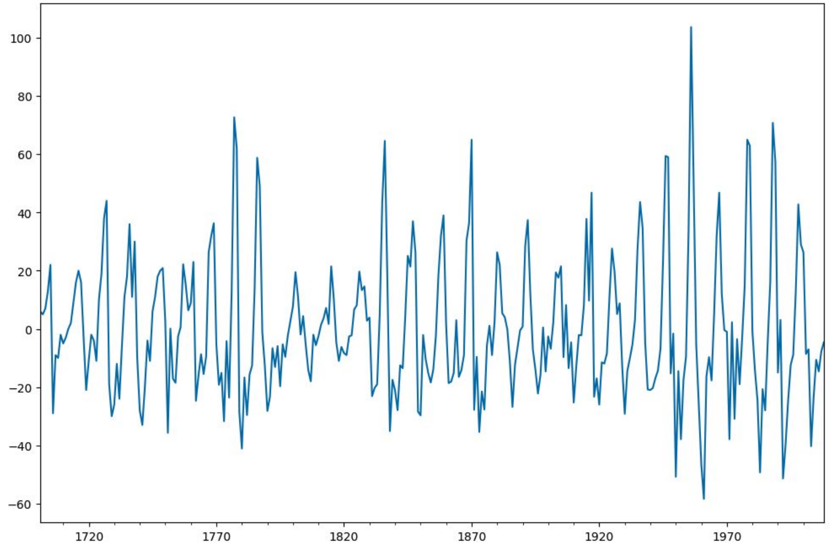

Content

- Stationarity
- Unit Root Test
- Differencing
- AR / MA / ARMA Models

AR / MA / ARMA Models

- After differencing, many time series become **stationary**
- Stationary series still exhibit **temporal dependence**
- AR / MA / ARMA models:
    - Capture **short-term memory**
    - Describe how the present depends on:
        - past values
        - past shocks
- They form the foundation of **ARIMA models**

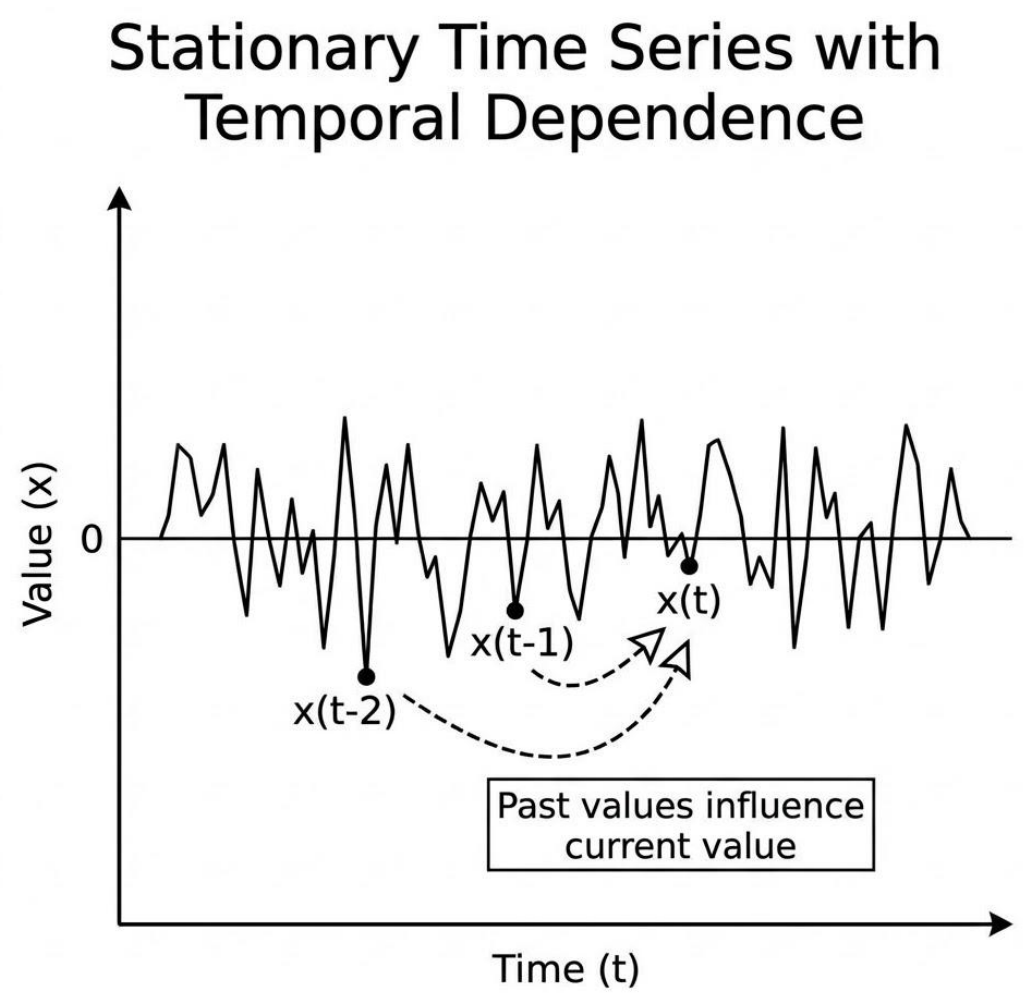

AR / MA / ARMA Models

Autoregressive (AR) Model

- An AR model explains the current value of a time series using its **own past values.**
- The idea is that the series has **memory:** past observations influence the present.
- Suitable for **stationary** time series.

AR(p) model:

$$x_t = c + \sum_{i=1}^{p} \phi_i x_{t-i} + \varepsilon_t$$

Where:

- x_t: value at time t
- p: order of the AR model
- ϕ_i: autoregressive coefficients
- ε_t: white noise

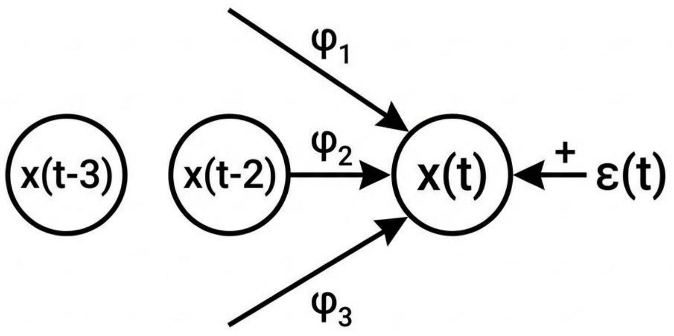

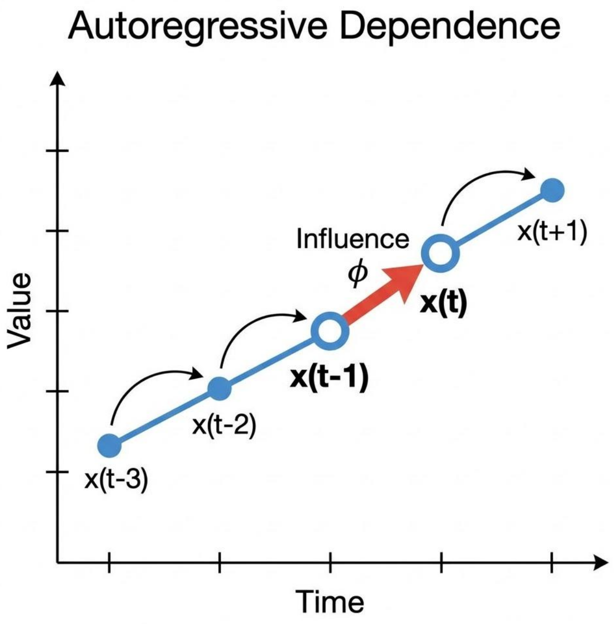

AR / MA / ARMA Models

Autoregressive (AR) Model

```python
from statsmodels.tsa.ar_model import AutoReg

# train
model = AutoReg(ts_diff, lags=1)
result = model.fit()

print(result.summary())
```

```
                            AutoReg Model Results
==============================================================================
Dep. Variable:            SUNACTIVITY   No. Observations:                  308
Model:                     AutoReg(1)   Log Likelihood               -1358.217
Method:               Conditional MLE   S.D. of innovations             20.191
Date:                Mon, 26 Jan 2026   AIC                           2722.434
Time:                        15:15:14   BIC                           2733.615
Sample:                             1   HQIC                          2726.905
                                  308
==================================================================================
                     coef    std err          z      P>|z|      [0.025      0.975]
----------------------------------------------------------------------------------
const             -0.0308      1.152     -0.027      0.979      -2.289       2.228
SUNACTIVITY.L1     0.5412      0.048     11.278      0.000       0.447       0.635
                                    Roots
=============================================================================
                  Real          Imaginary           Modulus         Frequency
-----------------------------------------------------------------------------
AR.1            1.8476           +0.0000j            1.8476            0.0000
-----------------------------------------------------------------------------
```

- Refer code example 2

Original:

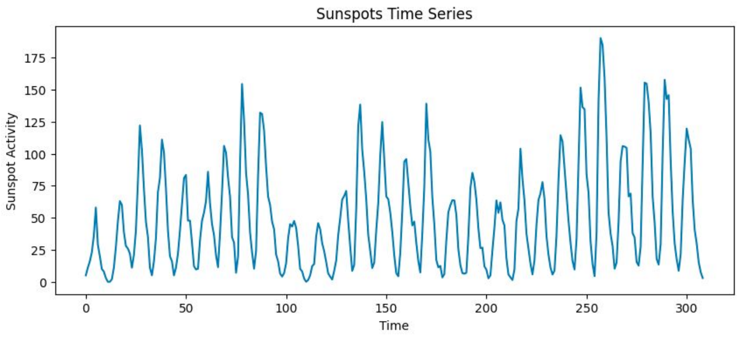

Differencing:

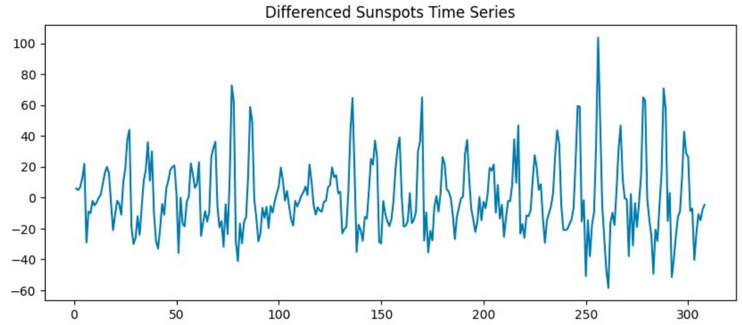

Forcasting:

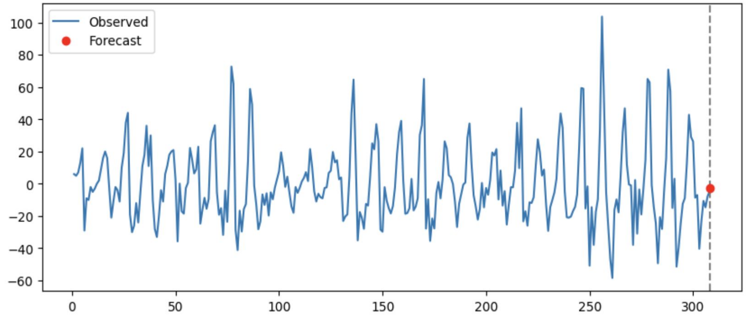

AR / MA / ARMA Models

Moving Average (MA) Model

- An MA model explains the current value of a time series using **past error terms (shocks)**.
- The idea is that the series reacts to **unexpected disturbances**, not past values.
- Suitable for **stationary time series**.

MA(q) model:

$$x_t = \mu + \varepsilon_t + \sum_{i=1}^{q} \theta_i \varepsilon_{t-i}$$

Where:

- x_t: value at time t
- μ: mean of the series
- ε_t: white noise (random shock)
- q: order of the MA model
- θ_i: moving average coefficients

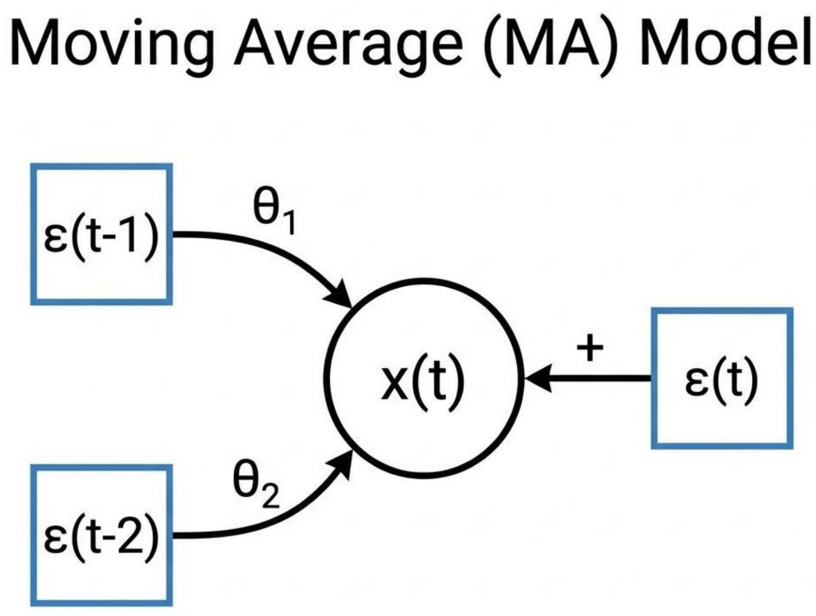

AR / MA / ARMA Models

Moving Average (MA) Model

```python
from statsmodels.tsa.arima.model import ARIMA

# MA(q) is ARIMA(0, 0, q)
model = ARIMA(ts_diff, order=(0, 0, 1))
result = model.fit()

print(result.summary())
```

MA trong statsmodels = ARIMA(0,0,q)

Trong ARIMA:

- p → AR part
- d → differencing
- q → MA part

Do đó:

MA(q)≡ ARIMA(0,0,q)

Nghĩa là:

- Không dùng giá trị quá khứ
- Không differencing
- Chỉ dùng q sai số quá khứ

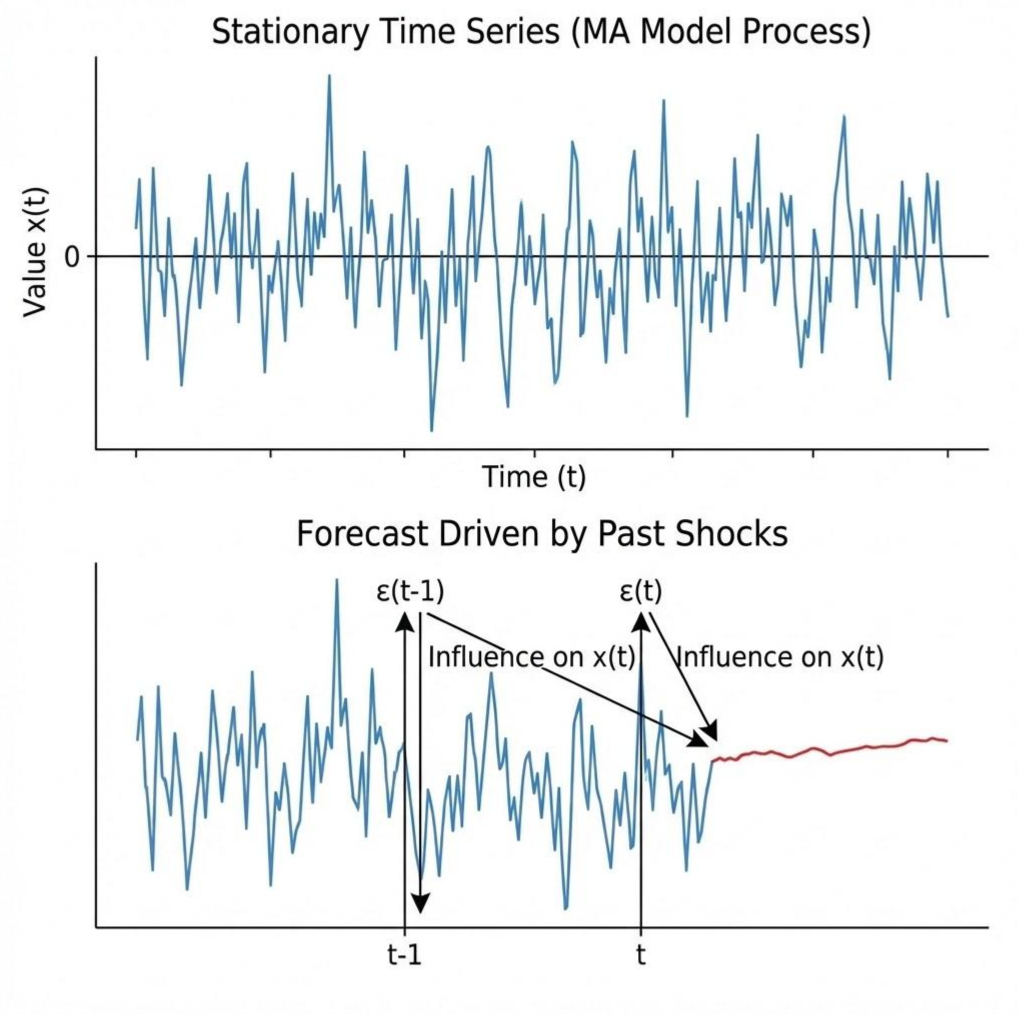

AR / MA / ARMA Models

Autoregressive (AR) vs Moving Average (MA)

Aspect	AR Model	MA Model
Core idea	Uses past values	Uses past shocks
Memory	Long-term memory	Short-term memory
Mathematical form	$x_t = \sum \phi_i x_{t-i} + \varepsilon_t$	$x_t = \varepsilon_t + \sum \theta_i \varepsilon_{t-i}$
Dependency	Past observations	Past errors
Interpretation	Persistence	Shock absorption
Typical use	Trend-like persistence	Noise-like series
ACF pattern	Tails off	Cuts off
PACF pattern	Cuts off	Tails off
Stationarity required	Yes	Yes

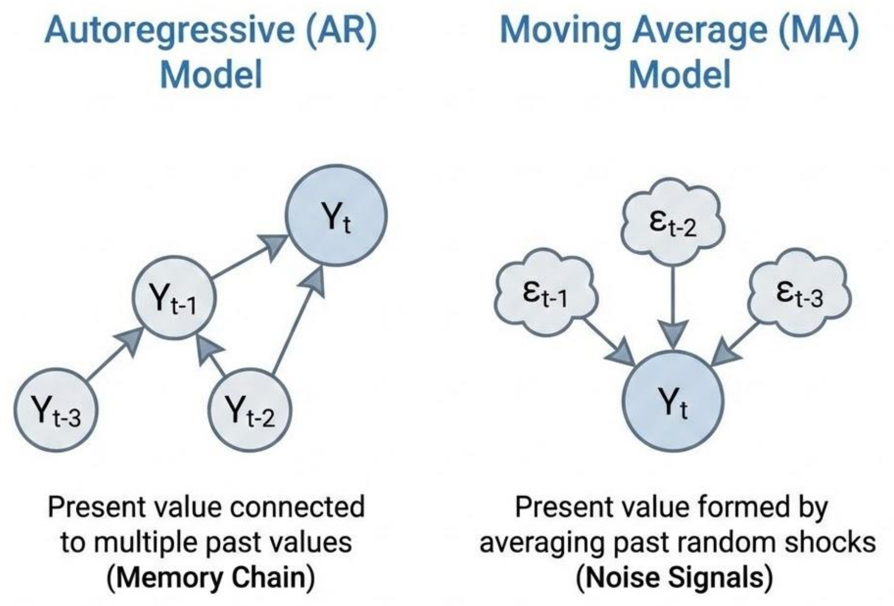

Thank you