Time Series

Dr. Meriem YOUSSEF

Pristini School of AI

24 février 2025

What is a time series ?

- A time series is a **chronological sequence of quantitative observations** x_t, observed over a period of time (weekly, monthly, quarterly, or yearly).
- Time series is a sequence

$$\{x_1, x_2, \ldots, x_T\} \quad \text{or} \quad x_t, t = 1, 2, \ldots, T$$

where t is an index denoting the period in time in which x occurs.

Main Objective of Time Series

- Predict/forecast the future given current and past observations.
- The properties of observed data are used to predict future observations of relevant variables.
- Aim to predict the response given the observed variables.

Types of Time Series

There are two types of time series data :

1. **Univariate** time series, are those where only one variable is measured over time.
2. **Multivariate** time series, are those where more than one variable are measured simultaneously.

Areas of Application of Time Series

Time series data provide useful information about the physical, biological, social or economic systems generating the time series, such as :

- ***Economics :*** profits, imports, exports, stock exchange indices,
- ***Sociology :*** school enrollments, unemployment, crime rate,
- ***Environment :*** amount of pollutants, such as CO_2 emissions,
- ***Medicine :*** blood pressure measurements over time for evaluating drugs to control hypertension.

Examples of Time series

Here are a few examples of plots of time series data :

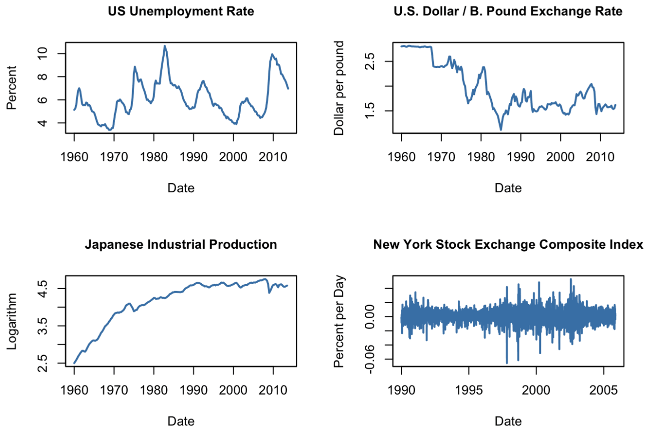

Examples of Time series

Here are a few examples of plots of time series data :

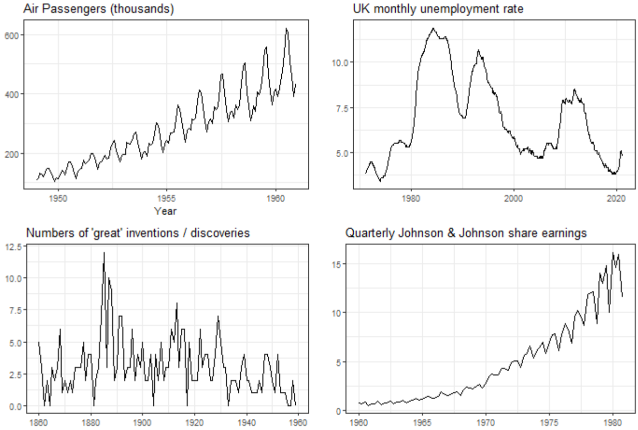

Components of a Time series

There are four components to a time series :

- Trend.
- Cycle.
- Seasonal Variations.
- Irregular Fluctuations.

Components of a Time series

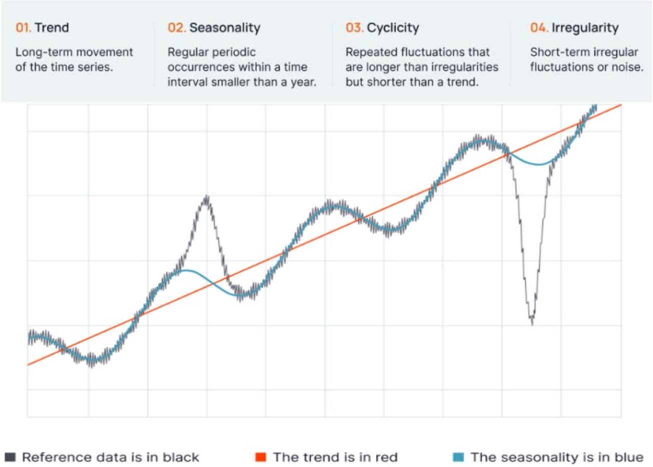

Components of a Time series

Basic framework for time series analysis

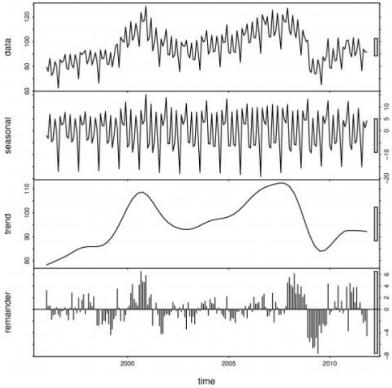

y_t (future forecast) data

=

S_t seasonal (fixed-period) fluctuations

+ or X

T_t underlying trend

+ or X

ε_t random error or varying-period cycle

Stationarity

Let $\{X_t\} = \{\dots, X_{t-1}, X_t, X_{t+1}, \dots\}$ denote a sequence of random variables (a time series) indexed by some time subscript t.
A Time Series is **stationary** if has the following conditions :

1. Constant μ (mean) for all t.
2. Constant σ (variance) for all t.
3. The autocovariance function between X_{t1} and X_{t2} only depends on the interval t_1 and t_2.

Stationarity

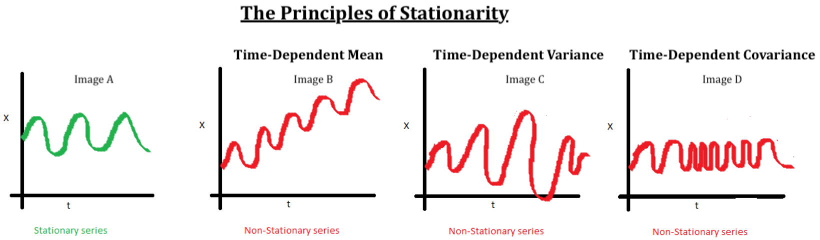

Stationarity Examples

The typical form of a stationary time series, commonly known as **white noise**.

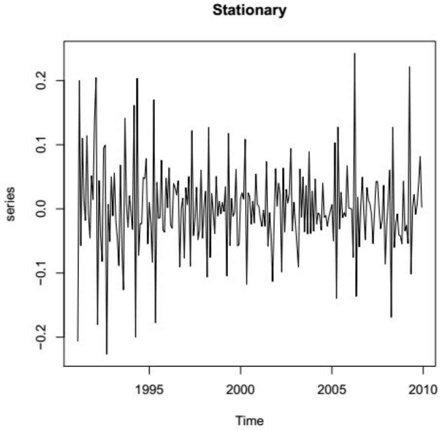

Stationarity Examples

Nonconstant variance series (**Heterocedasticity**).

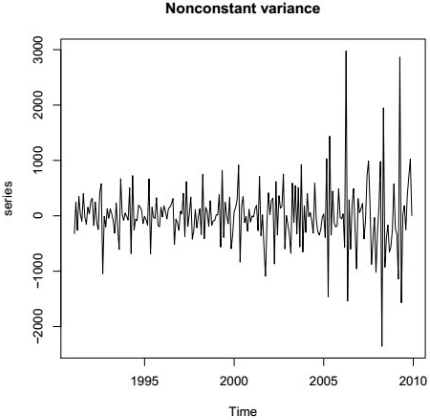

Stationarity Examples

Nonconstant mean series (**Trend**).

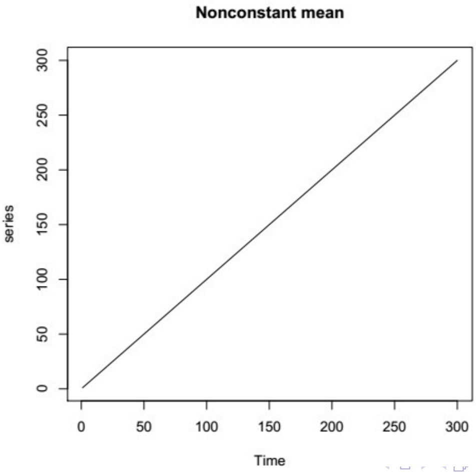

Stationarity Examples

Nonconstant mean series (**Trend**).

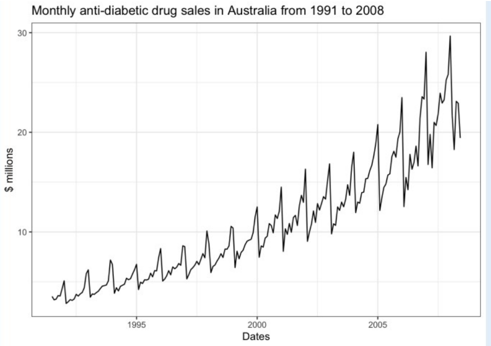

Stationarity Examples

Nonconstant mean and variance series.

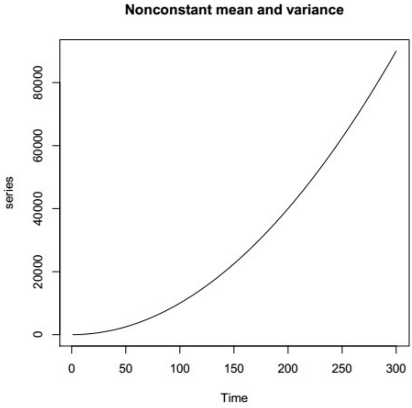

Why Stationarity is Important in Time Series

Key Reasons Why Stationarity Matters

- **Predictability and Consistency :** Ensures patterns are stable over time, aiding in reliable forecasting.
- **Model Validity :** Many time series models (e.g., ARMA, ARIMA, GARCH) assume stationarity for accurate estimation.
- **Statistical Inference :** Hypothesis testing and confidence intervals rely on stationarity assumptions.
- **Simplified Analysis :** Stationary series have simpler patterns that are easier to model.
- **Identifying True Relationships :** Reduces the risk of spurious correlations in multivariate analysis.

Testing for Stationarity

Common Methods to Test for Stationarity

- **Visual Inspection :** Plotting the data to check for trends and seasonality.
- **Augmented Dickey-Fuller (ADF) Test :** Formal test for the presence of unit roots.
- **Kwiatkowski-Phillips-Schmidt-Shin (KPSS) Test :** Checks for trend and level stationarity.
- **Correlogram Analysis :** Examining autocorrelation and partial autocorrelation plots.

Ensuring stationarity is often the first step in time series modeling for meaningful forecasting and analysis.

Autocorrelation ACF

- The coefficient of correlation between two values in a time series is called **the autocorrelation function (ACF)**.
- For example the ACF for a time series X_t is given by :

$$Corr(X_t, X_{t-k}).$$

- This value of k is the **time gap** being considered and is called the **lag**.
- A lag 1 autocorrelation (i.e., $k = 1$ in the above) is the correlation between values that are one time period apart.
- A lag k autocorrelation is the correlation between values that are k time periods apart.

Autocorrelation ACF

- ACF values fall between -1 and +1 calculated from the time series at different lags to measure the significance of correlations between the present observation and the past observations, and to determine how far back in time (i.e., of how many time-lags) are they correlated.
- The formula for Autocorrelation (ACF) :

$$\begin{aligned} \hat{\rho}_k &= \frac{\hat{cov}(X_t, X_{t-k})}{\hat{var}(X_t)} \\ &= \frac{\sum_{t=k+1}(X_t - \overline{X})(X_{t-k} - \overline{X})}{\sum_{t=1}^{n}(X_t - \overline{X})^2} \end{aligned}$$

Autocorrelation ACF

- Formula above essentially tells us that the autocorrelation coefficient for some lag k is calculated as the covariance between the original series and the series removed k lags, divided by the variance of the original series.

Autoregressive Models

An **autoregressive model of order p** (AR(p), $p = 1, 2, \ldots$) is given by :

$$\begin{aligned} X_t &= \sum_{j=1}^{p} \phi_j X_{t-j} + \varepsilon_t, \quad t \in \mathbb{N}^* \qquad (1) \\ &= \phi_0 + \phi_1 X_{t-1} + \phi_2 X_{t-2} + \ldots + \phi_p X_{t-p} + \varepsilon_t. \end{aligned}$$

Where $\{\varepsilon_t\}$ are independent, identically distributed, zero mean random variables.

Autoregressive Models

- The order of an autoregression is the number of immediately preceding values in the series that are used to predict the value at the present time.
- A p^{th}-order autoregression, written as $AR(p)$, is a multiple linear regression in which the value of the series at any time t is a (linear) function of the values at times $t-1, t-2, \ldots, t-p$.

Autoregressive Models Examples

$AR(1)$ MODELS :

- An autoregressive model is when a value from a time series is regressed on previous values from that same time series.
- for example, X_t on X_{t-1} :

$$X_t = \phi_0 + \phi_1 X_{t-1} + \varepsilon_t$$

- Where : $\varepsilon_i \sim N(0, \sigma^2)$,
  $cov(\varepsilon_i, \varepsilon_j) = 0 \quad \text{for} \quad i \neq j \quad \text{and}$
  $cov(\varepsilon_i, X_j) = 0 \quad \forall i, j$

Autoregressive Models Examples

$AR(1)$ MODELS :

- When $|\phi_1| < 1$, such a process, $AR(1)$ is **stationary**.
- When $|\phi_1| = 1$, we have a **random walk**.

Autoregressive Models Properties

- Property 1 : The mean of the X_i in a stationary $AR(p)$ process is :

$$\mu = \frac{\phi_0}{1 - \sum_{j=1}^{p} \phi_j}$$

- Property 2 : The variance of the X_i in a stationary $AR(1)$ process is :

$$var(X_i) = \frac{\sigma^2}{1 - \phi_1^2}$$

Autoregressive Models Properties

- Property 3 : The lag k autocorrelation in a stationary AR(1) process is :

$$\rho_k = \phi_1^k$$

- Property 4 : For any stationary $AR(p)$ process. The **autocovariance** at lag $k > 0$ can be calculated as :

$$\gamma_k = \phi_1\gamma_{k-1} + \phi_2\gamma_{k-2} + \ldots + \phi_p\gamma_{k-p}$$

- Similarly the **autocorrelation** at lag $k > 0$ can be calculated as :

$$\rho_k = \phi_1\rho_{k-1} + \phi_2\rho_{k-2} + \ldots + \phi_p\rho_{k-p}$$

Autoregressive Models Application

Example 1 : Simulate a sample of 100 elements from the $AR(1)$ process

$$X_t = 5 + 0.4X_{t-1} + \varepsilon_t$$

where $\varepsilon_i \sim N(0, 1)$ and calculate ACF.

Autoregressive Models Examples

$AR(2)$ MODELS :

- If we want to predict this year X_t using measurements of global temperature in the previous two years (X_{t-1}, X_{t-2}), then the autoregressive model for doing so would be :

$$X_t = \phi_0 + \phi_1 X_{t-1} + \phi_2 X_{t-2} + \varepsilon_t$$

- This model is a second-order autoregression, written as $AR(2)$, since the value at time t is predicted from the values at times $t-1$ and $t-2$.

Autoregressive Models Examples

Example 2 : Repeat Example 1 for the AR(2) process.

$$X_t = 5 + 0.4X_{t-1} + 0.1X_{t-2} + \varepsilon_t$$

where $\varepsilon_i \sim N(0, 1)$, and calculate ACF.

Moving Average Models

- The $MA(q)$ model say :

$$\begin{aligned} X_t &= \mu + \sum_{k=1}^{q} \theta_k \varepsilon_{t-k} + \varepsilon_t \\ &= \mu + \theta_1 \varepsilon_{t-1} + \ldots + \theta_q \varepsilon_{t-q} + \varepsilon_t \end{aligned}$$

- Where : $\{\varepsilon_t\}$ are independent, identically distributed, zero mean random variables.
  $\varepsilon_i \sim N(0, \sigma^2)$,
  $cov(\varepsilon_i, \varepsilon_j) = 0 \quad$ for $\quad i \neq j \quad$ and
  $cov(\varepsilon_i, X_j) = 0 \quad \forall i, j$
- The value of X depends on ε_t and q past values of ε.
- The value of X at time t is a linear function of past errors.

Autoregressive Models Properties

- Property 1 : The mean of an $MA(q)$ process is : μ
- Property 2 : The variance of an $MA(q)$ process is :

$$var(y_t) = \sigma^2(1 + \theta_1^2 + \ldots + \theta_q^2)$$

Autoregressive Models Properties

- Property 3 : The autocorrelation function of an $MA(1)$ process is :

$$\rho_1 = \frac{\theta_1}{1+\theta_1^2}, \quad \rho_k = 0 \quad \text{for} \quad k > 1$$

- Property 4 : The autocorrelation function of an $MA(q)$ process is :

$$\rho_k \quad = \quad \frac{\theta_k + \sum_{j=1}^{q-k} \theta_j \theta_{j+k}}{1+\sum_{j=1}^{q} \theta_j^2}$$

Moving Average Models Examples

$MA(1)$ models :

$$y_t = \mu + \theta_1 \varepsilon_{t-1} + \varepsilon_t$$

- Where ε_t and ε_{t-1} are independent of each other and $E(\varepsilon_t) = 0$.
- $E(y_t) = \mu$ and $var(y_t) = (1 + \theta_1^2)\sigma_\varepsilon^2$

The ARMA Models

- We can combine AR and MA models.
- ARMA(1,1) is defined as :

$$X_t = \phi_0 + \phi_1 X_{t-1} + \theta_1 \varepsilon_{t-1} + \varepsilon_t.$$

The ARMA Models

- The mean of the ARMA(1,1) times series is :

$$E(X_t) = \phi_0 + \phi_1 E(X_{t-1}) = \phi_0 + \phi_1 E(X_t)$$

- When $|\phi_1| < 1$, $E(X_t) = \dfrac{\phi_0}{1 - \phi_1}$
- X_t will tend to fluctuate around the mean.
- X_t is mean-reverting in this case.

The ARMA Models

- The general ARMA(p,q) model is :

$$X_t = \phi_0 + \Sigma_{i=1}^{p} \phi_i X_{t-i} + \Sigma_{i=1}^{q} \theta_i \varepsilon_{t-i} + \varepsilon_t$$

- The $MA(q)$ average has the feature that after q lags there isn't any correlation between two random variables.
- There are correlations at all lags for an $AR(p)$ model.

The ARIMA Family Models

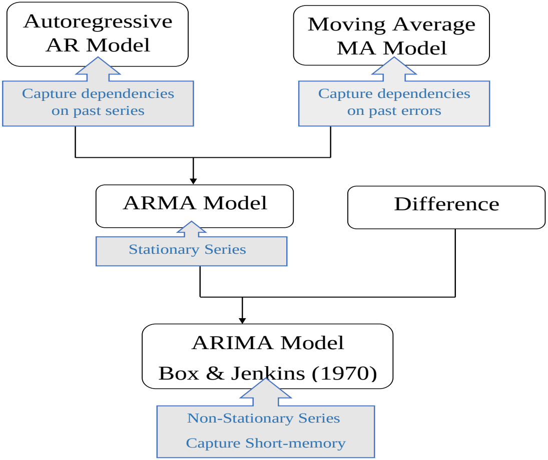

General ARIMA(p,d,q) Model

- The general ARIMA(p,d,q) model is given by :

$$\Delta^d X_t = \phi_0 + \sum_{i=1}^{p} \phi_i X_{t-i} + \sum_{j=1}^{q} \theta_j \varepsilon_{t-j} + \varepsilon_t,$$

  where $\Delta^d X_t$ represents differencing applied d times.

- **ARIMA Components :**
  - ϕ_i : Coefficients of the AR terms.
  - θ_j : Coefficients of the MA terms.
  - ε_t : Error term at time t.

ARIMA(1,1,1)

- An **ARIMA(1,1,1)** model is defined as :

$$\Delta X_t = \phi_0 + \phi_1 X_{t-1} + \theta_1 \varepsilon_{t-1} + \varepsilon_t,$$

  where $\Delta X_t = X_t - X_{t-1}$ (first differencing).

- It combines :
  - **AR(1)** : Dependence on the previous value X_{t-1}.
  - **I(1)** : Differencing to achieve stationarity.
  - **MA(1)** : Dependence on the previous error ε_{t-1}.

Random Walk with Drift

- A **Random Walk with Drift** is a model where :

$$X_t = X_{t-1} + \mu + \varepsilon_t,$$

where :

  - X_t : Value of the series at time t,
  - μ : Drift term (constant trend),
  - ε_t : Random error (white noise).

- Characteristics :
  - Without drift ($\mu = 0$), it is a simple random walk.
  - With drift, the series has a systematic upward or downward trend.

Example : Random Walk with Drift

- Suppose monthly sales data follows :

$$X_t = X_{t-1} + 5 + \varepsilon_t.$$

- Interpretation :
  - $\mu = 5$: Sales increase by 5 units per month on average.
  - ε_t : Random variations (e.g., market shocks).
- This model can be used to forecast future sales while accounting for the trend and random variations.

Practical Applications of ARIMA

- Forecasting stock prices or market indices.
- Predicting monthly sales or revenue.
- Analyzing economic indicators (e.g., inflation rates, GDP growth).

ARIMA models provide a robust framework for time series forecasting in diverse fields.

Steps to Build an ARIMA Model

Workflow

1. **Visualize the Data** : Plot the time series.
2. **Stationarize the Series** :
   - Use transformations (e.g., log).
   - Differencing.
3. **Identify Parameters (p, d, q)** :
   - Use ACF and PACF plots.
4. **Fit the Model** :
   - Use software packages like Python's statsmodels or R.
5. **Evaluate the Model** :
   - Analyze residuals.
   - Use performance metrics.
6. **Forecast** :
   - Generate predictions.

ARCH (Autoregressive Conditional Heteroskedasticity) Model

ARCH Model, introduced by Engle in 1982, models time series data with volatility clustering.
The ARCH(q) model is defined as :

$$X_t = \sigma_t \epsilon_t, \quad \epsilon_t \sim \mathcal{N}(0, 1)$$

where :

- X_t : Time series value at time t,
- σ_t^2 : Conditional variance at time t, given by :

$$\sigma_t^2 = \alpha_0 + \sum_{i=1}^{q} \alpha_i X_{t-i}^2$$

- α_0 : Constant ensuring positive variance,
- α_i : ARCH coefficients indicating the influence of past

The GARCH Model

Generalized Autoregressive Conditional Heteroskedasticity (GARCH) model, proposed by Bollerslev, extends the ARCH model by incorporating lagged values of the conditional variance itself, allowing for more persistent volatility patterns without requiring a large number of parameters.

GARCH(p,q) Model

A GARCH model of order (p,q), denoted as GARCH(p,q), is given by :

$$X_t = \sigma_t \epsilon_t, \quad \epsilon_t \sim \mathcal{N}(0,1)$$

- X_t : The value of the time series at time t,
- σ_t^2 : Conditional variance at time t,

$$\sigma_t^2 = \alpha_0 + \sum_{i=1}^{q} \alpha_i X_{t-i}^2 + \sum_{j=1}^{p} \beta_j \sigma_{t-j}^2$$

- α_0 : Constant term,
- α_i : ARCH coefficients,

Advantages and Limitations

Advantages :

- More parsimonious than high-order ARCH models.
- Better fit for financial data exhibiting volatility clustering and persistence.

Limitations :

- May not capture asymmetries in volatility.
- Negative shocks may have different impacts than positive shocks.

Extensions of the GARCH Model

To address limitations of the basic GARCH model, several extensions have been developed :

- EGARCH (Exponential GARCH)
- GJR-GARCH Model
- IGARCH (Integrated GARCH)

EGARCH (Exponential GARCH)

The **Exponential GARCH (EGARCH)** model, proposed by Nelson in 1991, allows for asymmetric effects of positive and negative shocks :

$$\log(\sigma_t^2) = \alpha_0 + \sum_{i=1}^{q} \alpha_i \frac{\epsilon_{t-i}}{\sigma_{t-i}} + \sum_{j=1}^{p} \beta_j \log(\sigma_{t-j}^2)$$

Log transformation ensures positivity of σ_t^2 *without coefficient constraints.*

GJR-GARCH Model

The **GJR-GARCH** model incorporates an indicator variable to capture asymmetric effects :

$$\sigma_t^2 = \alpha_0 + \sum_{i=1}^{q} \alpha_i X_{t-i}^2 + \sum_{j=1}^{p} \beta_j \sigma_{t-j}^2 + \sum_{i=1}^{q} \gamma_i X_{t-i}^2 I_{t-i}$$

where $I_{t-i} = 1$ if $X_{t-i} < 0$ and 0 otherwise.

Allows negative shocks to have different effects than positive shocks.

IGARCH (Integrated GARCH)

The **Integrated GARCH (IGARCH)** model is a special case where :

$$\sigma_t^2 = \alpha_0 + \sum_{i=1}^{q} \alpha_i X_{t-i}^2 + \sum_{j=1}^{p} \beta_j \sigma_{t-j}^2$$

with $\sum_{i=1}^{q} \alpha_i + \sum_{j=1}^{p} \beta_j = 1.$

IGARCH exhibits unit root properties in variance, implying persistence of shocks.

Progression of GARCH Models

1. **ARCH Model :** Captures volatility clustering but may require high order.
2. **GARCH Model :** Adds lagged variances for a parsimonious fit.
3. **Extensions (EGARCH, GJR-GARCH, IGARCH) :** Capture asymmetries and long memory effects.

GARCH models are essential in financial econometrics for modeling volatility dynamics in asset returns, interest rates, and other economic variables.

THANK YOU FOR YOUR ATTENTION