








Time Series Analysis

Chapter 0: Introduction: components and exponential smoothing



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

Acronyms used in this material (1/2)

- **ACF**: Autocorrelation Function
- **AI**: Artificial Intelligence
- **API**: Application Programming Interface
- **AR**: AutoRegressive (model)
- **ARCH**: AutoRegressive Conditional Heteroskedasticity
- **ARFIMA**: AutoRegressive Fractionally Integrated Moving Average (model)
- **ARIMA**: AutoRegressive Integrated Moving Average (model)
- **ARMA**: AutoRegressive Moving Average (model)
- **ASE**: Academia de Studii Economice din București (the Bucharest University of Economic Studies)
- **BET**: Bucharest Exchange Trading (index)
- **BNR**: Banca Națională a României (the National Bank of Romania)
- **BVB**: Bursa de Valori București (the Bucharest Stock Exchange)
- **COVID-19**: COronaVIrus Disease 2019
- **CSIE**: Facultatea de Cibernetică, Statistică și Informatică Economică (the Faculty of Cybernetics, Statistics and Economic Informatics)
- **ECTS**: European Credit Transfer and Accumulation System
- **EOD**: End Of Day (daily closing data)
- **EODHD**: EOD Historical Data (market-data provider)
- **ETS**: Error, Trend, Seasonality (exponential smoothing state space models)

Acronyms used in this material (2/2)

- **EU**: European Union
- **FRED**: Federal Reserve Economic Data
- **GARCH**: Generalised AutoRegressive Conditional Heteroskedasticity
- **GDP**: Gross Domestic Product
- **HICP**: Harmonised Index of Consumer Prices
- **INS**: Institutul Național de Statistică — the National Institute of Statistics (Romania)
- **LLM**: Large Language Model
- **LOESS**: LOcally Estimated Scatterplot Smoothing
- **LPPL**: Log-Periodic Power Law
- **MA**: Moving Average
- **MAE**: Mean Absolute Error
- **MAPE**: Mean Absolute Percentage Error
- **MASE**: Mean Absolute Scaled Error
- **RMSE**: Root Mean Squared Error
- **SARIMA**: Seasonal ARIMA (model)
- **SES**: Simple Exponential Smoothing
- **STL**: Seasonal and Trend decomposition using Loess
- **TBATS**: Trigonometric seasonality, Box–Cox, ARMA errors, Trend, Seasonal components (model)
- **TRAMO-SEATS**: Time series Regression with ARIMA noise, Missing values and Outliers; Signal Extraction in ARIMA Time Series
- **US**: United States
- **VAR**: Vector AutoRegression
- **VAT**: Value Added Tax
- **VECM**: Vector Error Correction Model

Route of this chapter

- **Guiding questions**
  - what is a time series, and where do we meet one?
  - which patterns repeat in economic, financial, energy and climate data?
  - how do we forecast a series, and how do we know the forecast is good?
- **Part I**
  - how the course works
  - time series on real data and a short history
  - components: trend, seasonality, cycle, noise
- **Part II**
  - notation, transformations and a first look at the ACF
  - exponential smoothing; forecasts and their evaluation
  - tools, and the possible contribution of AI

- **Learning outcomes**: after this chapter you can
  - explain how the course is organised and graded
  - recognise trend, seasonality, cycle and noise in a chart
  - compute growth rates, a moving average and a sample autocorrelation
  - produce naive and exponential smoothing forecasts
  - evaluate forecasts on a test set with MAE, RMSE and MASE
- Seminar 0 comes **before** this lecture and practises the first tools on data

Course overview

- **Programme**
  - Bachelor's programmes Economic Informatics and Economic Cybernetics, year 3, semester 2
  - Faculty of Cybernetics, Statistics and Economic Informatics (CSIE), ASE Bucharest
  - 4 ECTS credits; 2 hours of lecture and 1 hour of seminar a week; academic year 2026/2027
- **Lecture**
  - Prof. Daniel Traian Pele, danpele@ase.ro
- **Prerequisites**
  - probability and statistics, econometrics (linear regression), Python
- **Course website**
  - danpele.github.io/Time-Series-Analysis: slides, seminars, notebooks, Quantlets, quizzes



The Bucharest University of Economic Studies (ASE)

Photo: Joe Mabel (2014); CC BY 3.0; Wikimedia Commons

Evaluation

- **Written exam: 70%**
  - reading and interpreting software output, short derivations, choosing a model for a given series
  - covers the lectures and seminars of Chapters 0–10 and 15
- **Team project: 20%**
  - 2–4 students analyse and forecast real time series with the methods of the course
  - a GitHub repository, a short report, a presentation and an oral defence
- **Attendance and quizzes: 10%**
  - lectures and seminars
  - one quiz per chapter on the website: 20 questions drawn from a bank of 24, each answer explained
  - taken with the ASE account; name and score are recorded for the grade

Textbooks

- **Main textbook**
  - Huang & Petukhina (2022), *Applied Time Series Analysis and Forecasting with Python*, Springer
  - the chapters of the course follow its order: components and smoothing, stationarity, ARMA, ARIMA, seasonality, volatility, multivariate models
- **Free companion**
  - Hyndman & Athanasopoulos (2021), *Forecasting: Principles and Practice* (3rd edition, online)
  - the reference for decomposition, exponential smoothing and forecast evaluation (Chapter 0)
- **Theory**
  - Brockwell & Davis (2016), *Introduction to Time Series and Forecasting*
  - Hamilton (1994), *Time Series Analysis*: the standard reference of time series econometrics

Course map: chapters 0–15

No.	Chapter
0	Introduction: components and exponential smoothing
1	Stochastic processes and stationarity
2	ARMA models
3	Unit roots and ARIMA models
4	Seasonality and forecasting: SARIMA, TBATS, Prophet
5	Conditional volatility: ARCH and GARCH
6	VAR models and Granger causality
7	Cointegration and VECM

No.	Chapter
8	Long memory and ARFIMA
9	Machine learning for time series
10	State space models, Kalman filter and Markov switching
11	Foundation models for time series*
12	Spectral analysis*
13	Speculative bubbles: LPPL models*
14	Multivariate GARCH models*
15	Review and exam preparation

- Univariate models (0–5), several series (6, 7), extensions (8–10), review (15)
- Self-study, marked *: Chapters 11–14 (shorter materials and a quiz, no compulsory seminar)

Materials and tools

- **For every chapter**
  - lecture and seminar slides, in Romanian and in English
  - a lecture notebook and a seminar notebook (Python), opened in Google Colab with one click
  - a quiz of 20 questions
- **Quantlets**: every chart has a public folder with its code, a description file (`Metainfo.txt`) and the chart
  - the icon under each chart opens its Quantlet on GitHub
- **Data**: public sources, read directly in the code
  - Eurostat, INS (National Institute of Statistics), BNR (National Bank of Romania), FRED (Federal Reserve Bank of St. Louis)
  - daily market data from EODHD (EOD Historical Data), saved in the course repository until 18 September 2026

Seminars

- **Three parts**
  - A: short calculations on paper; B: real data, each task ending with an interpretation question; C: an open idea and the critique of an AI answer
- **Two kinds of exercises**
  - **[Solved]**: full solution in the slides and in the notebook, a model to follow
  - **[Proposed]**: you solve it; the solution is discussed in class
- **Nothing is handed in**: the seminar is practice for the exam and for the project

The team project

- **Question**: one forecasting or modelling question on real time series, chosen by the team
- **Methods**: those of the course
  - decomposition and smoothing, ARIMA and SARIMA, GARCH, VAR or VECM, Granger causality, impulse responses, forecast evaluation
  - Romanian data (INS, BNR, Eurostat, BVB) are encouraged
- **Deliverables**
  - a GitHub repository whose code reproduces every number and every chart
  - a short report, a presentation and the file `AI_USE.md`
- **Grading**
  - a clear question, correct methods and diagnostics, an honest out-of-sample evaluation
  - **oral defence**: each member explains the code and the results

AI policy

- **AI tools are allowed and must be declared**
  - AI: artificial intelligence; here, assistants based on an LLM (large language model) that write text and code
  - every use goes into `AI_USE.md`: the tool, the prompt, what was kept, what was corrected
- **You are responsible for every number and every reference**
  - typical AI errors: invented references, wrong formulas, code that runs but computes something else
  - every reference: open its DOI (Digital Object Identifier) and check the title
  - every number: recompute it with your own code
- **Oral defence**
  - a line of code you cannot explain does not count as your work

Definition

- **Time series**: observations of one variable, recorded at successive moments and ordered in time
- **Notation**: $y_1, y_2, \ldots, y_T$
  - t: the time index (quarter, month, day); T: the number of observations
  - **frequency**: the number of observations per year (4 quarterly, 12 monthly, about 252 for trading days)
- **Example**: Romanian real GDP, quarterly, billion EUR at 2010 prices
  - 2024 Q1: 38.7; 2024 Q2: 45.6; 2024 Q3: 52.8; 2024 Q4: 59.2; 2025 Q1: 38.9
  - the order matters: shuffling the five numbers destroys the information
- GDP: gross domestic product; Q1–Q4: the four quarters of a year

Three kinds of data

Type	Structure	Example
Cross-section	many units, one moment	GDP of the 27 EU countries in 2025
Time series	one unit, many moments	Romanian GDP, every quarter since 1995
Panel	many units, many moments	GDP of the 27 EU countries, every quarter

- **What changes in a time series**
  - observations are **not independent**: this quarter resembles the previous one
  - this dependence is the information we use to forecast
  - the classical formulas for independent samples (standard errors, tests) must be adapted
- EU: European Union

Why time series matter

Field	Series	Typical question
Economics	GDP, inflation, unemployment	Where will inflation be in a year?
Finance	exchange rates, stock indices, interest rates	How risky is the BET tomorrow?
Energy	electricity production and consumption, prices	How much electricity is needed next January?
Climate	CO2 concentration, temperatures	How fast does CO2 rise, beyond the seasonal swing?

- Decisions depend on forecasts: the central bank sets the interest rate on its inflation forecast; the grid operator plans production on its demand forecast

Four goals of time series analysis

- **Describe**: trend, seasonality, cycles, unusual observations
  - charts, decomposition, autocorrelation (this chapter and Chapter 1)
- **Model**: a probability model that could have produced the data
  - ARMA, ARIMA, GARCH (Chapters 2–5)
- **Forecast**: values not yet observed, with their uncertainty
  - the goal that runs through the whole course
- **Explain and control**: how one series responds to another
  - VAR models, Granger causality, cointegration (Chapters 6 and 7)
- Box & Jenkins (1970) framed the same goals: identify, estimate, check, forecast

Frequency, flows and stocks, adjusted data

- **Flow**: measured over a period; **stock**: measured at a moment
  - flows: GDP of a quarter, electricity produced in a month; they add up over time
  - stocks: exchange rate on a day, CO2 concentration, an index level; we average them, we do not add them
- **Changing the frequency**
  - monthly to quarterly: the sum for flows, the average or the last value for stocks
- **Seasonally adjusted data**
  - the statistical office removes the regular within-year pattern and the calendar effects (number of working days)
  - the unadjusted series is what was measured; the adjusted one is an estimate, revised when new data arrive

Data used in this chapter

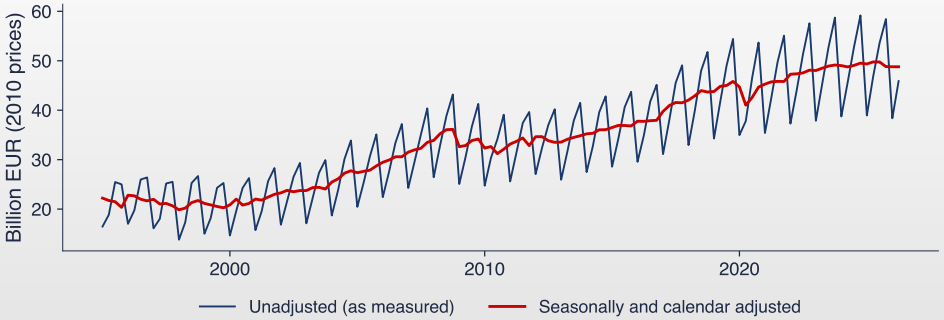
Series	Source	Frequency	Sample
Romanian real GDP	Eurostat	quarterly	1995 Q1 – 2026 Q2
Romanian HICP (consumer prices)	Eurostat	monthly	1996 – August 2026
Electricity generation, Romania	Eurostat	monthly	January 2008 – July 2026
EUR/RON	BNR reference rate	daily	2005 – 2026
BET, S&P 500	EODHD	daily	2000 – 2026
CO2 at Mauna Loa	statsmodels data set	weekly, averaged monthly	1958 – 2001

- HICP: Harmonised Index of Consumer Prices, computed with the same method in every EU country
- BET: the reference index of the BVB (Bucharest Stock Exchange); S&P 500: 500 large US companies
- statsmodels: the Python library for statistical models used throughout the course

Recap: What a time series is

- A time series is one variable observed over time; the order of the observations carries information
- Consecutive observations depend on each other: this is the basis of forecasting
- Goals: describe, model, forecast, explain
- Flows add up, stocks are averaged; adjusted series are estimates

Romanian real GDP



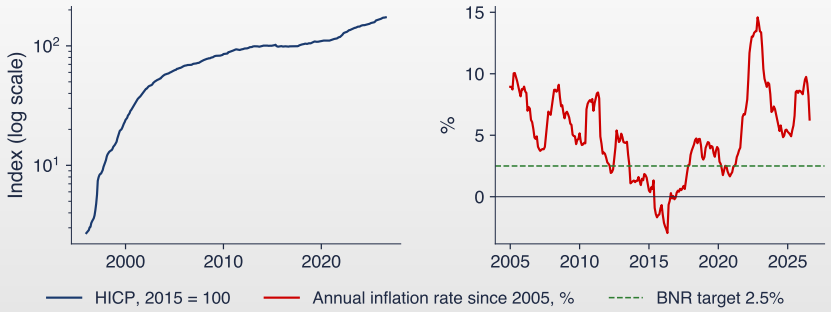
- Quarterly, 1995 Q1 – 2026 Q2, chain-linked volumes at 2010 prices (billion EUR); blue: as measured; red: seasonally and calendar adjusted (Eurostat)

TSA_ch0_examples

Interpretation: GDP

- **Trend**: the economy grew about 2.3 times from 1995 to 2025
  - an average real growth of 2.8% a year
- **Seasonality**: every year, Q1 is far below Q4
  - winter: less construction and agriculture; autumn: the harvest and year-end activity
  - the swing grows with the level of GDP: a multiplicative pattern
- **Cycle**: recessions interrupt the trend
  - annual GDP fell by 5.5% in 2009 (global financial crisis) and by 3.6% in 2020 (COVID-19 pandemic)
- Last value, 2026 Q2: 45.9 billion EUR unadjusted, 48.8 adjusted

Romanian consumer prices and inflation



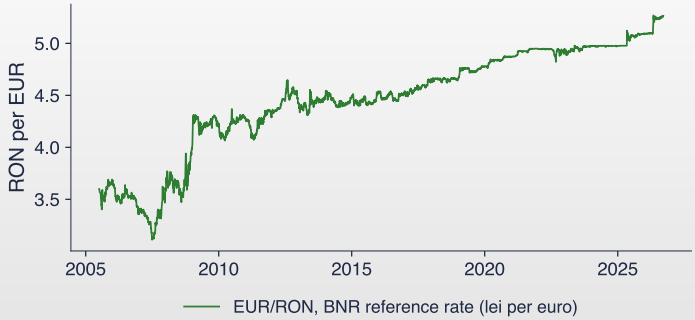
- Left: HICP, 2015 = 100, log scale; right: annual inflation rate $100\,(P_t/P_{t-\mathbf{12}} - 1)$ since 2005, against the BNR target of 2.5%; P_t: the HICP in month t

TSA_ch0_examples

Interpretation: inflation

- **Level and rate tell different stories**
  - the price index almost never falls: its trend is the accumulated inflation
  - the inflation rate rises and falls: it is the series the central bank targets
- **Episodes**
  - June 1997: 178%, after the price liberalisation of 1997
  - May 2016: -3.0%, after the VAT cuts of 2015–2016
  - November 2022: 14.6%, the energy price shock
  - August 2026: 6.3%
- A rate over 12 months removes the seasonal pattern but reacts with a delay: it reflects the whole past year
- VAT: value added tax

EUR/RON, BNR reference rate



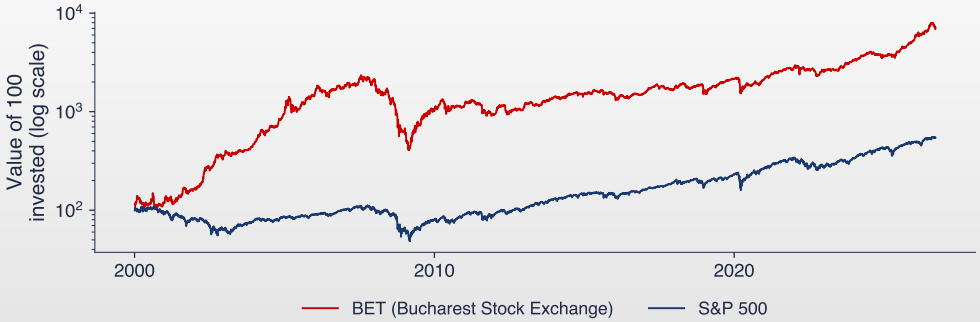
- Lei per euro, every working day from 4 July 2005 to 18 September 2026 (5 353 observations)

TSA_ch0_examples

Interpretation: EUR/RON

- **A trend without a fixed shape**
  - from 3.5986 to 5.2644 lei; lowest value 3.1112 on 2 July 2007
  - long calm periods, interrupted by jumps (autumn 2008, 2012, 2025)
- **No seasonality**
  - if the euro were predictably cheaper every spring, traders would buy it in spring and the pattern would disappear
- **Daily changes**
  - standard deviation of the daily log change: 0.28%; the largest move: 2.93% on 3 October 2008
- A trend like this one, built from accumulated shocks, is called a **stochastic trend** (Chapters 1 and 3)

BET and S&P 500 since 2000



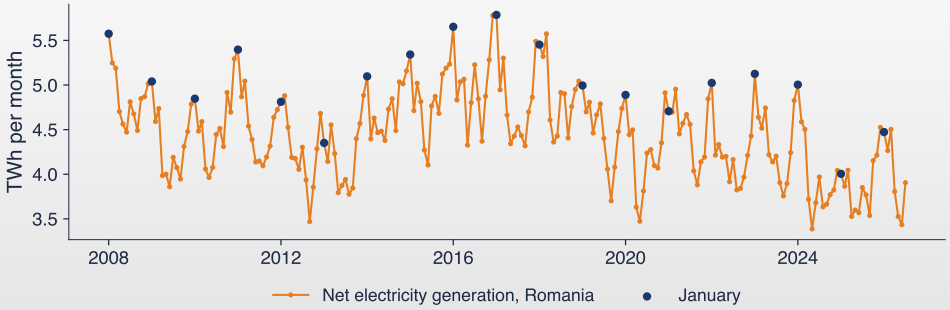
- Value of 100 invested on 5 January 2000, closing prices, log scale: equal vertical distances are equal percentage changes

TSA_ch0_examples

Interpretation: stock indices

- **Long-run growth with deep falls**
  - 100 invested grew about 71-fold in the BET and 5.5-fold in the S&P 500 (price indices, without dividends)
  - 2008–2009: the BET lost more than three quarters of its value; it took years to recover
- **Why the log scale**
  - on a linear scale, the early years look flat and recent moves look huge
  - the logarithm turns constant percentage growth into a straight line
- Prices of traded assets are hard to forecast; their **risk** is easier to forecast (Chapter 5)

Electricity generation in Romania



Net electricity generation, TWh per month, January 2008 – July 2026 (Eurostat); dots: January of each year

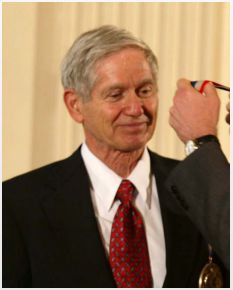
TSA_ch0_examples

Interpretation: electricity

- **Strong seasonality**
  - highest average month: January (5.03 TWh); lowest: September (4.06 TWh)
  - winter demand for heating and lighting; the ratio of the two averages is 1.24
- **A falling level**
  - 58.5 TWh in 2008, 46.7 TWh in 2025: more imports, less industrial demand
- **Irregular years**
  - dry years reduce hydropower; mild winters reduce demand
- TWh: terawatt-hour, 10^{12} watt-hours; GWh: gigawatt-hour, 10^{9} watt-hours

The Keeling curve

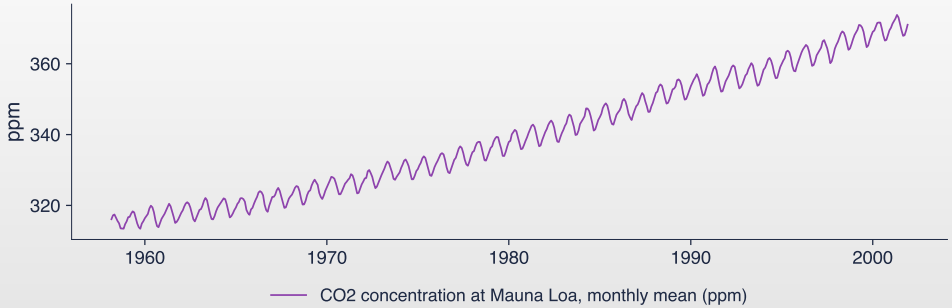
- **1958**: Charles David Keeling starts measuring CO2 at Mauna Loa, Hawaii, far from local sources (Keeling (1960))
  - the longest continuous record of atmospheric CO2
- **Two patterns in one series**
  - a rising trend: fossil fuel emissions accumulate
  - a yearly wave: plants of the Northern Hemisphere absorb CO2 in summer and release it in winter
- ppm: parts per million, molecules of CO2 per million molecules of dry air
- CO2: carbon dioxide



Charles David Keeling (1928–2005), 2001

Photo: National Science Foundation (2001); public domain; Wikimedia Commons

Atmospheric CO2 at Mauna Loa



- Monthly averages of the weekly measurements, March 1958 – December 2001 (526 months; the classic data set of statsmodels)

TSA_ch0_examples

Interpretation: CO2

- **Trend**: from 316.1 to 371.0 ppm
  - on average 1.26 ppm a year; the slope itself increases over time
- **Seasonality**: a regular yearly wave of a few ppm
  - the wave keeps roughly its size while the level rises: an additive pattern
- **Almost no noise**
  - a careful physical measurement, unlike GDP or exchange rates
- A smooth series with a clear trend and season: the easiest case for decomposition and smoothing

Patterns in six series

Series	Trend	Season	Cycle	Noise
GDP (unadjusted)	up	strong, multiplicative	recessions	small
Inflation rate	down after 1997	removed by the 12-month rate	shocks	medium
EUR/RON	stochastic	–	–	medium
BET	up, with crashes	–	booms and busts	large
Electricity	down	strong, winter peak	–	medium
CO2	up, accelerating	regular, additive	–	very small

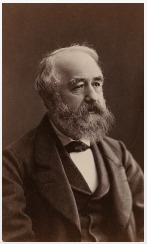
- **Question for the room**: which of these series would you forecast most accurately one year ahead?

Patterns in six series

Series	Trend	Season	Cycle	Noise
GDP (unadjusted)	up	strong, multiplicative	recessions	small
Inflation rate	down after 1997	removed by the 12-month rate	shocks	medium
EUR/RON	stochastic	–	–	medium
BET	up, with crashes	–	booms and busts	large
Electricity	down	strong, winter peak	–	medium
CO2	up, accelerating	regular, additive	–	very small

- **Question for the room**: which of these series would you forecast most accurately one year ahead?
- **Answer**: CO2: a smooth trend and a stable season, little noise; the BET is the hardest: no season and large noise

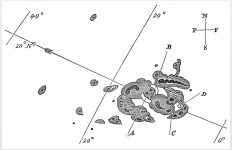
Sunspots: the first famous time series



Rudolf Wolf (1816–1893)

Photo: Emil Gassler (1888–1893); public domain; ETH-Bibliothek, Wikimedia Commons

- **1848**: the Swiss astronomer Rudolf Wolf defines a daily index of solar activity from the number of sunspots
- he reconstructs it back to 1749: the Wolf (later Wolfer) sunspot numbers
- the series rises and falls in waves of about 11 years, but neither the length nor the height of the waves is constant
- the question: a hidden periodic signal plus noise, or something else?



Sunspots drawn by Richard Carrington, 1 September 1859

Drawing: Richard Carrington (1859); public domain; Wikimedia Commons

Yule (1926, 1927): nonsense correlations and the autoregressive model

- **1926**: *nonsense correlations* between time series (Yule (1926))
  - two unrelated series with trends can show a correlation close to 1
  - his example: the share of Church of England marriages and the mortality rate in England and Wales
  - today: spurious regression (Chapter 3) and cointegration (Chapter 7)
- **1927**: Wolfer's sunspot numbers explained by their own past (Yule (1927))
  - instead of hidden sine waves: $y_t = \phi_1 y_{t-1} + \phi_2 y_{t-2} + \varepsilon_t$, with random disturbances ε_t
  - y_t: the sunspot number in year t; ϕ_1, ϕ_2: coefficients estimated from the data, which set the length and the damping of the waves
  - the first **autoregressive** model: AR(2) (Chapter 2)
  - disturbances shift the waves, so the period drifts, as in the data

Slutsky (1937): cycles from random shocks

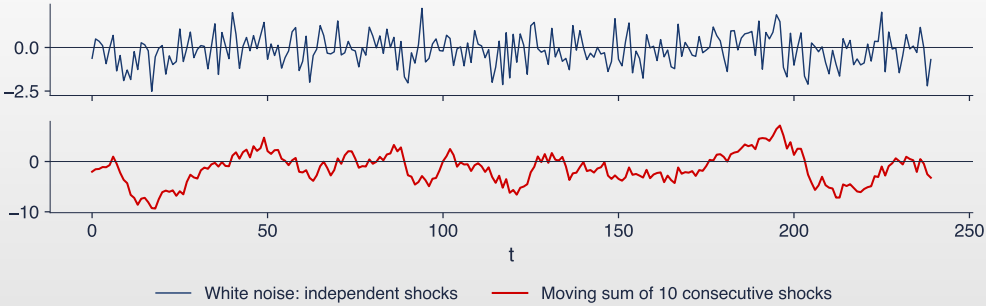


Eugen Slutsky (1880–1948)

Photo: unknown author; public domain; Wikimedia Commons

- **1927** (Moscow), **1937** in English: *The summation of random causes as the source of cyclic processes* (Slutzky (1937))
- **The experiment**
  - take independent random numbers $\varepsilon_{\mathbf{1}}, \varepsilon_{\mathbf{2}}, \ldots$ (the digits of a lottery draw)
  - replace each one by the sum of the last 10: $y_t = \varepsilon_t + \varepsilon_{t-\mathbf{1}} + \cdots + \varepsilon_{t-\mathbf{9}}$
  - the result looks like business cycles
- Today this is a **moving-average** process: MA(9) (Chapter 2)

Slutsky's experiment, simulated



Top: 240 independent shocks with the standard Normal distribution; bottom: the moving sum of 10 consecutive shocks

TSA_ch0_slutsky

Interpretation: Slutsky's experiment

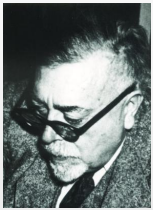
- **No cycle was put in, yet waves appear**
  - the moving sum crosses zero upwards 16 times in 240 periods: waves of irregular length
  - neighbouring values share 9 of their 10 shocks, so they are strongly correlated
- **Implications**
  - a cycle in a chart is not proof of a periodic cause
  - smoothing (moving averages) can *create* apparent cycles: be careful with smoothed data
- Yule (AR) and Slutsky (MA) gave the two building blocks of the ARMA models

Kolmogorov and Wiener: the best linear forecast



Andrei Kolmogorov (1903–1987)

Photo: Konrad Jacobs; CC BY-SA 2.0 de; Oberwolfach, Wikimedia Commons



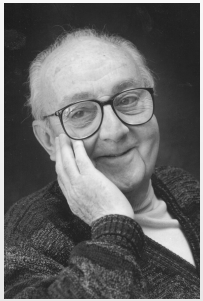
Norbert Wiener (1894–1964)

Photo: Konrad Jacobs; CC BY-SA 2.0 de; Oberwolfach, Wikimedia Commons

- **1941**: the theory of interpolation and extrapolation of stationary sequences (Kolmogorov (1941))
- **1942** (classified report), **1949**: the Wiener filter, developed for anti-aircraft fire control (Wiener (1949))
- Both answer: which weighted sum of past values forecasts the future with the smallest mean squared error? The ancestor of the Kalman filter (Chapter 10)

Box and Jenkins (1970): a method for practice

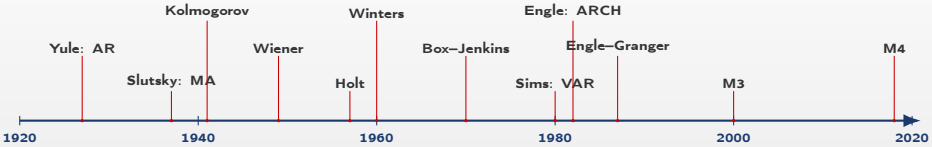
- **1970**: George Box and Gwilym Jenkins publish *Time Series Analysis: Forecasting and Control* (Box & Jenkins (1970))
- **The Box–Jenkins cycle**
  - **identify**: choose a model from charts and the ACF
  - **estimate**: fit its parameters
  - **check**: are the residuals unpredictable noise?
  - **forecast**: only when the checks pass
- ARIMA models and their seasonal version SARIMA come from this book (Chapters 2–4)
- Box: *all models are wrong, but some are useful*



George E.P. Box (1919–2013)

Photo: DavidMCEddy (2011); CC BY-SA 3.0; Wikimedia Commons

A timeline of time series analysis



- 1927 Yule, 1937 Slutsky, 1941 Kolmogorov, 1949 Wiener: the foundations (Chapters 1 and 2)
- 1957 Holt (1957), 1960 Winters (1960): exponential smoothing, born in industry for sales and inventories (this chapter)
- 1970 Box–Jenkins (Chapters 2–4); 1980 Sims (1980) (Chapter 6); 1982 Engle (1982) (Chapter 5); 1987 Engle & Granger (1987) (Chapter 7)
- 2000 and 2018: the M3 and M4 forecasting competitions compare methods on thousands of real series (Makridakis & Hibon (2000); Makridakis, Spiliotis & Assimakopoulos (2020))

Four components

- **Trend** T_t: the long-run movement of the level
  - GDP growth, the rise of CO2; it can change direction slowly
- **Seasonality** S_t: a pattern that repeats with a **fixed, known period** m
  - $m = 4$ for quarterly data, $m = 12$ for monthly data, $m = 7$ for daily data with a weekly pattern
- **Cycle** C_t: rises and falls **without a fixed period**, usually longer than a year
  - business cycles: expansions and recessions of 2–10 years; usually merged with the trend into a **trend-cycle**
- **Remainder** R_t (irregular component, noise): what is left
  - it should look unpredictable; if it does not, the decomposition missed something

Cycle or seasonality?

- **Question for the room**: electricity generation is high every January
  - is this a cycle or seasonality?

Cycle or seasonality?

- **Question for the room**: electricity generation is high every January
  - is this a cycle or seasonality?
- **Answer**: seasonality: the period is fixed (12 months) and known in advance
  - a recession arrives at an unknown moment and lasts an unknown time: that is a cycle
- **Why the difference matters**
  - seasonality can be forecast years ahead: next January will again be high
  - cycles are hard to forecast: their timing changes (Slutsky)
  - sunspots: an 11-year *cycle*, not a season, because its length varies

Additive and multiplicative models

- **Additive**: $y_t = T_t + S_t + R_t$
  - the seasonal swing has a constant size in units of y (CO2: a few ppm every year)
- **Multiplicative**: $y_t = T_t \times S_t \times R_t$
  - the seasonal swing is a constant **percentage** of the level (GDP: Q1 about a quarter below the trend)
  - S_t and R_t are factors around 1: $S_t = 0.8$ means 20% below the trend
- **The logarithm links them**: $\ln y_t = \ln T_t + \ln S_t + \ln R_t$
  - a multiplicative series becomes additive after taking logs
- **How to choose**: look at the chart
  - does the seasonal swing grow with the level? Multiplicative (or logs); otherwise additive

Centred moving averages

- **Moving average of order** k (odd): the mean of k neighbouring values
  - $\hat{T}_t = \frac{1}{k}\sum_{j=-q}^{q} y_{t+j}$, with $k = 2q + 1$; it smooths out the noise
  - q: the number of values on each side of t; j: the position relative to t; the hat marks an estimate
- **The key property**: a moving average over one full season removes the seasonal pattern
  - every quarter enters once, so the seasonal effects cancel out
- **Even period** $m = 4$: a 2×4 moving average, to stay centred on a quarter
  - $\hat{T}_t = \frac{1}{8}y_{t-2} + \frac{1}{4}y_{t-1} + \frac{1}{4}y_t + \frac{1}{4}y_{t+1} + \frac{1}{8}y_{t+2}$
  - five quarters, the two ends with half weight; for monthly data, 2×12
- The first and the last $m/2$ values have no centred average: the trend estimate is missing at both ends

Worked example: a 2×4 moving average

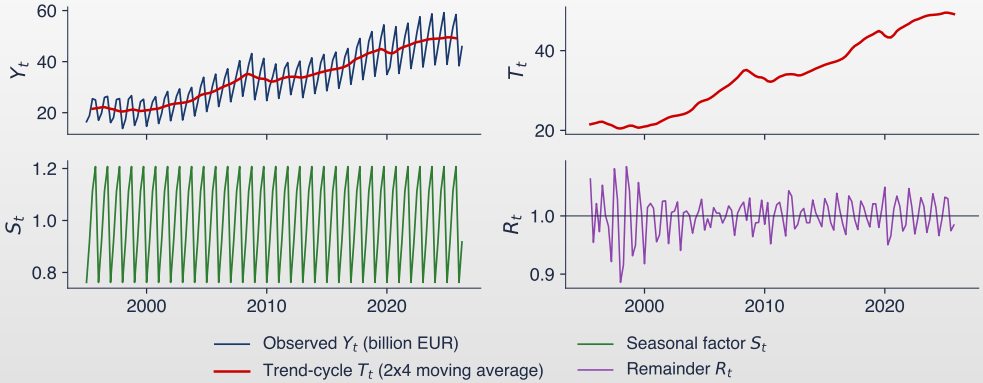
Quarter	y_t
2024 Q1	38.7
2024 Q2	45.6
2024 Q3	52.8
2024 Q4	59.2
2025 Q1	38.9

- Romanian real GDP, billion EUR (Eurostat, unadjusted)
- **Trend in 2024 Q3**
  - $\hat{T}_t = \frac{1}{4}\left(\frac{38.7}{2} + 45.6 + 52.8 + 59.2 + \frac{38.9}{2}\right) = 49.10$
- **Seasonal ratio**
  - $y_t / \hat{T}_t = 52.8/49.10$: the third quarter is above the trend
- Averaging such ratios over all years gives the seasonal factor of each quarter

Classical decomposition, step by step

- **Step 1**: trend-cycle $\hat{T}_t$ by a centred $2 \times m$ moving average
- **Step 2**: detrend
  - multiplicative: $y_t / \hat{T}_t$; additive: $y_t - \hat{T}_t$
- **Step 3**: seasonal factors $\hat{S}_1, \ldots, \hat{S}_m$
  - average the detrended values of each season over all years
  - normalise: they average to 1 (multiplicative) or to 0 (additive)
- **Step 4**: remainder $\hat{R}_t = y_t / (\hat{T}_t \hat{S}_t)$ or $y_t - \hat{T}_t - \hat{S}_t$
- **Seasonally adjusted series**: $y_t / \hat{S}_t$ or $y_t - \hat{S}_t$
  - limits: the same season every year, no trend at the ends, sensitive to outliers

Decomposition of the Romanian GDP



- Classical multiplicative decomposition, $m = 4$: observed series and trend-cycle, trend-cycle, seasonal factors, remainder

TSA_ch0_decomposition

Interpretation: seasonal factors of GDP

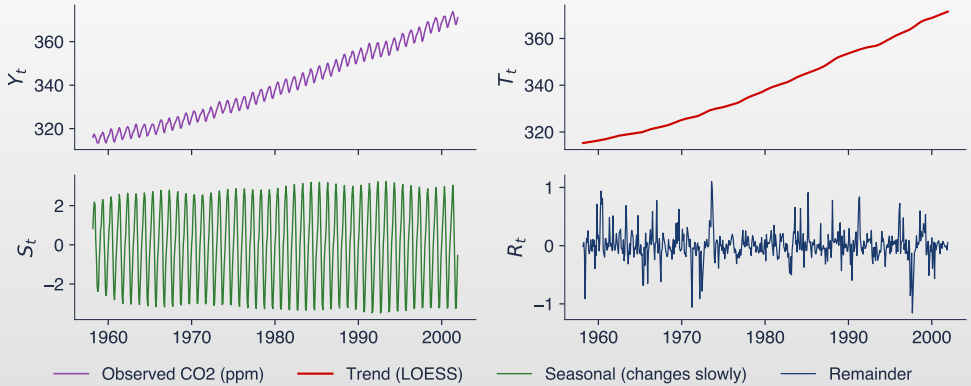
Quarter	$\hat{S}_q$	relative to trend
Q1	0.762	-23.8%
Q2	0.918	-8.2%
Q3	1.113	$+11.3\%$
Q4	1.207	$+20.7\%$

- **Reading the factors**
  - $\hat{S}_q$: the estimated seasonal factor of quarter q
  - a typical first quarter is 23.8% below the trend, a typical fourth quarter 20.7% above it
  - Q4 is about 58% larger than Q1 for purely seasonal reasons
- **Remainder**
  - standard deviation about 3.3%; larger in the 1990s and in 2020
- Consequence: compare a quarter with the same quarter of the previous year, never with the previous quarter of the raw series

STL: a flexible decomposition

- **STL**: Seasonal and Trend decomposition using LOESS (Cleveland et al. (1990))
  - LOESS (locally estimated scatterplot smoothing): a weighted regression fitted around each point
- **Advantages over the classical method**
  - the seasonal pattern may change slowly from year to year
  - a trend estimate up to the last observation
  - a robust version: outliers go into the remainder and do not distort the season
- **Limits**
  - additive only: use it on $\ln y_t$ for multiplicative series
  - no calendar effects (working days, Easter)
- Statistical offices use X-13ARIMA-SEATS or TRAMO-SEATS, which build on ARIMA models (Chapter 4)

STL decomposition of CO2



- Robust STL with period 12 on the monthly CO2 series: observed, trend, seasonal, remainder (ppm)

TSA_ch0_decomposition

Interpretation: STL of CO2

- **Trend**: smooth, slightly convex: the rise accelerates
- **Seasonal**: maximum in May, minimum in October
  - winter and spring: decay of fallen leaves releases CO2; summer: photosynthesis absorbs it
  - the yearly range grows from 4.9 ppm (1959) to 6.3 ppm (2001): STL lets the season change
- **Remainder**: standard deviation 0.25 ppm
  - small compared with the season: trend and season explain almost everything
  - check the larger spikes against the history of the station before calling them noise

Recap: Components and decomposition

- Trend, seasonality (fixed period), cycle (no fixed period), remainder
- Additive if the seasonal swing is constant, multiplicative if it grows with the level; logs turn one into the other
- A centred $2 \times m$ moving average removes the season and estimates the trend-cycle
- Romanian GDP: Q1 about 23.8% below the trend, Q4 about 20.7% above it
- STL lets the seasonal pattern change and is robust to outliers

Notation used in the course

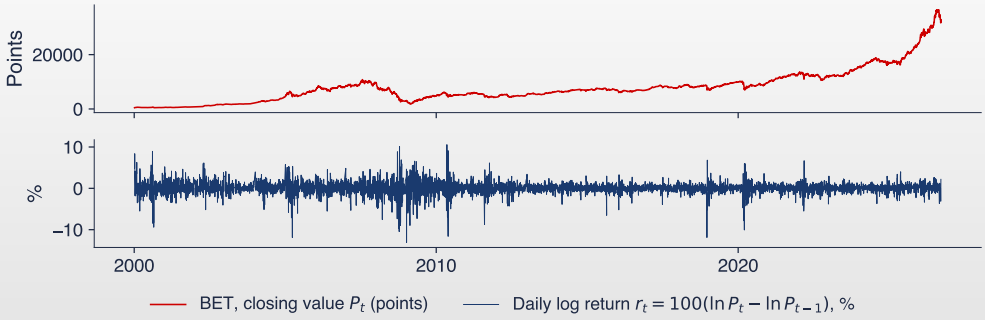
Symbol	Meaning
y_t, $t = 1, \ldots, T$	the observation at time t; T observations in the sample
$\bar{y} = \frac{1}{T}\sum_t y_t$	the sample mean
$Ly_t = y_{t-1}$	the **lag operator**: shifts the series back one period; $L^k y_t = y_{t-k}$
$\Delta y_t = y_t - y_{t-1}$	the **first difference**, $\Delta = 1 - L$
$\Delta_m y_t = y_t - y_{t-m}$	the **seasonal difference**, $\Delta_m = 1 - L^m$
$\hat{y}_{T+h\|T}$	the forecast of y_{T+h} made with the data up to T; h is the **horizon**
$e_t = y_t - \hat{y}_t$	the forecast error

- Hats ($\hat{}$) mark estimates and forecasts; Greek letters (ϕ, θ, α) mark parameters

Transformations: logs, differences, growth rates

- **Logarithm** $\ln y_t$: stabilises a variance that grows with the level
- **Growth rate over one period**: $g_t = 100\,(y_t / y_{t-\mathbf{1}} - 1)$
  - log approximation: $100\,\Delta \ln y_t \approx g_t$ for small changes
- **Annual growth rate** of a quarterly series: $100\,(y_t / y_{t-\mathbf{4}} - 1)$
  - it compares the same quarters, so the season cancels out
- **Worked example**: Romanian GDP, 2024 Q1 = 38.7, 2024 Q4 = 59.2, 2025 Q1 = 38.9
  - over one quarter: $100\,(38.9/59.2 - 1) \approx -34\%$: not a collapse, only winter
  - over one year: $100\,(38.9/38.7 - 1) \approx +0.5\%$: the economy barely grew
- **Log return** of a price: $r_t = 100\,(\ln P_t - \ln P_{t-\mathbf{1}})$, in %; P_t: the price on day t

From prices to returns: the BET



- Top: closing value of the BET since 2000; bottom: daily log returns in % (6 670 trading days)

TSA_ch0_acf

Interpretation: prices and returns

- **The price has a trend; the returns do not**
  - returns fluctuate around a mean of 0.064% a day, with a standard deviation of 1.40% (about 22% a year)
- **Extreme days**
  - worst: −13.1% on 7 January 2009; best: +10.6% on 10 May 2010
- **Calm and turbulent periods alternate**
  - large moves cluster in 2008–2010 and in March 2020: *volatility clustering* (Chapter 5)
- Differencing the log price removes the trend: most models of the course work on such transformed series

The sample autocorrelation function

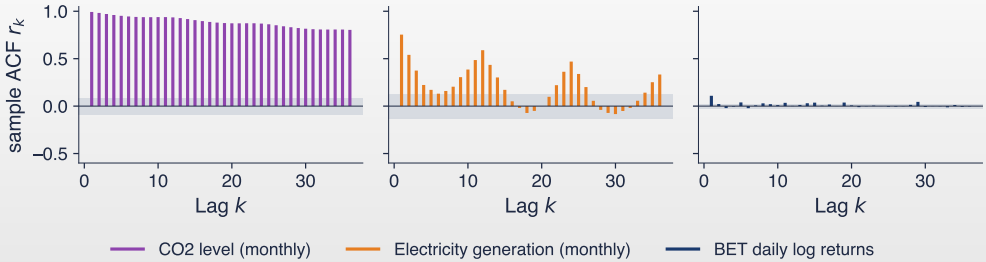
- **Question**: how strongly is y_t related to its own value k periods earlier?
- **Sample autocovariance** at lag k
  - $c_k = \frac{1}{T}\sum_{t=k+1}^{T}(y_t - \bar{y})(y_{t-k} - \bar{y})$
  - k: the lag, in periods; $\bar{y}$: the sample mean; the product is positive when y_t and y_{t-k} lie on the same side of the mean
- **Sample autocorrelation**: $r_k = c_k / c_0$, between -1 and 1
  - c_0: the sample variance (the autocovariance at lag 0)
  - the **ACF** (autocorrelation function) is the sequence $r_1, r_2, \ldots$ plotted against k (the **correlogram**)
- **Reference band** $\pm 1.96/\sqrt{T}$
  - if the series were independent noise, about 95% of the r_k would fall inside it; 1.96: the 97.5% quantile of the standard Normal distribution
  - the theory (stationarity, white noise, the Ljung–Box test) comes in Chapter 1

Worked example: r_1 by hand

t	y_t	$y_t - \bar{y}$	$(y_t - \bar{y})(y_{t-1} - \bar{y})$
1	2	-2	–
2	4	0	0
3	3	-1	0
4	5	1	-1
5	4	0	0
6	6	2	0
sum	24	0	-1

- $T = 6$, $\bar{y} = 24/6 = 4$
- sum of squared deviations: $4 + 0 + 1 + 1 + 0 + 4 = 10$
- $r_1 = \dfrac{-1/6}{10/6} = -0.1$
- at lag 2: products $(-1)(-2) + (1)(0) + (0)(-1) + (2)(1) = 4$, so $r_2 = 0.4$
- **Interpretation**
  - band $\pm 1.96/\sqrt{6} = \pm 0.80$: with 6 observations, neither value is distinguishable from 0

Three correlograms



- Sample ACF for lags 1–36; shaded: the band $\pm 1.96/\sqrt{T}$. Left: CO2 level; centre: monthly electricity generation; right: BET daily log returns

TSA_ch0_acf

Interpretation: three correlograms

- **Trend**: CO2, $r_1 = 0.993$, still 0.80 at lag 36
  - a slow, almost linear decay is the signature of a trend: the level remembers its past for years
- **Seasonality**: electricity, $r_6 = 0.13$, $r_{12} = 0.59$
  - peaks at 12, 24 and 36 months: this January resembles last January
- **Almost no memory**: BET returns, $r_1 = 0.109$, $r_2 = 0.022$, band ± 0.024
  - small but significant autocorrelations: hard to exploit after trading costs
  - the *absolute* returns have $r_1 = 0.40$: the size of moves is predictable, their sign much less (Chapter 5)
- The ACF of a level with a trend says little: transform first, then look at the ACF

Recap: Notation, transformations and the ACF

- $Ly_t = y_{t-1}$, $\Delta = 1 - L$, $\Delta_m = 1 - L^m$; $\hat{y}_{T+h|T}$ is the forecast for horizon h
- Logs stabilise the variance; differences remove the trend; annual rates remove the season
- $r_k = c_k/c_0$, compared with the band $\pm 1.96/\sqrt{T}$
- Slow decay: trend; peaks at $m, 2m, \ldots$: seasonality; all close to 0: little linear memory

Simple exponential smoothing (SES)

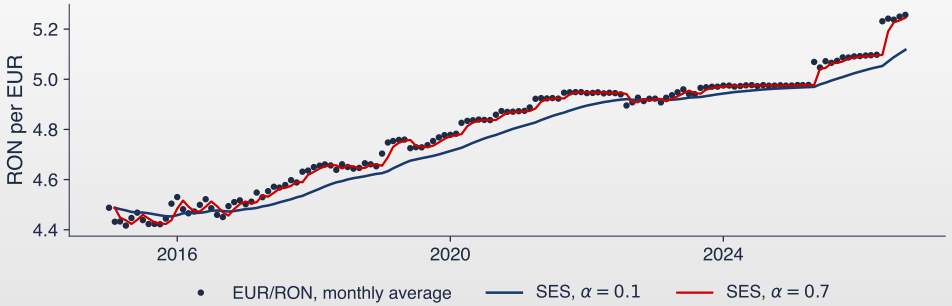
- **Idea**: forecast with a weighted average of the past, recent values weighing more
- **Level equation**: $\ell_t = \alpha y_t + (1-\alpha)\,\ell_{t-1}$, with $0 < \alpha \leq 1$
  - ℓ_t: the smoothed level at time t; α: the smoothing constant; ℓ_t moves from ℓ_{t-1} towards y_t by the fraction α
  - forecast for every horizon: $\hat{y}_{T+h|T} = \ell_T$, a flat line
- **Why "exponential"**: substituting repeatedly
  - $\ell_T = \alpha y_T + \alpha(1-\alpha) y_{T-1} + \alpha(1-\alpha)^2 y_{T-2} + \cdots$
  - the weights fall geometrically; with $\alpha = 0.5$: 0.5, 0.25, 0.125, ...
- **The smoothing constant** α
  - small α: long memory, smooth forecasts, slow reaction; $\alpha = 1$: the naive forecast $\hat{y}_{T+1} = y_T$
  - estimated by minimising the sum of squared one-step forecast errors

Worked example: SES with $\alpha = 0.5$

t	y_t	$\hat{y}_{t\mid t-\mathbf{1}} = \ell_{t-\mathbf{1}}$	$\ell_t = 0.5\, y_t + 0.5\, \ell_{t-\mathbf{1}}$
0	–	–	$\ell_{\mathbf{0}} = 10$
1	10	10	10
2	12	10	11
3	11	11	11
4	13	11	12
5	?	**12**	–

- initial level $\ell_{\mathbf{0}} = 10$ (the first observation)
- $t = 2$: $0.5 \times 12 + 0.5 \times 10 = 11$
- $t = 4$: $0.5 \times 13 + 0.5 \times 11 = 12$
- **Interpretation**
  - the forecast for $t = 5$ and every later period is 12: SES does not extrapolate trends
  - one-step errors: 0, 2, 0, 2: the forecasts lag behind the rising data

SES on the monthly EUR/RON



- Dots: monthly averages of the BNR reference rate, January 2015 – September 2026; lines: one-step SES forecasts $\hat{y}_{t|t-1}$ with $\alpha = 0.1$ and $\alpha = 0.7$

TSA_ch0_smoothing

Interpretation: the choice of α

- $\alpha = 0.1$: smooth, but always below the rising rate
  - it remembers several years: after the jump of 2025 it needs many months to catch up
- $\alpha = 0.7$: follows the data closely, one month late
- **Estimated** $\hat{\alpha} = 1.00$
  - the best SES forecast of next month is this month's average: SES reduces to the naive forecast
  - typical of exchange rates and prices: the last value contains almost all the information (a random walk, Chapter 1)
- An estimated α near 1 is a finding about the series, not a failure of the method

Holt and Holt–Winters (1/2)

- **Holt (1957)**: level plus a **trend** b_t (Holt (1957))
  - $\ell_t = \alpha y_t + (1-\alpha)(\ell_{t-1} + b_{t-1})$, $b_t = \beta(\ell_t - \ell_{t-1}) + (1-\beta) b_{t-1}$
  - the level is updated towards y_t from the previous level plus the previous slope
  - b_t: the estimated slope, the change of the level per period; it is updated towards the latest change $\ell_t - \ell_{t-1}$
  - $\beta \in (0, 1]$: the smoothing constant of the slope; a small β gives a slope that changes slowly
- **Forecast**: $\hat{y}_{T+h|T} = \ell_T + h\, b_T$, a straight line
  - the last level plus h times the last slope
  - a **damped** trend multiplies the slope by ϕ^j, $0 < \phi < 1$, at step j: the line flattens for long horizons

Holt and Holt–Winters (2/2)

- **Holt–Winters (1960)**: adds seasonal factors s_t with period m (Winters (1960))
  - multiplicative form: $\hat{y}_{T+h|T} = (\ell_T + h\, b_T)\, s_{T+h-m}$
  - s_{T+h-m}: the latest seasonal factor of the same season (for $h \leq m$); $s = 1.1$ means 10% above the trend line
  - each component has its own smoothing constant: α (level), β (slope), γ (season)
- **ETS**: Error, Trend, Seasonal (Hyndman, Koehler, Snyder & Grose (2002))
  - each method is a statistical model: N (none), A (additive), M (multiplicative), A_d (damped)
  - e.g. ETS(M,A_d,M): multiplicative errors, damped trend, multiplicative season; it gives forecast intervals and an information criterion (Chapter 2)
- A review of 25 years of exponential smoothing: Gardner (1985)

Recap: Exponential smoothing

- SES: $\ell_t = \alpha y_t + (1-\alpha)\ell_{t-1}$; weights decline geometrically; flat forecasts
- Small α: smooth and slow; $\alpha = 1$: the naive forecast
- EUR/RON: $\hat{\alpha} = 1.00$, the last value is the best SES forecast
- Holt adds a trend, Holt–Winters a season; ETS turns them into statistical models

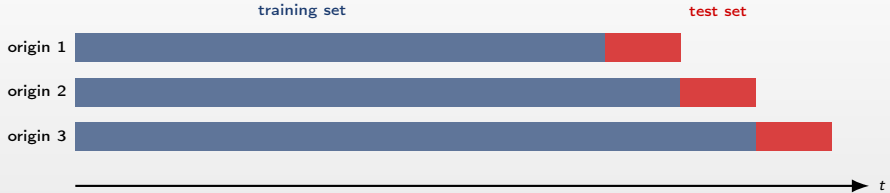
Forecasting: the setting

- **Forecast** $\hat{y}_{T+h|T}$: a statement about y_{T+h} made at time T
  - it may use only the information available at T: the **information set**
  - $h = 1, 2, \ldots$: the **horizon**; uncertainty grows with h
- **Point and interval forecasts**
  - point: one number; interval: a range that contains y_{T+h} with a stated probability (e.g. 95%)
- **What makes a series forecastable**
  - a stable pattern (CO2) rather than one driven by news (exchange rates)
  - a forecast that does not change the outcome: a price forecast published to millions of traders changes the price

Four benchmark methods

- **Mean**: $\hat{y}_{T+h|T} = \bar{y}$
  - the future looks like the average past
- **Naive**: $\hat{y}_{T+h|T} = y_T$
  - the last value; optimal for a random walk (exchange rates, prices)
- **Seasonal naive**: $\hat{y}_{T+h|T} = y_{T+h-m}$ (for $h \leq m$)
  - next January = last January; strong for highly seasonal series
- **Drift**: $\hat{y}_{T+h|T} = y_T + h\,\dfrac{y_T - y_1}{T-1}$
  - the last value plus the average past change per period
- A new model is worth using only if it beats these benchmarks on data it has not seen

Training set and test set



- **Split by time, never at random**
  - training set: the first observations, used to choose and fit the method
  - test set: the last h observations, kept hidden and used only to measure the errors
  - a random split lets the model see the future: the errors look too small
- **Rolling origin** (time series cross-validation): repeat with the origin moved forward
  - average the errors over origins: a more stable estimate of accuracy

Measures of forecast error

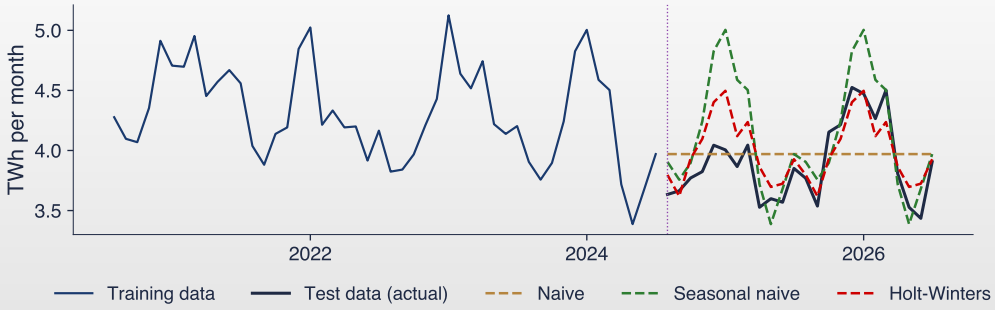
- Errors on the test set: $e_{T+j} = y_{T+j} - \hat{y}_{T+j|T}$, $j = 1, \ldots, h$
- **MAE** (mean absolute error) $= \frac{1}{h} \sum_j |e_{T+j}|$
  - in the units of the series; easy to explain
- **RMSE** (root mean squared error) $= \sqrt{\frac{1}{h} \sum_j e_{T+j}^2}$
  - penalises large errors more; always $\geq$ MAE
- **MAPE** (mean absolute percentage error) $= \frac{100}{h} \sum_j |e_{T+j}/y_{T+j}|$
  - without units, but useless when y is close to 0 (inflation, returns)
- **MASE** (mean absolute scaled error) = MAE / in-sample MAE of the naive method (Hyndman & Koehler (2006)); for seasonal series, of the seasonal naive method (as in M4)
  - below 1: better than that benchmark on the training data; comparable across series

Worked example: errors by hand

	Q1	Q2	Q3	Q4
year 1	10	14	18	12
year 2	11	15	20	13
year 3 (test)	12	16	21	14
naive	13	13	13	13
seasonal naive	11	15	20	13

- **Naive**: errors $-1, 3, 8, 1$
  - MAE $= (1 + 3 + 8 + 1)/4 = 3.25$
  - RMSE $= \sqrt{(1 + 9 + 64 + 1)/4} = \sqrt{18.75} = 4.33$
- **Seasonal naive**: errors $1, 1, 1, 1$
  - MAE $=$ RMSE $= 1$
- **Interpretation**
  - the naive forecast misses the season: its large Q3 error dominates the RMSE
  - on seasonal data, the seasonal naive method is the benchmark to beat

Forecasting electricity generation



- Test set: August 2024 – July 2026 (24 months); training data: all earlier months; dashed: forecasts made once, at the end of the training set

TSA_ch0_forecast

Interpretation: comparison of forecast accuracy

Method	MAE	RMSE	MAPE	MASE
Naive	0.28	0.33	7.3%	0.82
Seasonal naive	0.28	0.38	7.2%	0.83
SES	0.28	0.32	7.1%	0.81
Holt–Winters	0.17	0.21	4.5%	0.51

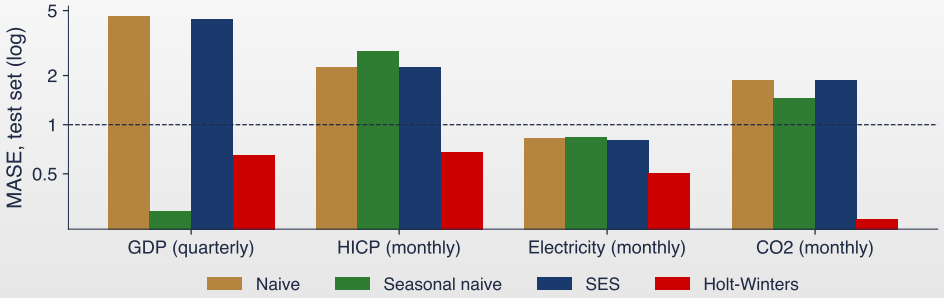
- MAE and RMSE in TWh per month

- **Holt–Winters** has the smallest error on every measure
  - it combines a falling level with the winter peak
- **The seasonal naive method is not better than the naive one here**
  - it copies the high winter of 2023–2024 into a lower 2024–2026
- One test period of 24 months is one sample: a rolling origin gives a firmer answer

Case study: the M4 competition

- **The question**: which forecasting methods work best on real data?
  - the M competitions, started by Spyros Makridakis in 1982: forecasts submitted blind, scored on hidden test data
- **M4 (2018)**: 100 000 series (yearly to hourly), 61 methods (Makridakis, Spiliotis & Assimakopoulos (2020))
  - 12 of the 17 most accurate methods were **combinations** of several methods
  - the winner was a hybrid of exponential smoothing and a neural network
  - pure machine learning methods did poorly: none beat a simple combination of exponential smoothing methods
- **Why it matters for this course**
  - benchmarks first; accuracy measured out of sample, on many series rather than one
  - the same lesson appeared in M3 (Makridakis & Hibon (2000)): complex is not automatically better

A small M-Competition on Romanian data



- MASE on the test set (log scale) of four methods on four seasonal series; test sets: the last 8 quarters (GDP) or 24 months; dashed line: MASE = 1

TSA_ch0_forecast

Interpretation: the mini-competition

- **Holt–Winters wins on 3 of 4 series**
  - CO2: MASE 0.26 against 1.45 for the seasonal naive method; HICP: 0.68 against 2.24 for the naive one
- **GDP: the seasonal naive method wins**, MASE 0.29 against 0.65
  - growth stalled in 2025–2026: the trend learnt by Holt–Winters overshoots
- **Naive and SES fail on seasonal series**
  - GDP: MASE 4.62 (naive) and 4.39 (SES): they ignore the season
- As in M4: no method wins everywhere, and a simple benchmark can beat a better-looking model

Recap: Forecasts and their evaluation

- $\hat{y}_{T+h|T}$ uses only the information up to T; uncertainty grows with h
- Benchmarks: mean, naive, seasonal naive, drift
- Split by time; measure MAE, RMSE, MAPE and MASE on the test set only
- M4 and our mini-competition: simple methods are strong; combinations and Holt–Winters do well

Python for time series

- **The libraries of the course**
  - **NumPy**: arrays and linear algebra (Harris et al. (2020))
  - **pandas**: tables indexed by date, resampling, lags, differences (McKinney (2010))
  - **matplotlib**: charts
  - **statsmodels**: decomposition, exponential smoothing, ARIMA, VAR, unit-root tests (Seabold & Perktold (2010))
  - **arch**: GARCH models (Chapter 5)
- **Google Colab**: notebooks in the browser, nothing to install
  - open a course notebook with one click; *File* → *Save a copy in Drive* keeps your changes
  - every chart of the slides has a Quantlet: the same code, runnable in Colab

A first script

```
import pandas as pd
from statsmodels.tsa.seasonal import seasonal_decompose
from statsmodels.tsa.holtwinters import ExponentialSmoothing

# Romanian real GDP, quarterly (Eurostat)
url = ('https://ec.europa.eu/eurostat/api/dissemination/'
       'sdmx/2.1/data/namq_10_gdp/Q.CLV10_MEUR.NSA.B1GQ.RO'
       '?format=SDMX-CSV')
t = pd.read_csv(url)
y = pd.Series(t['OBS_VALUE'].values / 1000,
              index=pd.PeriodIndex(t['TIME_PERIOD'], freq='Q'))
y = y.to_timestamp()

dec = seasonal_decompose(y, model='multiplicative', period=4)
print(dec.seasonal.iloc[:4])            # seasonal factors
yoy = 100 * (y / y.shift(4) - 1)        # annual growth, %
hw = ExponentialSmoothing(y, trend='add', seasonal='mul',
                          seasonal_periods=4).fit()
print(hw.forecast(4))                   # next four quarters
```

- **Line by line**
- read the series directly from the Eurostat API (no key)
- `PeriodIndex`: quarters as the time index
- classical multiplicative decomposition with $m = 4$
- `shift(4)`: the value of the same quarter a year earlier
- Holt–Winters with additive trend and multiplicative season
- API: application programming interface

Possible contribution of AI

- **Where an AI assistant helps**
  - a first draft of the code that reads a Eurostat or BNR series and plots it
  - explaining an error message or a statsmodels function
  - listing candidate explanations of a pattern (a drop in GDP, a jump in EUR/RON)
- **Time-series errors that AI code makes often**
  - a random train/test split, which lets the model see the future
  - growth over one quarter called "annual growth"; `shift(1)` instead of `shift(4)`
  - an RMSE without the square root; a MAPE on a series that crosses zero
- Seminar 0, exercise C2: find three such errors in an AI answer

Checks before trusting a result

- **The data**
  - plot the series first; check the frequency, the units, the first and the last date
  - adjusted or unadjusted? A model for the season needs the unadjusted series
- **The evaluation**
  - the test set comes after the training set; it is used once, at the end
  - compare with the naive and seasonal naive forecasts before claiming success
- **The claims**: the checks of the AI policy (references, numbers, `AI_USE.md`)

Key takeaways

- A time series is ordered in time; its dependence on the past is what makes forecasting possible
- Components: trend, seasonality (fixed period), cycle (no fixed period), remainder; additive or multiplicative
- Yule and Slutsky: dependence and sums of shocks, the roots of AR and MA models
- The ACF reveals trend (slow decay), season (peaks at m) and memory
- Exponential smoothing weighs the recent past more; evaluate every forecast against naive benchmarks on a test set

Self-assessment, project idea and next: Chapter 1

- **Self-assessment**
  - Why is the quarter-on-quarter growth of unadjusted GDP misleading?
  - What does a seasonal factor of 0.76 for Q1 mean?
  - What does an estimated α close to 1 say about a series?
  - Why must the test set come after the training set?
- **Project idea**
  - forecast monthly electricity generation in Romania one year ahead; compare the seasonal naive method, Holt–Winters, and STL followed by SES, with a rolling origin
- **Chapter 1: stochastic processes and stationarity**
  - What is a stochastic process?
  - When does a series have a stable mean and variance?
  - What are white noise and a random walk?
  - How do we test that autocorrelations are zero?

References (1/3)

- Box, G.E.P., Jenkins, G.M. (1970). *Time Series Analysis: Forecasting and Control*. Holden-Day; 4th ed. with G.C. Reinsel: Wiley (2008).
- Brockwell, P.J., Davis, R.A. (2016). *Introduction to Time Series and Forecasting* (3rd ed.). Springer.
- Cleveland, R.B., Cleveland, W.S., McRae, J.E., Terpenning, I. (1990). STL: a seasonal-trend decomposition procedure based on loess. *Journal of Official Statistics*, 6(1), 3–73.
- Engle, R.F. (1982). Autoregressive conditional heteroscedasticity with estimates of the variance of United Kingdom inflation. *Econometrica*, 50(4), 987–1007.
- Engle, R.F., Granger, C.W.J. (1987). Co-integration and error correction: representation, estimation, and testing. *Econometrica*, 55(2), 251–276.
- Gardner, E.S. (1985). Exponential smoothing: the state of the art. *Journal of Forecasting*, 4(1), 1–28.
- Hamilton, J.D. (1994). *Time Series Analysis*. Princeton University Press.
- Harris, C.R., Millman, K.J., van der Walt, S.J., et al. (2020). Array programming with NumPy. *Nature*, 585, 357–362.

References (2/3)

- Holt, C.C. (1957; reprinted 2004). Forecasting seasonals and trends by exponentially weighted moving averages. *International Journal of Forecasting*, 20(1), 5–10.
- Huang, C., Petukhina, A. (2022). *Applied Time Series Analysis and Forecasting with Python*. Springer.
- Hyndman, R.J., Athanasopoulos, G. (2021). *Forecasting: Principles and Practice* (3rd ed.). OTexts.
- Hyndman, R.J., Koehler, A.B. (2006). Another look at measures of forecast accuracy. *International Journal of Forecasting*, 22(4), 679–688.
- Hyndman, R.J., Koehler, A.B., Snyder, R.D., Grose, S. (2002). A state space framework for automatic forecasting using exponential smoothing methods. *International Journal of Forecasting*, 18(3), 439–454.
- Keeling, C.D. (1960). The concentration and isotopic abundances of carbon dioxide in the atmosphere. *Tellus*, 12(2), 200–203.
- Kolmogorov, A.N. (1941). Interpolation and extrapolation of stationary random sequences. English translation in *Selected Works of A.N. Kolmogorov*, Vol. II, Springer (1992), 272–280.
- Makridakis, S., Hibon, M. (2000). The M3-Competition: results, conclusions and implications. *International Journal of Forecasting*, 16(4), 451–476.
- Makridakis, S., Spiliotis, E., Assimakopoulos, V. (2020). The M4 Competition: 100,000 time series and 61 forecasting methods. *International Journal of Forecasting*, 36(1), 54–74.

References (3/3)

- McKinney, W. (2010). Data structures for statistical computing in Python. *Proceedings of the 9th Python in Science Conference*, 56–61.
- Seabold, S., Perktold, J. (2010). Statsmodels: econometric and statistical modeling with Python. *Proceedings of the 9th Python in Science Conference*, 92–96.
- Sims, C.A. (1980). Macroeconomics and reality. *Econometrica*, 48(1), 1–48.
- Slutzky, E. (1937). The summation of random causes as the source of cyclic processes. *Econometrica*, 5(2), 105–146.
- Wiener, N. (1949). *Extrapolation, Interpolation, and Smoothing of Stationary Time Series*. MIT Press.
- Winters, P.R. (1960). Forecasting sales by exponentially weighted moving averages. *Management Science*, 6(3), 324–342.
- Yule, G.U. (1926). Why do we sometimes get nonsense-correlations between time-series? *Journal of the Royal Statistical Society*, 89(1), 1–63.
- Yule, G.U. (1927). On a method of investigating periodicities in disturbed series, with special reference to Wolfer's sunspot numbers. *Philosophical Transactions of the Royal Society A*, 226, 267–298.