








Statistics of Financial Markets

Chapter 16: Review



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
- **AD**: Anderson–Darling (test)
- **AI**: Artificial Intelligence
- **AIC**: Akaike Information Criterion
- **ARCH**: AutoRegressive Conditional Heteroskedasticity
- **ARCH-LM**: ARCH Lagrange Multiplier (test)
- **AUC**: Area Under the (ROC) Curve
- **BET**: Bucharest Exchange Trading (index)
- **BIC**: Bayesian Information Criterion
- **CAGR**: Compound Annual Growth Rate
- **CDF**: Cumulative Distribution Function
- **CI**: Confidence Interval
- **DFA**: Detrended Fluctuation Analysis
- **EAD**: Exposure At Default
- **EGARCH**: Exponential GARCH
- **EODHD**: EOD Historical Data (market-data provider)
- **ES**: Expected Shortfall
- **EVT**: Extreme Value Theory
- **EWMA**: Exponentially Weighted Moving Average
- **FHS**: Filtered Historical Simulation
- **FPR**: False Positive Rate
- **GARCH**: Generalised AutoRegressive Conditional Heteroskedasticity
- **GBM**: Geometric Brownian Motion

Acronyms used in this material (2/2)

- **GEV**: Generalised Extreme Value (distribution)
- **GJR**: Glosten–Jagannathan–Runkle (asymmetric GARCH)
- **GPD**: Generalised Pareto Distribution
- **GPH**: Geweke–Porter-Hudak (log-periodogram estimator)
- **HAR**: Heterogeneous AutoRegressive (model)
- **HS**: Historical Simulation
- **IGARCH**: Integrated GARCH
- **JB**: Jarque–Bera (test)
- **KS**: Kolmogorov–Smirnov (test)
- **LDA**: Linear Discriminant Analysis
- **LGD**: Loss Given Default
- **LR**: Likelihood Ratio
- **MDD**: Maximum DrawDown
- **MES**: Marginal Expected Shortfall
- **NIG**: Normal Inverse Gaussian (distribution)
- **PD**: Probability of Default
- **QLIKE**: Quasi-LIKElihood (loss function)
- **QQ**: Quantile–Quantile (plot)
- **SE**: Standard Error
- **SR**: Sharpe Ratio
- **SRISK**: Systemic RISK (capital shortfall in a crisis)
- **TPR**: True Positive Rate
- **US**: United States
- **VR**: Variance Ratio

Today's Question and Route

- **Question**: what does the whole course say about the returns of the BET, the S&P 500 and Bitcoin, and how is it examined?
  - one chapter that ties Chapters 0–15 together
- **Route**
  - Part I: the course map, two recap slides per chapter, an overview of tests and models
  - Part II: the exam (format, grading, eight solved problems) and the team project
  - Part III: how AI could help in your own research



The Stock Exchange Palace, Bucharest

Photo: Chainwit. (2023); CC BY 4.0; Wikimedia Commons

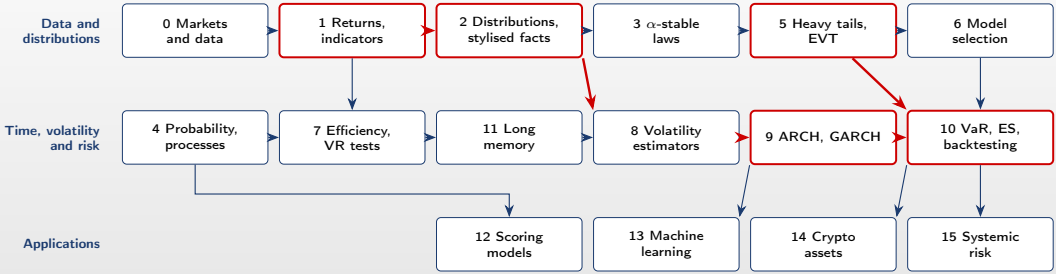
Learning Outcomes

- Place every chapter on the path from prices to risk numbers and name the question it answers
- Recall the key formula, the key empirical fact and the most common mistake of each chapter
- Choose the right test or model for a given question about a return series
- Solve exam-type problems step by step and interpret the result in three to five sentences
- Plan the team project, its oral defence and the declaration of AI use

Reading and Tools

- Textbook: Franke, Härdle & Hafner (2019); exercises with solutions: Borak, Härdle & López-Cabrera (2013)
  - time series and volatility: Tsay (2010); market efficiency: Campbell, Lo & MacKinlay (1997); risk measures: McNeil, Frey & Embrechts (2015)
- Python Quantlets of this chapter: Quantlets/Ch_16
  - the course in one table: S&P 500, BET and Bitcoin on one common window
- Lecture notebook, the course in one notebook: open in Google Colab
- Video course: Statistics of Financial Markets
- Self-assessment: the quiz of this chapter, 20 questions from all chapters

The Course Map: Chapters 0–15

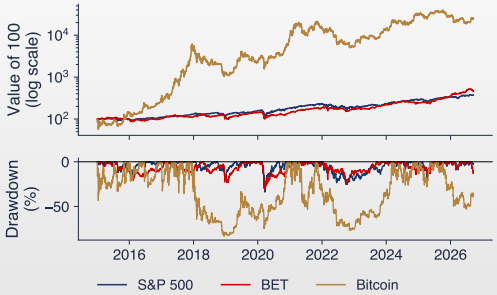


- Red: the main thread, from returns and their distribution to a risk number we can test (Chapter 10)
- Row 2: dependence in time; the sign of returns is hard to predict (Chapters 7, 11), their size is predictable (Chapters 8–9)
- Row 3: the same tools applied to credit, prediction, crypto assets and the banking system

The Structure of the Course

- **Data** (Chapters 0–1): prices from one source, adjusted, on the right calendar
  - returns: $r_t = \ln(P_t/P_{t-1})$, in % for one day
- **Distribution** (Chapters 2–6): heavy tails, not the Normal distribution
  - Student-t, stable laws, EVT; the model is chosen with likelihood, AIC and tail tests
- **Dependence in time** (Chapters 7–11): the sign is almost unpredictable, the size is predictable
  - variance ratio and Hurst for the mean; EWMA and GARCH for the variance
- **Risk** (Chapter 10): VaR 1% and ES 2.5%, with a backtest
  - the applications (Chapters 12–15) reuse the same chain: data, distribution, dependence, risk, evaluation

S&P 500, BET and Bitcoin over a Common Period



- 100 invested on 2 January 2015 (log scale) and the drawdown $P_t / \max_{s \le t} P_s - 1$
- Maximum drawdown: S&P 500 −33.9%, BET −31.1% (both in March 2020), Bitcoin −83.4%

SFM_ch16_course_summary

Summary Table of the Course

Daily log returns	S&P 500	BET	Bitcoin
Observations (per year)	2 943 (252)	2 925 (250)	4 278 (365)
Mean log return, % a year	11.2	13.2	47.4
Volatility, % a year	17.7	15.5	66.7
Sharpe ratio ($r_f = 0$) $\pm$ SE	0.72 ± 0.33	0.93 ± 0.35	1.04 ± 0.36
Skewness; excess kurtosis	-0.64; 15.8	-1.41; 17.8	-0.73; 12.2
Hill tail index of losses ($k = 2.5\% \, n$)	2.75	2.19	2.67
Ljung–Box $Q(10)$: r_t; r_t^2	236; 3166	55; 701	21; 211
VR(5); robust $Z^*(5)$	0.84; -1.34	1.17; 1.93	0.99; -0.22
GARCH(1,1)-t: $\alpha + \beta$	0.993	0.925	1.000
VaR 1%: historical; Normal (%)	3.29; 2.55	2.72; 2.23	10.60; 7.99
ES 2.5%, historical (%)	3.53	3.21	10.93
Hurst (R/S): r_t; $\lvert r_t \rvert$	0.53; 0.84	0.59; 0.75	0.59; 0.78

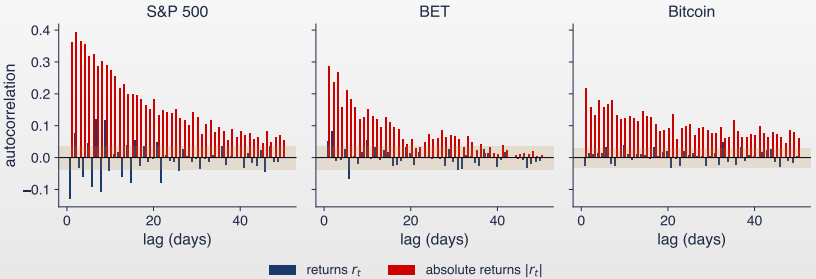
- Window: first return 5 January 2015 (S&P 500), 6 January 2015 (BET), 2 January 2015 (Bitcoin) – 18 September 2026; each series on its own calendar (Bitcoin: 365 days a year); critical value of $\chi^2(10)$ at 5%: 18.31

SFM_ch16_course_summary

Interpreting the Table

- **Return and risk**: Bitcoin has about four times the volatility of the two indices
  - the Sharpe ratios are close once their standard errors are taken into account
- **Distribution**: negative skewness, excess kurtosis above 10, Hill indices between 2 and 3
  - the Normal VaR 1% is too low by 22% (S&P 500), 18% (BET) and 25% (Bitcoin)
- **Dependence**: Ljung–Box on r_t^2 is many times larger than on r_t
  - the robust VR test rejects no random walk at 5% here; H of $|r_t|$ is far above 0.5
  - volatility is persistent: $\alpha + \beta$ close to 1 (Bitcoin: an IGARCH, $\alpha + \beta = 1$)

Stylised Facts: ACF of Returns and Absolute Returns



- Returns: autocorrelations small and mostly inside the band $\pm 1.96/\sqrt{T}$; absolute returns: positive at 50, 49 and 50 of the 50 lags (BET lag 43: -0.001)
- Uncorrelated is not independent: $\hat{\rho}_1(|r|) = 0.36$ (S&P 500), 0.29 (BET), 0.22 (Bitcoin)

SFM_ch16_course_summary

Chapter 0: Introduction (1/2)

- **Key formulas**
  - annualised volatility $\sigma_{\text{ann}} = \sqrt{A}\,\sigma$, A = observations per year (252 for shares, 365 for crypto)
  - drawdown $D_t = P_t / \max_{s \le t} P_s - 1$; maximum drawdown MDD $= \min_t D_t$
- **Notation**
  - σ: the standard deviation of daily returns; A: the number of observations per year
  - P_t: the price on day t; $\max_{s \le t} P_s$: the highest price up to day t; $D_t \le 0$

Chapter 0: Introduction (2/2)

- **Empirical facts**
  - since 2015: volatility 17.7% a year for the S&P 500, 66.7% for Bitcoin
  - BET: −82.5% from 24 July 2007 to 25 February 2009; back to the peak only on 16 March 2021
- **Common mistakes**
  - prices from several sources mixed, or unadjusted prices for shares with dividends and splits
  - a referenced number or paper produced by AI and never checked

Chapter 1: Data, Returns and Indicators (1/2)

- **Key formulas**
  - $R_t = P_t/P_{t-1} - 1$, $r_t = \ln(1 + R_t)$; log returns add over time, simple returns add across assets
  - CAGR $= (P_T/P_0)^{1/Y} - 1$; Sharpe $= (\mu - r_f)/\sigma$, SE $\approx \sqrt{(1 + \mathrm{SR}^2/2)/Y}$; volatility drag $\approx \sigma^2/2$
- **Notation**
  - R_t: the simple return, r_t: the log return of period t; P_0, P_T: the first and the last price; Y: the number of years
  - μ, σ: the annual mean and volatility; r_f: the risk-free rate; SR: the Sharpe ratio; SE: its standard error

Chapter 1: Data, Returns and Indicators (2/2)

- **Empirical facts**
  - BET since 2015: CAGR 14.1%, volatility 15.5%, Sharpe 0.93 with 95% CI [0.24, 1.61]
  - Bitcoin: arithmetic mean 69.5% a year, CAGR only 60.6%: the volatility drag
- **Common mistakes**
  - adding simple returns over time; averaging returns instead of computing the CAGR
  - annualising with 252 for Bitcoin; reporting a Sharpe ratio without its standard error

Chapter 2: Classical Distributions and Stylised Facts (1/2)

- **Key formulas**
  - $S = m_3/m_2^{3/2}$, $K = m_4/m_2^2 - 3$; JB $= \frac{n}{6}(S^2 + K^2/4) \sim \chi^2(2)$ Jarque & Bera (1980)
  - Student-$t(\nu)$: excess kurtosis $6/(\nu - 4)$ for $\nu > 4$; Ljung–Box $Q(m) = n(n+2)\sum_h \hat{\rho}(h)^2/(n-h)$
- **Notation**
  - $m_k = \frac{1}{n}\sum_t (r_t - \bar{r})^k$: the k-th central moment; S: skewness; K: excess kurtosis (0 for the Normal distribution)
  - JB: Jarque–Bera, $\chi^2(2)$ under normality; ν: degrees of freedom; $\hat{\rho}(h)$: the autocorrelation at lag h; m: the number of lags

Chapter 2: Classical Distributions and Stylised Facts (2/2)

- **Empirical facts**
  - BET since 2010: $S = -0.98$, $K = 17.9$, JB $= 56\,762$; fitted Student-t: $\hat{\nu} = 2.96$
  - $\hat{\rho}_1(r) = 0.03$ but $\hat{\rho}_1(|r|) = 0.34$: the six stylised facts of Cont (2001)
- **Common mistakes**
  - reading the kurtosis $K + 3$ as the excess kurtosis (the Normal distribution has kurtosis 3, excess 0)
  - concluding independence from a non-significant ACF of r_t

Chapter 3: α-Stable Distributions (1/2)

- **Key formulas**
  - $X_1 + \cdots + X_n \stackrel{d}{=} n^{1/\alpha} X + d_n$; tails $P(X > x) \sim c\, x^{-\alpha}$; $E|X|^p < \infty$ only for $p < \alpha$
  - four parameters: α (tails), β (skewness), γ (scale), δ (location), in S0 or S1 Nolan (2020)
- **Notation**
  - $X_1, \ldots, X_n$: independent copies of X; $\stackrel{d}{=}$: same distribution; d_n: a shift; $\alpha \in (0, 2]$: the stability index
  - c: a constant; p: the order of a moment; S0, S1: two parametrisations of the same family

Chapter 3: α-Stable Distributions (2/2)

- **Empirical facts**
  - maximum likelihood since 2000: $\hat{\alpha} = 1.49$ (BET), 1.54 (S&P 500), 1.41 (Bitcoin), as Mandelbrot (1963) found for cotton
  - but $\hat{\alpha}$ rises with the horizon (S&P 500: 1.82 on monthly data) and the sample variance settles down: heavy tails with finite variance; the BET is the exception ($\hat{\alpha} = 1.50$ on monthly data, no clear rise)
- **Common mistakes**
  - using the standard deviation as a scale when $\alpha < 2$ (the variance does not exist)
  - comparing δ from two programs that use different parametrisations

Chapter 4: Probability (1/2)

- **Key formulas**
  - $\mathrm{Var}(aX + bY) = a^2\sigma_X^2 + b^2\sigma_Y^2 + 2ab\,\mathrm{Cov}(X, Y)$; Monte Carlo error $\sqrt{p(1-p)/N}$
  - GBM: $S_t = S_0 \exp((\mu - \sigma^2/2)t + \sigma W_t)$, mean $S_0 e^{\mu t}$, median $S_0 e^{(\mu - \sigma^2/2)t}$
- **Notation**
  - a, b: weights; σ_X, σ_Y: standard deviations; p: the estimated probability; N: the number of Monte Carlo draws
  - GBM (geometric Brownian motion): S_t: price at time t (years); μ: drift; σ: volatility; $W_t \sim N(0, t)$: a Brownian motion

Chapter 4: Probability (2/2)

- **Empirical facts**
  - S&P 500 since 2000: a day below -5% has frequency 0.34%; the Normal distribution gives 0.0017%, 199 times less
  - mean return 6.2% a year with SE 3.7%: about 86 years of data would be needed for $t = 3$
- **Common mistakes**
  - uncorrelated taken for independent; the mean of S_t taken for its median
  - a Monte Carlo estimate without its standard error

Chapter 5: Heavy Tails and Extreme Value Theory (1/2)

- **Key formulas**
  - Hill: $\hat{\alpha}_k = [\frac{1}{k}\sum_{i=1}^{k} \ln(L_{(i)}/L_{(k+1)})]^{-1}$, SE $\approx \hat{\alpha}/\sqrt{k}$ Hill (1975); $\xi = 1/\alpha$
  - GPD above u: VaR $= u + \frac{\beta}{\xi}[(n\alpha/N_u)^{-\xi} - 1]$; GEV for block maxima
- **Notation**
  - $L_{(i)}$: the i-th largest loss; k: the number of tail observations; $\xi = 1/\alpha$: the shape (tail) parameter
  - GPD: u: the threshold; β: the scale; N_u: the losses above u out of n; in the VaR formula α is the level (1%), not the tail index

Chapter 5: Heavy Tails and Extreme Value Theory (2/2)

- **Empirical facts**
  - tail index of the losses: BET 2.57 (95% CI [2.20, 2.95], since 1997), S&P 500 2.97, Bitcoin 2.71
  - BET, VaR 0.1%: EVT 9.6%, Normal only 4.5%
- **Common mistakes**
  - choosing k without looking at the Hill plot; reporting $\hat{\alpha}$ without an interval
  - applying the Hill estimator to returns instead of losses $L = -r$

Chapter 6: Model Selection and Risk Management (1/2)

- **Key formulas**
  - $LR = 2(\ell_1 - \ell_0) \approx \chi^2(k_1 - k_0)$ for nested models; AIC $= -2\ell + 2k$, BIC $= -2\ell + k \ln n$
  - KS: $D_n = \sup_x |F_n(x) - F(x)|$; Anderson–Darling weights the tails
- **Notation**
  - ℓ_0, ℓ_1: the maximised log-likelihoods of the smaller and of the larger model; k_0, k_1, k: numbers of parameters; n: sample size
  - F_n: the empirical distribution function; F: the model; D_n: their largest vertical distance; smaller AIC or BIC is better

Chapter 6: Model Selection and Risk Management (2/2)

- **Empirical facts**
  - BET since 2000, Anderson–Darling: Normal 161, Student-t 0.54, NIG 0.40
  - VaR 1% of the BET: Normal 3.19%, heavy-tailed models 3.74–4.94%
- **Common mistakes**
  - the LR test at the boundary (e.g. $\nu = \infty$) without the corrected critical value
  - rejecting every model with a large n and stopping there, instead of ranking the models

Chapter 7: Efficient Markets, Random Walk and VR Tests (1/2)

- **Key formulas**
  - weak-form efficiency: past prices do not predict excess returns Fama (1970); null: martingale differences
  - $\mathrm{VR}(q) = 1 + 2\sum_{k=1}^{q-1}(1 - k/q)\rho(k)$; robust $Z^*(q) = (\widehat{\mathrm{VR}} - 1)/\sqrt{\hat{\theta}/T}$ Lo & MacKinlay (1988)
- **Notation**
  - $\rho(k)$: the autocorrelation at lag k; q: the horizon of the variance ratio; $\mathrm{VR} = 1$ under a random walk
  - T: the number of returns; $\hat{\theta}$: a variance estimate robust to volatility clustering; $Z^* \approx N(0,1)$ under the null

Chapter 7: Efficient Markets, Random Walk and VR Tests (2/2)

- **Empirical facts**
  - BET since 1997: VR(5) $= 1.28$, $Z^* = 5.25$ (momentum)
  - S&P 500 since 1990: VR(5) $= 0.86$, $Z^* = -2.76$ (mean reversion); Bitcoin: $Z^* = -0.19$
- **Common mistakes**
  - the i.i.d. statistic $Z(q)$ on returns with volatility clustering (it rejects too often)
  - a unit root in prices read as proof of efficiency; predictability read as profit

Chapter 8: Volatility Estimators (1/2)

- **Key formulas**
  - EWMA: $\sigma_t^2 = \lambda\sigma_{t-1}^2 + (1-\lambda)r_{t-1}^2$, half-life $\ln 0.5 / \ln \lambda$ J.P. Morgan/Reuters (1996)
  - Parkinson: $(\ln H_t - \ln L_t)^2/(4 \ln 2)$; ARCH-LM $= nR^2 \sim \chi^2(q)$
- **Notation**
  - $\lambda \in (0,1)$: the decay factor; r_{t-1}: yesterday's return; H_t, L_t: the high and the low of day t
  - ARCH-LM: R^2 of the regression of r_t^2 on its q lags, times n; a large value: volatility clustering

Chapter 8: Volatility Estimators (2/2)

- **Empirical facts**
  - S&P 500: Ljung–Box $Q(10)$ is 91 on r_t and 8015 on r_t^2 (critical value 18.31)
  - EWMA with $\lambda = 0.94$: half-life of 11.2 days; no ghost effect
- **Common mistakes**
  - a rolling window of fixed length: one extreme day enters and leaves abruptly (ghost effect)
  - range estimators on prices with overnight gaps, without the Yang–Zhang correction

Chapter 9: ARCH and GARCH Models (1/2)

- **Key formulas**
  - GARCH(1,1): $\sigma_t^2 = \omega + \alpha\varepsilon_{t-1}^2 + \beta\sigma_{t-1}^2$, $\bar{\sigma}^2 = \omega/(1-\alpha-\beta)$ Engle (1982), Bollerslev (1986)
  - forecast $E_t\sigma_{t+h}^2 = \bar{\sigma}^2 + (\alpha+\beta)^{h-1}(\sigma_{t+1}^2 - \bar{\sigma}^2)$; half-life $\ln 0.5 / \ln(\alpha+\beta)$
- **Notation**
  - ω: the constant; α: the reaction to yesterday's shock ε_{t-1}; β: the memory; $\bar{\sigma}^2$: the long-run variance
  - E_t: the forecast made at day t; h: the horizon in days; $\alpha+\beta$: the speed at which the forecast returns to $\bar{\sigma}^2$

Chapter 9: ARCH and GARCH Models (2/2)

- **Empirical facts**
  - GARCH(1,1)-t since 2000: $\alpha + \beta = 0.995$ (S&P 500, half-life 131 days), 0.991 (BET, 74 days)
  - Bitcoin: $\alpha + \beta = 1.000$, an IGARCH; equity indices need asymmetry (GJR, EGARCH)
- **Common mistakes**
  - Normal innovations for returns with heavy tails; non-robust standard errors
  - a long-run variance $\omega/(1 - \alpha - \beta)$ reported when $\alpha + \beta \geq 1$

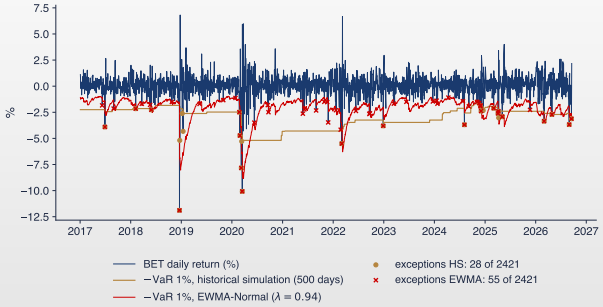
Chapter 10: VaR, ES and Backtesting (1/2)

- **Key formulas**
  - $\mathrm{VaR}_\alpha = -q_\alpha(X)$, $\mathrm{ES}_\alpha = -\frac{1}{\alpha}\int_0^\alpha q_u\,du$; Normal: $-(\mu+\sigma z_\alpha)$, $-\mu+\sigma\varphi(z_\alpha)/\alpha$
  - Kupiec: $LR_{uc} = -2\ln[(1-\alpha)^{n-x}\alpha^x/((1-\hat{\pi})^{n-x}\hat{\pi}^x)] \sim \chi^2(1)$ Kupiec (1995); ES is coherent Artzner et al. (1999)
- **Notation**
  - q_α: the α-quantile of the return X; z_α: the standard Normal quantile; φ: the standard Normal density
  - x: the exceptions in n days; $\hat{\pi} = x/n$: the observed rate; LR_{uc} compares it with the target α

Chapter 10: VaR, ES and Backtesting (2/2)

- **Empirical facts**
  - S&P 500, VaR 1%: historical 3.46%, Normal 2.80%, i.e. 19% too low
  - exception rates since 2005: Normal 2.71%, HS 1.41%, GARCH-t 1.74%, FHS 1.58% (target 1%)
- **Common mistakes**
  - a negative VaR or "VaR 99%": the course writes VaR 1%, a positive loss
  - $\sqrt{10}$ scaling as if returns were i.i.d. Normal; checking only the number of exceptions, not their clustering

VaR 1% Forecasts of the BET: Historical Simulation and EWMA-Normal



- Since 3 January 2017, 2 421 days, 24.2 expected exceptions: HS 28 (1.16%, Kupiec $p = 0.450$), EWMA-Normal 55 (2.27%, $p < 0.001$)
- EWMA follows the storms, but the Normal quantile is too small; HS has the right rate, but its exceptions come in clusters (at most 10 in 250 days)

SFM_ch16_course_summary

Chapter 11: Fractal Markets Hypothesis and Long Memory (1/2)

- **Key formulas**
  - $(R/S)_n \propto n^H$ Hurst (1951); $H = 0.5$: no memory, $H > 0.5$: persistence; $d = H - 0.5$
  - risk scaling: $\mathrm{sd}(r^{(h)}) = h^H \mathrm{sd}(r)$; fractal markets need many investment horizons Peters (1994)
- **Notation**
  - $(R/S)_n$: the rescaled range (range of the cumulated deviations divided by the standard deviation) on blocks of n days
  - H: the Hurst exponent; d: the fractional parameter; $r^{(h)}$: the h-day return; sd: the standard deviation

Chapter 11: Fractal Markets Hypothesis and Long Memory (2/2)

- **Empirical facts**
  - BET: H (R/S) $= 0.58$ for r_t, at the edge of the Monte Carlo band, and the Lo test rejects: weak persistence; 0.79 for $|r_t|$
  - S&P 500: 0.53 for returns, 0.86 for absolute returns: long memory is in the volatility
- **Common mistakes**
  - comparing $\hat{H}$ with 0.5 instead of a Monte Carlo band for the same N (R/S is biased upwards)
  - structural breaks or GARCH taken for long memory

Chapter 12: Scoring Models (1/2)

- **Key formulas**
  - EL $=$ PD $\times$ LGD $\times$ EAD; logit $\ln \frac{p}{1-p} = \beta_0 + x^\top \beta$, e^{β_j} = odds ratio
  - AUC $= P(s_{\text{bad}} > s_{\text{good}})$, Gini $= 2\,\text{AUC} - 1$; Altman Z and Merton distance to default Altman (1968)
- **Notation**
  - EL: expected loss; PD: probability of default; LGD: loss given default; EAD: exposure at default
  - p: the probability of default given x; β_j: a coefficient; s: the score of a bad or a good borrower; AUC $= 0.5$: no discrimination, 1: perfect

Chapter 12: Scoring Models (2/2)

- **Empirical facts**
  - German credit data (1000 loans), test sample of 300: AUC $= 0.801$, 95% CI [0.75, 0.85], Gini 0.60
  - discrimination is not calibration: check both, out of sample
- **Common mistakes**
  - accuracy on imbalanced classes (always "good" is already right in most cases)
  - choosing the cut-off without the costs of the two errors

Chapter 13: Machine Learning (1/2)

- **Key formulas**
  - bias–variance: $E(y_0 - \hat{f})^2 = \text{bias}^2 + \text{Var}(\hat{f}) + \sigma^2$; ridge, lasso, trees, boosting
  - out of sample: $R^2_{OOS} = 1 - \sum(y - \hat{y})^2 / \sum(y - \tilde{y})^2$ against a simple benchmark $\tilde{y}$ Gu, Kelly & Xiu (2020)
- **Notation**
  - y_0: a new observation; $\hat{f}$: the fitted model; σ^2: the noise variance; $\hat{y}$: the forecast
  - $\tilde{y}$: the benchmark forecast; sums over the test sample; $R^2_{OOS} < 0$: worse than the benchmark

Chapter 13: Machine Learning (2/2)

- **Empirical facts**
  - sign of S&P 500 returns since 2008: "always up" 54.4%, logit 53.8%, boosting 52.1%
  - realised volatility: HAR $R^2_{OOS} = 0.61$ against the mean; lasso lowers the squared error of HAR by only 7.1%
- **Common mistakes**
  - random cross-validation on time series (look-ahead bias and leakage)
  - reporting the best of many trials without deflating the Sharpe ratio

Chapter 14: Crypto Assets (1/2)

- **Key formulas**
  - annualise with 365 days: $\sigma_{\text{ann}} = \sqrt{365}\,\sigma$; supply of Bitcoin capped at 21 million Nakamoto (2008)
  - peg deviation of a stablecoin $d_t = 10\,000\,(P_t - 1)$ basis points
- **Notation**
  - σ: the standard deviation of daily returns; crypto trades 365 days a year
  - P_t: the price of the stablecoin in USD; 1 basis point $= 0.01\%$; $d_t < 0$: below the peg

Chapter 14: Crypto Assets (2/2)

- **Empirical facts**
  - Bitcoin: volatility 66.6% a year with 365 days (55.3% if one wrongly uses 252), maximum drawdown -83.4%
  - correlation with the S&P 500: 0.02 before 2020, 0.39 after: not a safe haven
- **Common mistakes**
  - removing weekends from crypto data, or joining crypto and equity returns on different calendars
  - a Normal VaR for a series with $\alpha + \beta = 1$ and tail index below 3

Chapter 15: Systemic Risk (1/2)

- **Key formulas**
  - CoVaR: the VaR of the system given an institution at its VaR; $\Delta\mathrm{CoVaR} = b\,(q_{50} - q_\alpha)$ from a quantile regression Adrian & Brunnermeier (2016)
  - MES, SRISK, Granger networks and the Diebold–Yilmaz connectedness table Diebold & Yilmaz (2012)
- **Notation**
  - b: the slope of the quantile regression of the system on the bank; q_α, q_{50}: the α-quantile and the median of the bank's return
  - MES: the bank's mean loss on the market's worst days; SRISK: the capital missing in a crisis

Chapter 15: Systemic Risk (2/2)

- **Empirical facts**
  - Banca Transilvania and the BET: VaR 1% 5.23%, ΔCoVaR 1% 3.31 pp (95% CI [2.35, 4.78])
  - 11 US and European banks: total connectedness of about 78%
- **Common mistakes**
  - ranking banks by their own VaR: a large VaR need not mean a large contribution to the system
  - correlations compared across calm and crisis periods without the Forbes–Rigobon adjustment

Overview of Methods (1/2): Returns, Distributions and Tails

Question	Tool	Decision	Ch.
How good was the investment?	CAGR, Sharpe, Sortino, MDD	compare with SE and a benchmark	1
Are returns Normal?	JB, QQ plot	reject if JB > 5.99	2
How heavy are the tails?	Hill plot, GPD, stable α	moments exist for $p < \alpha$	3, 5
Which distribution fits?	LR, AIC, BIC, KS, AD	smallest AIC; AD for the tails	6
How rare is a loss of x?	empirical CDF, EVT, Monte Carlo	report the SE of the estimate	4, 5
Is the variance finite?	running variance, aggregation of $\hat{\alpha}$	stable sample variance: finite	3

Overview of Methods (2/2): Dependence in Time, Volatility and Risk

Question	Tool	Decision	Ch.
Are returns predictable?	ACF, Ljung–Box, robust VR $Z^*(q)$	reject if $\lvert Z^*\rvert > 1.96$	7
Is there volatility clustering?	Ljung–Box on r_t^2, ARCH-LM	reject: model the variance	8
How volatile tomorrow?	EWMA, GARCH(1,1)-t, GJR	QLIKE against a benchmark	8, 9
How much can we lose?	VaR 1%, ES 2.5%: HS, t, EVT, FHS	Kupiec, Christoffersen, traffic light	10
Is there long memory?	R/S, DFA, GPH, Lo test	compare with a Monte Carlo band	11
Who will default?	logit, LDA, AUC, Gini	out of sample, with costs	12
Which bank matters for the system?	CoVaR, MES, SRISK	bootstrap CI of ΔCoVaR	15

Conventions Used Throughout the Course

- **Returns**: daily log returns in %, $r_t = 100\,(\ln P_t - \ln P_{t-1})$
  - portfolios: simple returns, prices joined on common days first
- **Annualisation**: mean $\times A$, volatility $\times\sqrt{A}$, with the actual A
  - A: the number of observations per year: about 252 for shares and indices, 365 for crypto
- **Risk**: the level is the tail probability: VaR 1%, ES 2.5%
  - $\mathrm{VaR}_\alpha = -q_\alpha > 0$, a loss; historical simulation with $k = \lceil n\alpha \rceil$
- **Tests**: state H_0, the statistic, its distribution, the critical value and the decision at 5%
  - then one sentence on what the decision means for the data

Common Methodological Errors

- **Data**: aligning returns instead of prices; unadjusted prices; weekends dropped from crypto
- **Distribution**: the Normal distribution for tail quantiles; excess kurtosis confused with kurtosis
- **Inference**: i.i.d. standard errors with volatility clustering; a statistic without its SE
- **Time**: look-ahead bias, parameters estimated on the whole sample and used in the past
- **Risk**: the sign and the level of VaR; ES compared with VaR at different levels
- **Interpretation**: significance taken for economic importance, predictability for profit

The Written Exam: Format

- **70% of the final grade**; written, 2 hours; a calculator is allowed
  - all subjects are compulsory; 9 points plus 1 point ex officio
- **Useful values are given** on the paper
  - quantiles of the Normal and Student-t distributions, critical values of χ^2, logarithms and square roots
- **Each subject**: a short context, numbers or a table or a chart, one or two questions
  - a computation, followed by an interpretation in three to five sentences



Bucharest University of Economic Studies

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

Content Assessed

- **Chapters 0–15**: lectures and seminars
  - the definitions and formulas of the "Prerequisites for Today" slides of every seminar
  - the computations of Part A of the seminars, on paper
- **Reading software output**
  - tables and charts like those of Part B: estimates, test statistics, p-values, backtests
  - no code is written at the exam
- **Understanding, not memory**
  - why a method fails on financial returns, which test answers which question

Grading Criteria

- **The method**: the formula, with the numbers substituted
  - a correct method with an arithmetic slip receives most of the points
- **The result**: the number, with its unit (%, days, RON) and its sign
  - for a test: H_0, the statistic, the critical value or the p-value, the decision
- **The interpretation**: three to five sentences, statistical and economic
  - state the assumptions of an approximate formula (i.i.d., Normal, large n)
  - a correct number with a wrong interpretation does not receive all the points

Typical Problem Types

Type	Chapters	Example in this chapter
returns, annualisation, performance	0, 1	Problem 1: CAGR, drawdown, volatility drag
moments, normality, tails	2, 3, 5, 6	Problem 2: JB, Student-t, Hill
probability and simulation	4	Problem 3: portfolio, Monte Carlo, GBM
efficiency and memory	7, 11	Problems 4 and 7: Ljung–Box, VR, Hurst
volatility	8, 9, 14	Problem 5: GARCH, EWMA, annualisation of Bitcoin
risk measures and backtesting	10	Problem 6: VaR, ES, Kupiec, traffic light
credit and systemic risk	12, 13, 15	Problem 8: confusion matrix, EL, ΔCoVaR

- The eight problems below are review examples with full solutions; the exam has its own subjects

Problem 1: Returns and Drawdown

Context

- An index is worth 100, 160, 64 and 96 at the end of years 0, 1, 2 and 3

Tasks

1. Compute the simple returns R_t and the log returns r_t of the three years.
2. Compute the total return in two ways: from $\sum_t r_t$ and from the prices.
3. Compute the arithmetic mean of R_t and the CAGR, and explain the difference.
4. Compute the maximum drawdown and the gain needed from the trough to the previous peak.

Problem 1: Solution

1. $R = 60\%, -60\%, 50\%$; $r = \ln 1.6 = 0.470$, $\ln 0.4 = -0.916$, $\ln 1.5 = 0.405$
2. $\sum_t r_t = -0.041 = \ln(96/100)$, so the total return is $e^{\mathbf{-0.041}} - 1 = -4\%$, the same as $96/100 - 1$
3. arithmetic mean $(60 - 60 + 50)/3 = 16.67\%$; CAGR $= (96/100)^{\mathbf{1/3}} - 1 = -1.35\%$
4. MDD $= 64/160 - 1 = -60\%$; to return from 64 to 160 the index needs $160/64 - 1 = 150\%$

Interpretation

- A positive average return and a loss of money: volatility lowers the compound growth (volatility drag)
- Log returns add over time; simple returns do not

Problem 2: Distribution and Tails

Context

- $n = 1000$ daily returns with skewness $S = -0.5$ and excess kurtosis $K = 6$
- The five largest losses, in %: 8.2, 6.5, 5.9, 5.1, 4.6; $\chi^2_{0.95}(2) = 5.99$

Tasks

1. Compute the Jarque–Bera statistic, its two parts, and decide at 5%.
2. Find the degrees of freedom ν of a Student-t distribution with the same excess kurtosis.
3. Compute the Hill estimate $\hat{\alpha}$ with $k = 4$.
4. Say which moments of the losses exist if $\alpha = \hat{\alpha}$, and whether a stable law with $\alpha < 2$ is consistent with it.

Problem 2: Solution

1. skewness part $\frac{1000}{6}(0.25) = 41.67$; kurtosis part $\frac{1000}{24}(36) = 1500.0$; JB $= 1541.7 > 5.99$: normality is rejected
2. $6/(\nu - 4) = 6 \Rightarrow \nu = 5$
3. $\ln(8.2/4.6) = 0.578$, $\ln(6.5/4.6) = 0.346$, $\ln(5.9/4.6) = 0.249$, $\ln(5.1/4.6) = 0.103$; mean 0.3190, so $\hat{\alpha} = 3.14$
4. moments of order $p < 3.14$ exist: mean, variance, skewness; the kurtosis does not; a stable law with $\alpha < 2$ (infinite variance) has heavier tails than these data

Interpretation

- The kurtosis part dominates JB: heavy tails, not asymmetry, reject normality
- $k = 4$ gives a very noisy $\hat{\alpha}$ (SE $\approx \hat{\alpha}/2$); in practice k is chosen from a Hill plot

Problem 3: Portfolio, Monte Carlo and GBM

Context

- Two assets: daily volatilities 1.5% and 2.5%, correlation 0.3, weights 60% and 40%
- A share price follows a GBM with $\mu = 8\%$, $\sigma = 25\%$ a year, $S_0 = 100$

Tasks

1. Compute the daily volatility of the portfolio and compare it with the weighted average of the volatilities.
2. A Monte Carlo study with $N = 10\,000$ draws estimates a probability of 1%; compute its SE and a 95% interval.
3. Compute the mean and the median of S_1, and $P(S_1 < S_0)$.

Problem 3: Solution

1. $\sigma_p^2 = 0.81 + 1.00 + 0.54 = 2.35$, $\sigma_p = 1.53\%$, below the weighted average 1.90%
2. SE $= \sqrt{0.01 \times 0.99/10\,000} = 0.099\%$; interval $[0.80\%, 1.20\%]$
3. mean $100e^{\mathbf{0.08}} = 108.33$; median $100e^{\mathbf{0.08-0.03125}} = 105.00$; $P(S_1 < S_0) = \Phi(-0.0488/0.25) = \Phi(-0.195) = 42.3\%$

Interpretation

- Diversification: with $\rho < 1$ the portfolio is less risky than the average of its parts
- A GBM investor loses money after one year with probability 42.3%, although the expected price rises

Problem 4: Autocorrelation and the Variance Ratio

- **Context**
  - $T = 2500$ daily returns; $\hat{\rho}(1) = 0.06$, $\hat{\rho}(2) = 0.02$, $\hat{\rho}(3) = -0.01$, $\hat{\rho}(4) = 0.01$; $\chi^2_{0.95}(4) = 9.49$
  - The heteroskedasticity-robust standard error of $\widehat{\mathrm{VR}}(5)$ is 0.065
- **Tasks**
  1. Compute the Ljung–Box statistic $Q(4)$ and decide at 5%.
  2. Compute $\mathrm{VR}(5) = 1 + 2\sum_{k=1}^{4}(1 - k/5)\hat{\rho}(k)$.
  3. Compute the i.i.d. statistic $Z(5)$ with $\mathrm{SE} = \sqrt{2(2q-1)(q-1)/(3qT)}$ and the robust $Z^*(5)$, and decide at 5%.

Problem 4: Solution

1. $Q(4) = T(T+2)\sum_{k=1}^{4} \hat{\rho}_k^2/(T-k) = 10.51$, close to $T\sum_k \hat{\rho}_k^2$ with $\sum_k \hat{\rho}_k^2 = 0.0042$; $10.51 > 9.49$: the absence of autocorrelation is rejected
2. $\mathrm{VR}(5) = 1 + 2(0.8 \times 0.06 + 0.6 \times 0.02 - 0.4 \times 0.01 + 0.2 \times 0.01) = 1 + 2 \times 0.058 = 1.116$
3. $\mathrm{SE} = \sqrt{72/37\,500} = 0.0438$, $Z(5) = 2.65 > 1.96$; robust $Z^*(5) = (1.116 - 1)/0.065 = 1.78 < 1.96$

Interpretation

- VR > 1: positive autocorrelation (momentum); the i.i.d. test rejects the random walk
- The robust test does not: volatility clustering inflates $Z(5)$; the robust decision is the one to report

Problem 5: Volatility Forecasts

Context

- GARCH(1,1) for daily returns in %: $\omega = 0.03$, $\alpha = 0.09$, $\beta = 0.89$; today $\sigma_t^2 = 1.8$ and $\varepsilon_t = -2.5$
- Bitcoin: daily standard deviation 3.5%

Tasks

1. Compute σ_{t+1}^2 and the one-day Normal VaR 1% (mean zero, $z_{0.01} = -2.326$).
2. Compute the persistence, the half-life, the long-run variance and the long-run annual volatility.
3. Compute $E_t[\sigma_{t+10}^2]$ and the EWMA forecast with $\lambda = 0.94$.
4. Annualise the volatility of Bitcoin.

Problem 5: Solution

1. $\sigma^2_{t+1} = 0.03 + 0.5625 + 1.602 = 2.1945$, $\sigma_{t+1} = 1.481\%$; VaR $1\% = 2.326 \times 1.481 = 3.45\%$
2. $\alpha + \beta = 0.98$; $h_{1/2} = \ln 0.5 / \ln 0.98 = 34.3$ days; $\bar{\sigma}^2 = 0.03/0.02 = 1.50$; $\sqrt{252 \times 1.50} = 19.44\%$ a year
3. $E_t\sigma^2_{t+10} = 1.50 + 0.98^9(2.1945 - 1.50) = 2.079$; EWMA: $0.94 \times 1.8 + 0.06 \times 6.25 = 2.067$
4. Bitcoin trades every day: $3.5 \times \sqrt{365} = 66.9\%$; with 252 the result, 55.6%, would be wrong

Interpretation

- A large negative shock raises tomorrow's variance; the forecast returns to $\bar{\sigma}^2$ at the speed $\alpha + \beta$
- EWMA has no long-run level: its forecast stays flat for every horizon

Problem 6: VaR, ES and Backtesting

Context

- A position of 1 000 000 RON; daily returns $N(0, 1.6^2)$ in %; $\varphi(1.96) = 0.05845$
- The seven smallest of 250 historical returns, in %: -5.1, -4.4, -3.9, -3.6, -3.2, -3.0, -2.9

Tasks

1. Compute the Normal VaR 1% and ES 2.5% in % and in RON, and the 10-day VaR 1% by the square-root-of-time rule.
2. Compute VaR 1% and ES 2.5% by historical simulation.
3. A model has 6 exceptions of VaR 1% in 250 days: compute the Kupiec statistic, decide at 5%, and give the Basel zone.

Problem 6: Solution

1. VaR 1% $= 2.326 \times 1.6 = 3.72\%$ (37 222 RON); ES 2.5% $= 1.6 \times 0.05845/0.025 = 1.6 \times 2.338 = 3.74\%$ (37 405 RON); 10 days: $\sqrt{10} \times 3.72 = 11.77\%$
2. $k = \lceil 2.5 \rceil = 3$: VaR 1% $= 3.9\%$; $k = \lceil 6.25 \rceil = 7$: ES 2.5% $= 26.1/7 = 3.73\%$
3. $\ell_{\mathbf{0}} = 244 \ln 0.99 + 6 \ln 0.01 = -30.08$, $\ell_{\mathbf{1}} = 244 \ln 0.976 + 6 \ln 0.024 = -28.31$; $LR_{uc} = 3.56 < 3.84$ ($p = 0.059$): not rejected; yellow zone (5–9 exceptions)

Interpretation

- For the Normal distribution ES 2.5% and VaR 1% are almost equal; the historical tail is heavier
- Kupiec does not reject 6 exceptions, but the traffic light already asks for a higher capital multiplier

Problem 7: Long Memory and Risk Scaling

Context

- The average R/S statistic of daily returns is 4.1 for blocks of $n = 20$ days and 60.0 for $n = 2000$
- One-day VaR 1% $= 2.5\%$

Tasks

1. Estimate H from the slope of $\log(R/S)$ on $\log n$, and the fractional parameter d.
2. Scale the VaR to 10 days with h^H and with $\sqrt{h}$, and compare.
3. Say what must be checked before concluding that there is long memory.

Problem 7: Solution

1. $H = \log_{10}(60.0/4.1)/\log_{10}(2000/20) = \log_{10}(14.63)/2 = 1.165/2 = 0.583$; $d = H - 0.5 = 0.083$
2. $10^H = 3.83$: VaR $= 9.56\%$; $\sqrt{10} = 3.16$: VaR $= 7.91\%$
3. a Monte Carlo band of $\hat{H}$ for i.i.d. data of the same length (R/S is biased upwards), the Lo test, split samples and shuffled data (breaks, GARCH)

Interpretation

- With $H > 0.5$ the square-root-of-time rule understates multi-day risk
- On the course data the long memory is in $|r_t|$, not in r_t (Chapter 11)

Problem 8: Credit Scoring and Systemic Risk

Context

- A test sample of 300 loans, 90 bad; at the chosen cut-off: 60 bad and 42 good loans are flagged as bad
- Quantile regression at 1% of the system return on a bank's return: $\hat{a} = -2.1$, $\hat{b} = 0.6$; bank: $q_{0.01} = -5.0\%$, median 0.1%

Tasks

1. Compute the accuracy, the true positive rate (TPR), the false positive rate (FPR) and the precision, and compare the accuracy with "always good".
2. Compute the Gini coefficient for AUC $= 0.80$ and the expected loss for PD $= 4\%$, LGD $= 45\%$, EAD $= 200\,000$ RON.
3. Compute the CoVaR 1% of the system with the bank at its VaR and at its median, and ΔCoVaR.

Problem 8: Solution

1. accuracy $(60 + 168)/300 = 76.0\%$; TPR $= 60/90 = 66.7\%$; FPR $= 42/210 = 20.0\%$; precision $60/102 = 58.8\%$; "always good": 70.0%
2. Gini $= 2 \times 0.80 - 1 = 0.60$; EL $= 0.04 \times 0.45 \times 200\,000 = 3\,600$ RON
3. CoVaR $= -(-2.1 + 0.6 \times (-5.0)) = 5.10\%$; at the median: $-(-2.1 + 0.6 \times 0.1) = 2.04\%$; ΔCoVaR $= 0.6 \times 5.1 = 3.06$ pp

Interpretation

- The model is barely more accurate than "always good", but it finds two thirds of the bad loans: judge it by AUC and costs
- The distress of this bank raises the VaR of the system by 3.06 percentage points

The Team Project: Content and Deliverables

- **20% of the final grade**: a team analysis of real market data with the methods of the course
  - one research question; returns and indicators, distributions and tails, efficiency tests, volatility and risk measures
  - ideas: the "Project Idea" slide at the end of every chapter and Part C of every seminar
- **Deliverables**
  - a GitHub repository whose code reproduces every number and every chart from the saved data
  - a short report and a presentation of the results
  - the file `AI_USE.md`, which declares every use of AI tools

Project Grading Criteria

- **The question**: clear, answerable with the data, relevant for investors or regulators
  - "Is the BET today less volatile than in 2008?" is a question; "an analysis of the BET" is not
- **The methods and the checks**
  - the right test for the question (overview of methods); robust standard errors; out-of-sample evaluation
  - robustness: another window, another series, another estimator
- **The interpretation**: what the numbers mean, and what the data cannot show
  - **Reproducibility**: the repository runs from the saved data to every chart

The Oral Defence

- **Each member explains the code and the results**
  - typical questions: what does this line compute? where does this number come from? why this test and not another?
  - what changes if the window, the series or the level α changes?
- **A line of code that cannot be explained does not count as your own work**
  - this holds for code written with or without AI
- **Preparation**
  - rerun the repository from scratch the day before; know the definitions of every number on your slides

The File AI_USE.md

- **AI tools are allowed and must be declared**
  - for every use: the tool, the prompt, what was kept and what was corrected
  - the errors of the AI that you found and how you found them
- **You are responsible for every number and every reference**
  - every number: recomputed with your own code from the course data
  - every reference: its DOI opened and its title checked
- **Typical AI errors in this course**
  - "VaR 99%" with a negative sign, Bitcoin annualised with 252 days, invented references, a test with the wrong distribution

An Open Research Question on the BET

- **Has the BET become a "normal" market?**
  - since 1997: VR(5) $= 1.28$ (momentum) and tail index 2.57
  - since 2015: VR(5) $= 1.17$, $Z^* = 1.93$, tail index 2.19
- **Why it is open**
  - the momentum fades, but the tails do not become thinner
  - one extreme day (-11.9% on 19 December 2018) moves the tail estimate; few crises, short samples
- Different teams with the same data reach different answers: nonstandard errors Menkveld et al. (2024)



Benoit Mandelbrot, who showed in 1963 that price changes have heavy tails

Photo: Rama (2007); CC BY-SA 2.0 fr; Wikimedia Commons

Steps of a Research Project

1. **Question**: one sentence, with a number that would answer it
2. **Data**: series, source, calendar, window, fixed before looking at the results
3. **Method**: the test chosen from the overview of methods, with its null hypothesis
4. **Checks**: robust SE, Monte Carlo band, subsamples, another estimator
5. **Interpretation**: the answer, its uncertainty, its limits
6. **Replication**: a second person runs the code and gets the same numbers

⊡ AI can speed up every step; it replaces none of the checks Wang et al. (2023)

Possible Uses of AI Tools

- **Literature**: a first list of studies on the efficiency and the tails of emerging markets, each to be checked by its DOI
- **Code**: a first draft of rolling VR(5), Hill and GARCH estimates on windows of 500 days
- **Robustness**: a grid of choices (window, k, level α, series) and one table of all results
- **Writing**: a clearer paragraph, after the numbers are final
- Example prompt
  - `Write Python code that computes, on rolling windows of 500 trading days moved by 21 days, the Lo-MacKinlay robust variance-ratio statistic Z*(5), the Hill tail index of the losses with k = 2.5% of the window, and the GARCH(1,1)-t persistence for the daily log returns of the BET, and plots the three series with their 95% bands.`

What to Check

- The formulas: the robust $Z^*(q)$, not $Z(q)$; the Hill estimator on losses $L = -r$, with k given as a share of the window
- The look-ahead: each window uses only the data up to its last day
- The bands: overlapping windows are not independent evidence; one result in twenty is significant by chance
- The numbers: recompute one window step by step and compare
- The references: every cited paper must exist; open the DOI and check the title

Project Idea

- **Question**: did the BET become more efficient, less volatile and thinner-tailed between 2000 and 2026?
  - data: BET since 1997, with the S&P 500 and the DAX as benchmarks (EODHD)
- Steps
  - rolling robust VR(5), Hill index, GARCH persistence and Hurst exponent (Chapters 5, 7, 9, 11)
  - bands from Monte Carlo or the bootstrap; the same analysis for the benchmarks
  - dates of change (e.g. 2008, 2020) checked against the bands, not chosen after looking
- Deliverable: one chart with four panels, one table, and a paragraph on what the data can and cannot show
- Declare any AI use, and list the errors of the AI that you corrected

Key Takeaways

- Returns, not prices: log returns in %, the right calendar, annualised with the actual frequency
- Returns have heavy tails (tail index 2–4), negative skewness and finite variance: the Normal distribution understates risk
- The sign of returns is almost unpredictable; their size is predictable (volatility clustering, GARCH)
- Risk: VaR 1% $= -q_{1\%}$ and ES 2.5%, always followed by a backtest
- Every result needs its uncertainty: SE, robust tests, Monte Carlo bands, out-of-sample checks
- The exam rewards method, result and interpretation; the project rewards a clear question answered honestly

Self-Assessment

- **Question**: Ljung–Box does not reject on r_t but rejects strongly on r_t^2. What does this mean?
  - **Answer**: returns are uncorrelated but not independent: volatility clustering; model the variance (GARCH)
- **Question**: a fitted stable law gives $\hat{\alpha} = 1.5$ and the Hill estimate is $\hat{\alpha} = 3$. Which is right about the variance?
  - **Answer**: Hill, which looks only at the tail: the variance is finite; the stable fit describes the body
- **Question**: a model has 1% exceptions, all in March 2020. Is it a good VaR model?
  - **Answer**: no: the rate is right, but the exceptions cluster; the Christoffersen test rejects it

References (1/2)

- Adrian, T., Brunnermeier, M.K. (2016). CoVaR. *American Economic Review*, 106(7), 1705–1741.
- Altman, E.I. (1968). Financial ratios, discriminant analysis and the prediction of corporate bankruptcy. *Journal of Finance*, 23(4), 589–609.
- Artzner, P., Delbaen, F., Eber, J.-M., Heath, D. (1999). Coherent measures of risk. *Mathematical Finance*, 9(3), 203–228.
- Bollerslev, T. (1986). Generalized autoregressive conditional heteroskedasticity. *Journal of Econometrics*, 31(3), 307–327.
- Borak, S., Härdle, W.K., López-Cabrera, B. (2013). *Statistics of Financial Markets: Exercises and Solutions*, 2nd ed. Springer.
- Campbell, J.Y., Lo, A.W., MacKinlay, A.C. (1997). *The Econometrics of Financial Markets*. Princeton University Press.
- Cont, R. (2001). Empirical properties of asset returns: stylized facts and statistical issues. *Quantitative Finance*, 1(2), 223–236.
- Diebold, F.X., Yilmaz, K. (2012). Better to give than to receive: predictive directional measurement of volatility spillovers. *International Journal of Forecasting*, 28(1), 57–66.
- Engle, R.F. (1982). Autoregressive conditional heteroscedasticity with estimates of the variance of United Kingdom inflation. *Econometrica*, 50(4), 987–1007.
- Fama, E.F. (1970). Efficient capital markets: a review of theory and empirical work. *Journal of Finance*, 25(2), 383–417.
- Franke, J., Härdle, W.K., Hafner, C.M. (2019). *Statistics of Financial Markets: An Introduction*, 5th ed. Springer.
- Gu, S., Kelly, B., Xiu, D. (2020). Empirical asset pricing via machine learning. *Review of Financial Studies*, 33(5), 2223–2273.
- Hill, B.M. (1975). A simple general approach to inference about the tail of a distribution. *Annals of Statistics*, 3(5), 1163–1174.

References (2/2)

- Hurst, H.E. (1951). Long-term storage capacity of reservoirs. *Transactions of the American Society of Civil Engineers*, 116(1), 770–799.
- J.P. Morgan/Reuters (1996). *RiskMetrics – Technical Document*, 4th ed. New York.
- Jarque, C.M., Bera, A.K. (1980). Efficient tests for normality, homoscedasticity and serial independence of regression residuals. *Economics Letters*, 6(3), 255–259.
- Kupiec, P.H. (1995). Techniques for verifying the accuracy of risk measurement models. *Journal of Derivatives*, 3(2), 73–84.
- Lo, A.W., MacKinlay, A.C. (1988). Stock market prices do not follow random walks: evidence from a simple specification test. *Review of Financial Studies*, 1(1), 41–66.
- Mandelbrot, B. (1963). The variation of certain speculative prices. *Journal of Business*, 36(4), 394–419.
- McNeil, A.J., Frey, R., Embrechts, P. (2015). *Quantitative Risk Management: Concepts, Techniques and Tools*, revised ed. Princeton University Press.
- Menkveld, A.J., Dreber, A., Holzmeister, F., Huber, J., Johannesson, M., et al. (2024). Nonstandard errors. *Journal of Finance*, 79(3), 2339–2390.
- Nakamoto, S. (2008). Bitcoin: a peer-to-peer electronic cash system. White paper.
- Nolan, J.P. (2020). *Univariate Stable Distributions: Models for Heavy Tailed Data*. Springer.
- Peters, E.E. (1994). *Fractal Market Analysis: Applying Chaos Theory to Investment and Economics*. Wiley.
- Tsay, R.S. (2010). *Analysis of Financial Time Series*, 3rd ed. Wiley.
- Wang, H., Fu, T., Du, Y., Gao, W., et al. (2023). Scientific discovery in the age of artificial intelligence. *Nature*, 620, 47–60.