

Modelling Financial Markets

Seminar 7: Value-at-Risk and Expected Shortfall — Hands-on



Daniel Traian PELE
Drd. Antoaneta Amza, antoaneta.amza@csie.ase.ro

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

How Large Can a One-Day Loss Be?

- Question of the seminar
  - **how large are tomorrow's tail losses, and how precisely can we tell?**
- Deliverables
  - the answers to the A exercises
  - one table and one figure per B exercise
  - the corrected C1 code
  - one interpretation sentence per exercise
- Notebook for today: open the seminar notebook in Colab
- Exercises, types and order: the Exercise Map after the definitions
  - B4, A5 and B5 use extreme value theory (EVT) and Euler contributions, defined in the notions below
- Seminar: Prof. Daniel Traian Pele, danpele@ase.ro; Antoaneta Amza, PhD candidate, antoaneta.amza@csie.ase.ro

Which Markets and Return Units Do We Use?

- BET is the Bucharest Exchange Trading index; BVB is Bursa de Valori București. BNR is Banca Națională a României, the source of the EUR/RON reference rate (Romanian lei per euro: RON for one EUR).
- An Exchange-Traded Fund (ETF) trades on an exchange. SPY (SPDR S&P 500 ETF Trust) tracks the S&P 500 (Standard & Poor's); TLT (iShares 20+ Year Treasury Bond ETF) holds long-term US Treasury bonds; GLD (SPDR Gold Shares) provides gold exposure. US means United States; USD denotes US dollars.
- Single-asset losses use log returns in percent. For 100 to 90, the simple return is -10%, but the log return is about -10.54%; the log loss is positive 10.54%. Portfolio returns use simple returns and asset weights.
- C1 critiques an artificial intelligence (AI) answer. A plausible explanation must still respect the sign, units and sample used in the exercise.

Data Used

Series	Price used	Sample	Calendar	Used in
BET	daily close	2000–2026	weekdays	B1, B2, B4, C1
S&P 500	daily close	2000–2026	US trading days	B2
Bitcoin	daily close (USD)	2014–2026	every day	B3
SPY, TLT, GLD	adjusted close	2014–2026	common days with Bitcoin	B5
Gold; EUR/RON	gold: daily close (USD per ounce); EUR/RON: BNR reference rate	2000–2026; 2005–2026	weekdays; Romanian working days	B6
BVB blue chips	Banca Transilvania, OMV Petrom and six other large BVB stocks, adjusted close	2016–2026	BVB trading days	C2

- Daily prices of the course data, one series per asset (field used: close or adjusted close, as in the table); all samples end on 18 September 2026
- EUR/RON: the official BNR reference rate, not the EUR/RON series of the market-data provider (it contains erroneous quotes, see Seminar 0)
- Two or more assets: align the *prices* on common trading days first, then compute returns

Setup: Run and Verify

```
import numpy as np, pandas as pd
URL = ('https://raw.githubusercontent.com/'
       'danpele/MFM/main/data/market/')
px = pd.read_csv(URL + 'BET.csv', index_col='date',
                 parse_dates=True)['close']
px = px.loc['2000-01-01':'2026-09-18']
px = px[px.index.dayofweek < 5]
r = 100 * np.log(px).diff().dropna()   # % log returns
L = -r                                 # losses
print(px.head(3)); print(len(r), r.index[-1])
```

- Code from the notebook section "Data"
- Expected output: $n = 6\,670$ returns, last date 2026-09-18

date	close	r_t (%)
2000-01-05	460.67	–
2000-01-06	484.04	4.95
2000-01-07	526.34	8.38

- Last close: 32 757.81 on 2026-09-18
- If your count or first value differs, check the series, the price column (close or adjusted close) and the weekday filter
- Other series: use the files and price columns on the slide "Which Files Define the Sample?"

VaR and ES

- Details: Chapter 7
- Return X, loss $L = -X$; the level α is a **tail probability** (e.g. 1% or 2.5%)
- **Value-at-Risk** (VaR): the loss exceeded with probability α, i.e. minus the α-quantile of the return

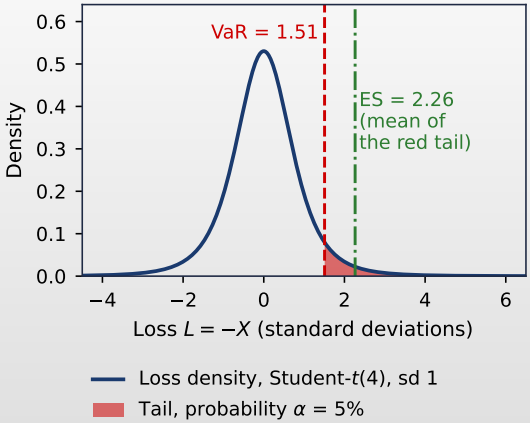
$$\mathrm{VaR}_\alpha = -q_\alpha(X), \qquad q_\alpha(X) = \inf\{x : \ P(X \le x) > \alpha\}$$

  - $q_\alpha(X)$: the α-quantile, the smallest x (infimum) at which the distribution function $P(X \le x)$ exceeds α; VaR says how often a loss is exceeded, not by how much
  - for a continuous loss $L = -X$, $\mathrm{VaR}_\alpha = q_{\mathbf{1}-\alpha}(L)$; for discrete losses (the example below) use the definition through returns
- **Expected Shortfall** (ES): the average loss in the worst α of outcomes

$$\mathrm{ES}_\alpha = -\frac{1}{\alpha}\int_{\mathbf{0}}^{\alpha} q_u(X)\,du \ = \ \mathbb{E}[L \mid L \ge \mathrm{VaR}_\alpha] \ \text{ (continuous } L)$$

  - An average of the VaRs at all levels $u \le \alpha$, so $\mathrm{ES}_\alpha \ge \mathrm{VaR}_\alpha$; $\mathbb{E}[\cdot \mid \cdot]$: the mean over the outcomes that satisfy the condition
  - Example: ten equally likely losses $-2, -1, 0, 0, 1, 1, 2, 3, 5, 9$, $\alpha = 20\%$: $P(L \ge 5) = 20\%$ is not above α, $P(L \ge 3) = 30\%$ is, so $\mathrm{VaR}_{\mathbf{20\%}} = 3$; the worst 20% are $\{9, 5\}$, so $\mathrm{ES}_{\mathbf{20\%}} = 7$
  - Regulation: the Basel market-risk rules replaced VaR 1% by ES 2.5% Basel Committee (2019)

VaR and ES on a Loss Density



- A loss with a Student-$t(4)$ distribution scaled to standard deviation 1; tail probability $\alpha = 5\%$ in red
- VaR: where the red tail starts, $P(L > \mathrm{VaR}_\alpha) = \alpha$
  - Here $\mathrm{VaR}_{\mathbf{5\%}} = 1.51$ standard deviations
- ES: the centre of mass of the red tail, so always to the right of VaR
  - Here $\mathrm{ES}_{\mathbf{5\%}} = 2.26$; two tails with the same VaR can have very different ES
- Signs: VaR and ES are reported as **positive** numbers when they are losses

Normal VaR and ES

- Normal returns: $X = \mu + \sigma Z$, $Z \sim N(0,1)$

$$\mathrm{VaR}_\alpha = -(\mu + \sigma z_\alpha), \qquad \mathrm{ES}_\alpha = -\mu + \sigma \frac{\varphi(z_\alpha)}{\alpha}$$

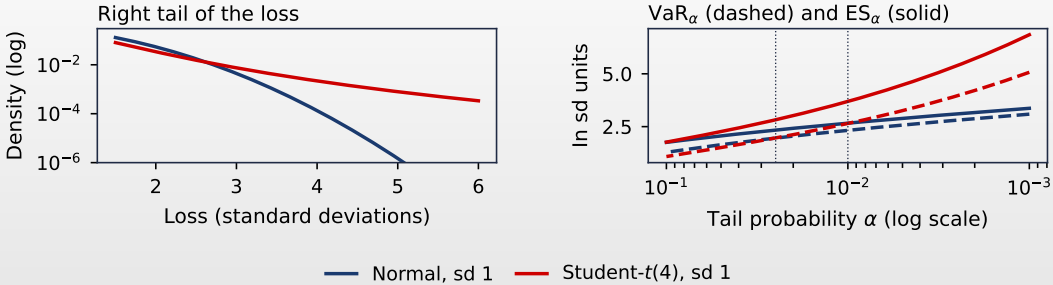
  - μ: mean; σ: standard deviation; Φ, φ: distribution function and density of $N(0,1)$; $z_\alpha = \Phi^{-1}(\alpha) < 0$: its α-quantile
  - VaR lies $|z_\alpha|$ standard deviations below the mean; ES uses the larger factor $\varphi(z_\alpha)/\alpha$, the mean of $-Z$ over the tail (A1 derives it)
  - $z_{1\%} = -2.326$, $\varphi(z_{1\%})/0.01 = 2.665$; $z_{2.5\%} = -1.960$, $\varphi(z_{2.5\%})/0.025 = 2.338$; with $\mu = 0$, $\sigma = 1\%$: VaR 1% $= 2.33\%$, ES 2.5% $= 2.34\%$

- Student-t returns: $X = m + sT$, $T \sim t_\nu$ with density g_ν and quantile $t_\alpha = t_\nu^{-1}(\alpha)$

$$\mathrm{VaR}_\alpha = -(m + s\, t_\alpha), \qquad \mathrm{ES}_\alpha = -m + s \frac{g_\nu(t_\alpha)}{\alpha} \frac{\nu + t_\alpha^2}{\nu - 1} \quad (\nu > 1)$$

  - m: location; $s > 0$: scale; ν: degrees of freedom, smaller ν = heavier tails; $\mathrm{Var}(T) = \nu/(\nu-2)$ for $\nu > 2$, so a standard deviation σ needs $s = \sigma\sqrt{(\nu-2)/\nu}$ (A2)

Normal Against Student-t



- Both distributions have standard deviation 1; left: the right tail of the loss on a log scale, where the Student-$t(4)$ density is many times higher beyond 3 standard deviations
- Right: VaR (dashed) and ES (solid) in standard deviations, for α from 10% to 0.1%; dotted: $\alpha = 2.5\%$ and 1%
  - At moderate levels the two models are close; the gap opens deep in the tail, and faster for ES than for VaR
  - Question for the class: at $\alpha = 5\%$ and the same standard deviation, which model gives the larger VaR?
  - Answer: the Normal distribution (1.64 vs 1.51 standard deviations); at 1% the order reverses (2.33 vs 2.65)

Coherent Risk Measures

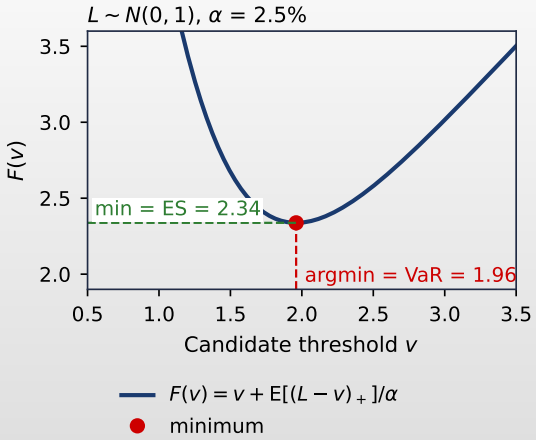
- **Coherent risk measure** ρ Artzner et al. (1999): a rule that maps a loss L to the capital it needs and satisfies four axioms

$$L_1 \leq L_2 \Rightarrow \rho(L_1) \leq \rho(L_2), \qquad \rho(L + c) = \rho(L) + c,$$
$$\rho(\lambda L) = \lambda\, \rho(L)\ (\lambda > 0), \qquad \rho(L_1 + L_2) \leq \rho(L_1) + \rho(L_2)$$

  - Monotone (larger losses, more capital); translation invariant (a sure loss c adds c); positively homogeneous (double the position, double the capital)
  - **Subadditive**: merging two positions L_1, L_2 never needs more capital than holding them apart, so diversification is never penalised

- ES is coherent; VaR satisfies the first three axioms, but can fail subadditivity
  - Typical cases: discrete losses (defaults, A4) and very heavy tails; for jointly Normal losses and a tail probability below 50% VaR is subadditive

ES as a Minimum



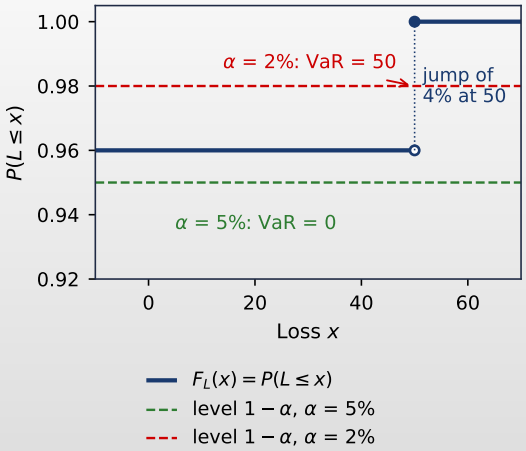
- **Rockafellar–Uryasev** Rockafellar & Uryasev (2000): ES is the minimum of a convex function of a candidate threshold v

$$F_\alpha(v) = v + \frac{\mathbb{E}[(L-v)_+]}{\alpha}, \qquad \mathrm{ES}_\alpha = \min_v F_\alpha(v)$$

  - $(x)_+ = \max(x, 0)$; $\mathbb{E}[(L-v)_+]$: mean loss in excess of v (zero below v)
  - Chart ($L \sim N(0,1)$, $\alpha = 2.5\%$): the minimum is reached at $v = \mathrm{VaR}_\alpha = 1.96$ and equals $\mathrm{ES}_\alpha = 2.34$

- Uses: ES estimation as an optimisation (A7), ES contributions by differentiation (A5 extended)

Losses with an Atom

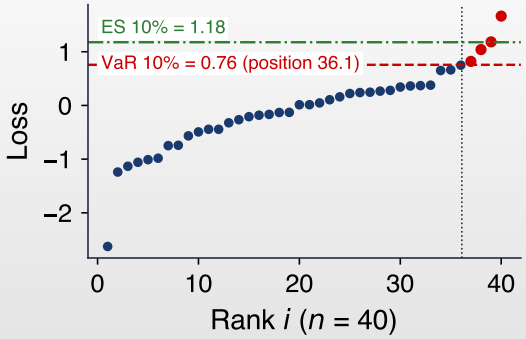


- An atom: a loss value with positive probability, e.g. a default; the distribution function jumps there
- VaR from the infimum definition: the first x where $P(L \le x) \ge 1 - \alpha$
  - Chart (loss 50 with probability 4%): $\alpha = 5\%$ gives VaR $= 0$; $\alpha = 2\%$ gives VaR $= 50$
- ES must fill a tail of probability exactly α Acerbi & Tasche (2002)

$$\mathrm{ES}_\alpha = \frac{1}{\alpha}\Big(\mathbb{E}\big[L\,1\{L > \mathrm{VaR}_\alpha\}\big] + \mathrm{VaR}_\alpha\big(\alpha - P(L > \mathrm{VaR}_\alpha)\big)\Big)$$

  - $1\{\cdot\}$: 1 if the condition holds, 0 otherwise; first term: losses strictly beyond VaR; second: the missing tail probability, placed at VaR
  - Chart, $\alpha = 5\%$: $(50 \times 0.04 + 0 \times 0.01)/0.05 = 40$

Historical Simulation



- sorted losses $L_{(1)} \leq \ldots \leq L_{(n)}$
- the $k = 4$ losses $\geq$ VaR

- Historical simulation (HS): no model, only the past losses sorted, $L_{(1)} \leq \cdots \leq L_{(n)}$

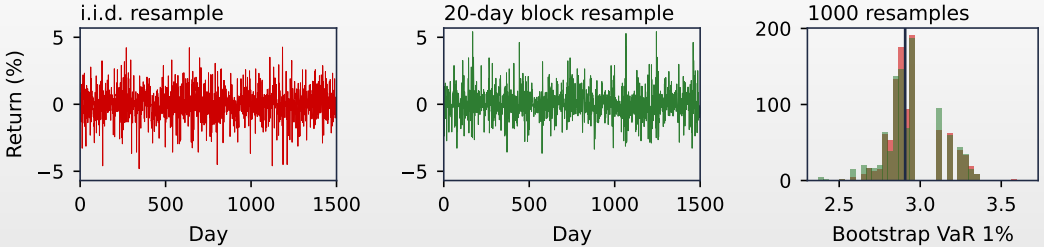
$$\widehat{\mathrm{VaR}}_\alpha = L_{(j)} + (p - j)\big(L_{(j+1)} - L_{(j)}\big),$$
$$p = 1 + (n-1)(1-\alpha), \quad j = \lfloor p \rfloor$$

  - p: position among the sorted losses (`np.quantile`, linear interpolation); $j = \lfloor p \rfloor$: its integer part

- $\widehat{\mathrm{ES}}_\alpha = \frac{1}{n\alpha}\big[\sum_{i=1}^{m} L_{(n-i+1)} + (n\alpha - m)L_{(n-m)}\big]$, $m = \lfloor n\alpha \rfloor$: the m largest losses with weight 1, the next with the fractional weight $n\alpha - m$
  - Chart: $n = 40$, $\alpha = 10\%$, $p = 36.1$, $m = 4$
  - The tail count $n\alpha$ drives precision: $n = 500$, $\alpha = 1\%$ leaves five losses

- HS is unconditional: a calm window gives a low VaR even if today is volatile

Bootstrap Intervals



- **Bootstrap** Efron (1979): draw n losses with replacement, recompute the statistic $\hat{\theta}^{(b)}$, repeat for $b = 1, \ldots, B$
  - 95% percentile Confidence Interval (CI): $[q_{2.5\%},\ q_{97.5\%}]$ of the $\hat{\theta}^{(b)}$; its width measures the sampling uncertainty
- Independent and identically distributed (i.i.d.) draws break volatility clustering (left); the **moving-block bootstrap** Künsch (1989) draws blocks of, e.g., 20 consecutive days, concatenated and cut to length n (middle)
  - Right: with clustered crises, the block distribution of VaR 1% is wider than the i.i.d. one

Precision of VaR

- **Asymptotic Standard Error (SE) of HS VaR** (i.i.d. losses)

$$\mathrm{SE}(\widehat{\mathrm{VaR}}_\alpha) \approx \frac{\sqrt{\alpha(1-\alpha)/n}}{f(\mathrm{VaR}_\alpha)}$$

  - f: density of L at VaR; few observations near VaR (small f) = imprecise VaR; $\sqrt{\alpha(1-\alpha)/n}$: standard error of the share of losses above VaR
  - $\hat{f}$: Gaussian kernel density, bandwidth $h = \hat{\sigma}\,(4/(3n))^{\mathbf{1/5}}$ (SciPy's Silverman rule), $\hat{\sigma}$: the sample standard deviation

- **HAC** (heteroskedasticity- and autocorrelation-consistent) version: replace $\alpha(1-\alpha)$ by the long-run variance of $I_t = 1\{L_t > \widehat{\mathrm{VaR}}\}$

$$\hat{\Omega} = \hat{\gamma}_{\mathbf{0}} + 2\sum_{j=\mathbf{1}}^{m}\Big(1 - \frac{j}{m+1}\Big)\hat{\gamma}_j$$

  - $\hat{\gamma}_j$: sample autocovariance of I_t at lag j; Bartlett weights $1 - j/(m+1)$; Newey–West lag $m = \lfloor 4(n/100)^{\mathbf{2/9}} \rfloor$; clustered exceedances give $\hat{\Omega} > \alpha(1-\alpha)$

Precision of ES

- **ES**: influence function Chen (2008) and its variance

$$\mathrm{IF}(x) = \mathrm{VaR}_\alpha + \frac{(x - \mathrm{VaR}_\alpha)_+}{\alpha} - \mathrm{ES}_\alpha, \qquad \mathrm{Var}(\widehat{\mathrm{ES}}_\alpha) \approx \frac{\mathrm{Var}(\mathrm{IF}(L))}{n}$$

  - $\mathrm{IF}(x)$: the effect of one extra observation x on the estimate; only losses beyond VaR move it, each with weight $1/\alpha$
  - So the precision of $\widehat{\mathrm{ES}}_\alpha$ is driven by the few largest losses: the tail count $n\alpha$ and how heavy the tail is beyond VaR
- Rule of thumb: relative standard error of an i.i.d. ES estimate roughly $\propto 1/\sqrt{n\alpha}$
  - $n = 250$, $\alpha = 2.5\%$: only 6.25 tail losses; a block bootstrap widens the interval further when the tail days cluster

Filtered Historical Simulation

- Filtered Historical Simulation (FHS) Barone-Adesi et al. (1999): fit an AR(1) (autoregressive) mean and a GARCH(1,1) (Generalised AutoRegressive Conditional Heteroskedasticity, Chapter 5) variance by Normal quasi-maximum likelihood (QML), with an AR intercept, then standardise the returns

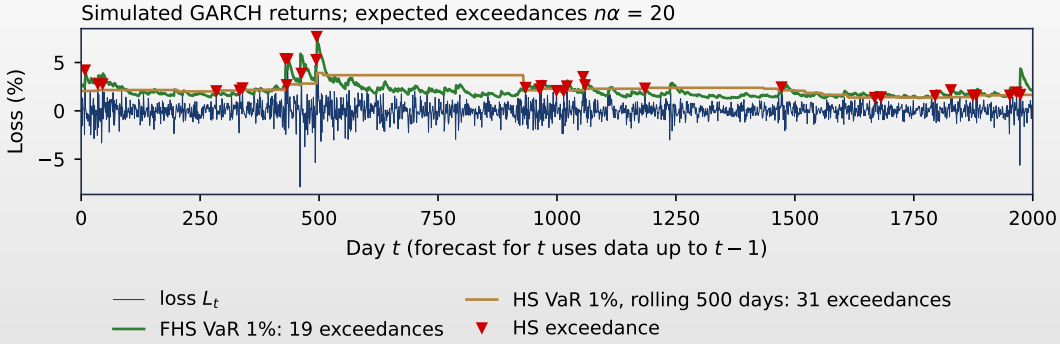
$$r_t = \mu_t + \sigma_t z_t, \quad \mu_t = c + \phi\, r_{t-\mathbf{1}}, \quad \sigma_t^{\mathbf{2}} = \omega + \alpha_{\mathbf{1}}(r_{t-\mathbf{1}} - \mu_{t-\mathbf{1}})^{\mathbf{2}} + \beta_{\mathbf{1}}\sigma_{t-\mathbf{1}}^{\mathbf{2}}, \quad \hat{z}_t = \frac{r_t - \hat{\mu}_t}{\hat{\sigma}_t}$$

  - μ_t: conditional mean (AR(1): c constant, ϕ autoregressive coefficient); σ_t: conditional volatility; z_t: i.i.d. shock with mean 0, variance 1; $\hat{z}_t$: standardised residual
  - $\omega > 0$; $\alpha_{\mathbf{1}}$: reaction to yesterday's shock; $\beta_{\mathbf{1}}$: memory; $\alpha_{\mathbf{1}} + \beta_{\mathbf{1}} < 1$
- Tomorrow's VaR and ES: forecast scale times a residual tail statistic, minus the forecast mean

$$\mathrm{VaR}_{t+\mathbf{1}} = -\hat{\mu}_{t+\mathbf{1}} + \hat{\sigma}_{t+\mathbf{1}}\,\hat{\xi}_\alpha, \qquad \hat{\xi}_\alpha = -q_\alpha(\hat{z}), \qquad \mathrm{ES}_{t+\mathbf{1}} = -\hat{\mu}_{t+\mathbf{1}} + \hat{\sigma}_{t+\mathbf{1}}\,\overline{(-\hat{z})}_{\geq \hat{\xi}_\alpha}$$

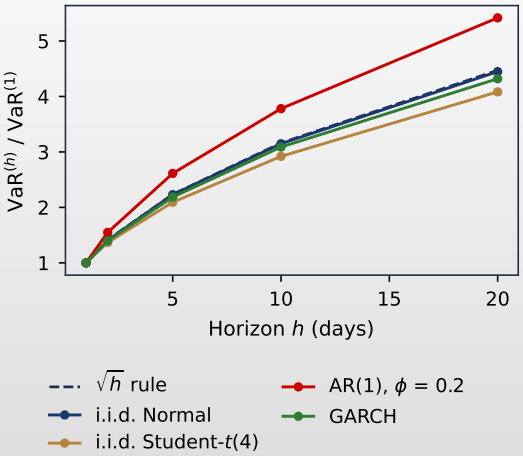
  - $\overline{(-\hat{z})}_{\geq \hat{\xi}_\alpha}$: the average of the worst fraction α of the residual losses $-\hat{z}_t$, with the fractional weight of historical simulation
  - Example: $\hat{\mu}_{t+\mathbf{1}} = 0.1\%$, $\hat{\sigma}_{t+\mathbf{1}} = 3\%$, $\hat{\xi}_{\mathbf{1\%}} = 2.8$: $\mathrm{VaR}_{\mathbf{1\%}} = -0.1 + 3 \times 2.8 = 8.3\%$

Filtered VaR and Exceedances



- Exceedance on day t: $L_t > \mathrm{VaR}_t$, with VaR_t computed from data up to day $t-1$ only (out of sample)
  - Under a correct model the count N over n days is $\mathrm{Bin}(n, \alpha)$: mean $n\alpha$, standard deviation $\sqrt{n\alpha(1-\alpha)}$; $n = 2\,000$, $\alpha = 1\%$: 20 ± 4.4
  - Chart (simulated GARCH): FHS follows the volatility and gives 19 exceedances; the rolling HS VaR reacts late and gives 31, grouped after volatility jumps

Longer Horizons



- Horizon of h days: log losses add up

$$L_t^{(h)} = L_t + L_{t-\mathbf{1}} + \cdots + L_{t-h+\mathbf{1}}$$

  - $\mathrm{VaR}^{(h)}$: quantile of these sums; consecutive sums share $h-1$ days (overlap)
  - Simple returns compound: $100\big[1 - \prod_{i=\mathbf{1}}^{h}(1 + R_{t+i}/100)\big]$, R in %

- Square-root-of-time rule: $\mathrm{VaR}^{(h)} = \sqrt{h}\,\mathrm{VaR}^{(\mathbf{1})}$
  - Exact, in particular, for i.i.d. zero-mean Normal losses: the variance of a sum of h independent days is h times the daily one
  - Chart (simulated): positive autocorrelation pushes the ratio above $\sqrt{h}$; fat tails pull it below

Peaks Over Threshold: the GPD

- Excesses $Y = L - u$ over a high threshold u follow approximately a **generalised Pareto distribution** (GPD) McNeil, Frey & Embrechts (2015)

$$P(Y \le y \mid L > u) \approx G_{\xi,\beta}(y) = 1 - \Big(1 + \frac{\xi y}{\beta}\Big)^{-1/\xi}, \qquad y \ge 0$$

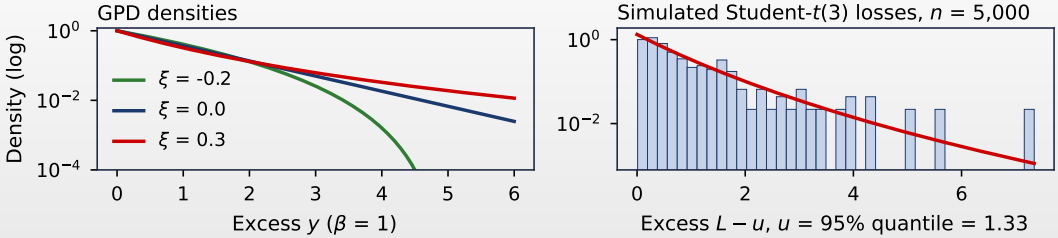
  - ξ: shape; $\xi > 0$: heavy, power-law tail (moments of order $\ge 1/\xi$ are infinite); $\xi = 0$: Gumbel domain, the normalised excesses have an exponential limit ($1 - e^{-y/\beta}$), e.g. for the Normal distribution; $\xi < 0$: bounded tail; $\beta > 0$: scale

- Tail VaR and ES (n: sample size, N_u: number of exceedances, $\xi < 1$)

$$\mathrm{VaR}_\alpha = u + \frac{\beta}{\xi}\Big[\Big(\frac{n}{N_u}\,\alpha\Big)^{-\xi} - 1\Big], \qquad \mathrm{ES}_\alpha = \frac{\mathrm{VaR}_\alpha + \beta - \xi u}{1 - \xi}$$

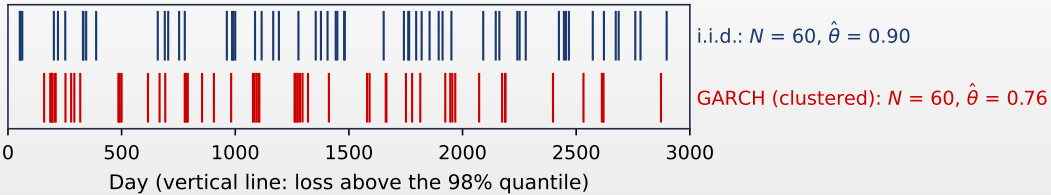
  - N_u/n estimates $P(L > u)$; the GPD extrapolates beyond the largest observed losses, where HS has no data
  - i.i.d. standard errors of the maximum-likelihood estimates: $\mathrm{SE}(\hat\xi) = (1 + \hat\xi)/\sqrt{N_u}$, $\mathrm{SE}(\hat\beta) = \hat\beta\sqrt{2(1 + \hat\xi)/N_u}$; choice of u: bias (too low) against variance (too high)

The GPD on Simulated Losses



- Left: GPD densities with $\beta = 1$ on a log scale: $\xi = -0.2$ ends at $y = 5$, $\xi = 0$ is a straight line (exponential), $\xi = 0.3$ decays as a power
- Right: 5 000 Student-t(3) losses, threshold at the 95% quantile, 250 excesses; the fitted $\hat{\xi} = 0.13$ is below the true tail value $1/3$: with 250 excesses $\mathrm{SE}(\hat{\xi}) \approx 0.07$, and a threshold that is not high enough biases $\hat{\xi}$

Clustered Exceedances: the Extremal Index

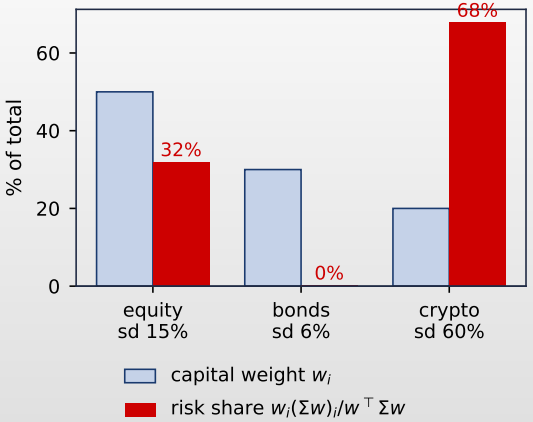


- **Extremal index** $\theta \in (0, 1]$: the reciprocal of the mean cluster size of exceedances; $\theta = 1$: exceedances do not cluster asymptotically (e.g. i.i.d.; it does not imply independence of the series), $\theta = 0.5$: clusters of two on average

$$\hat{\theta} = \min\left\{1,\ \frac{2\left(\sum_i (T_i - 1)\right)^2}{(N-1)\sum_i (T_i - 1)(T_i - 2)}\right\} \quad \text{if } \max_i T_i > 2$$

  - Intervals estimator Ferro & Segers (2003): $S_{\mathbf{1}} < \cdots < S_N$: days of the N exceedances; $T_i = S_{i+\mathbf{1}} - S_i$: gaps between them; the chart: exceedances of the 98% quantile, i.i.d. against GARCH
  - About θN independent clusters: i.i.d. standard errors computed from N exceedances are too small when $\theta < 1$

Risk Contributions (Euler Allocation)

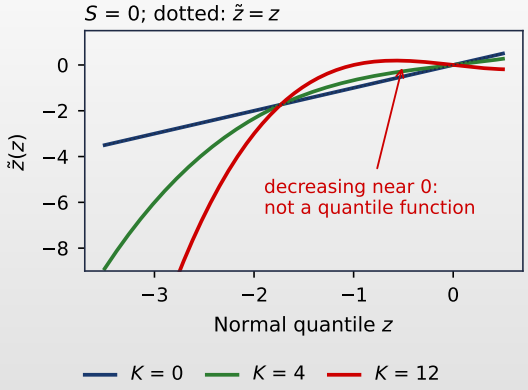


- Normal portfolio: exposures w, covariance matrix Σ, mean neglected

$$\sigma_p = \sqrt{w^\top \Sigma w}, \qquad \mathrm{VaR} = z_{\mathbf{1}-\alpha}\sigma_p,$$
$$\frac{\partial \mathrm{VaR}}{\partial w_i} = z_{\mathbf{1}-\alpha}\frac{(\Sigma w)_i}{\sigma_p}, \qquad C_i = w_i \frac{\partial \mathrm{VaR}}{\partial w_i}$$

  - Marginal VaR $\partial \mathrm{VaR}/\partial w_i$: change of VaR per extra unit in asset i; component VaR C_i Tasche (2008): $\sum_i C_i = \mathrm{VaR}$ (Euler: VaR is homogeneous of degree 1)
  - ES contributions: $\mathbb{E}[-w_i R_i \mid L_p \geq \mathrm{VaR}]$, the mean of asset i's loss on the tail days of the portfolio
- Diversification benefit: $\sum_i z_{\mathbf{1}-\alpha} w_i \sigma_i - \mathrm{VaR} \geq 0$
  - Risk parity: weights with equal shares, $w_i(\Sigma w)_i = w_j(\Sigma w)_j$ for all i, j
  - Inputs (equity, bonds, crypto): weights 50/30/20%, SDs 15/6/60% on the same horizon; correlations $\rho_{\mathbf{12}} = -0.2$, $\rho_{\mathbf{13}} = 0.3$, $\rho_{\mathbf{23}} = 0$. Crypto: 20% of capital, about two thirds of risk

The Cornish–Fisher (CF) Expansion



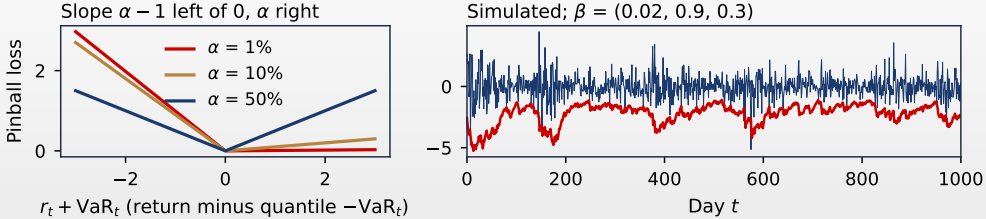
- Adjusts the Normal quantile for skewness S and excess kurtosis K Cornish & Fisher (1938)

$$\tilde{z} = z + \frac{(z^2-1)S}{6} + \frac{(z^3-3z)K}{24} - \frac{(2z^3-5z)S^2}{36}$$

  - $S = \mathbb{E}[(X-\mu)^3]/\sigma^3$; $K = \mathbb{E}[(X-\mu)^4]/\sigma^4 - 3$ (0 for the Normal distribution); $\mathrm{VaR}_\alpha = -(\mu + \sigma\tilde{z}(z_\alpha))$
  - $K > 0$ pushes the far-tail quantile further out; $S < 0$ lengthens the left tail

- A valid quantile function must be increasing in z; chart ($S = 0$): with large K, $\tilde{z}$ decreases near 0, so the expansion stops being a quantile function

Quantile Regression, the Pinball Loss and CAViaR



- **Pinball loss**: the loss whose expected value is minimised by the true α-quantile $-\mathrm{VaR}_t$

$$\ell_\alpha(r_t, \mathrm{VaR}_t) = \left(\alpha - 1\{r_t < -\mathrm{VaR}_t\}\right)(r_t + \mathrm{VaR}_t)$$

  - An exceedance costs $(1-\alpha)$ per unit, a day above the quantile costs α per unit (left); a lower average pinball loss = better VaR forecasts
- **CAViaR** (Conditional Autoregressive VaR), SAV specification (Symmetric Absolute Value) Engle & Manganelli (2004): a recursion for VaR, fitted by minimising $\sum_t \ell_\alpha$
  - $\mathrm{VaR}_t = \beta_\mathbf{0} + \beta_\mathbf{1}\mathrm{VaR}_{t-\mathbf{1}} + \beta_\mathbf{2}|r_{t-\mathbf{1}}|$; $|\beta_\mathbf{1}| < 1$: persistence (in the lecture and in B7: $0 < \beta_\mathbf{1} < 1$); $\beta_\mathbf{2}$: reaction to $|r_{t-\mathbf{1}}|$ (right, simulated)

Comparing Forecasts and Regulatory Capital

- **Diebold–Mariano (DM) test** Diebold & Mariano (1995): are two forecasts equally accurate?

$$d_t = \ell_t^A - \ell_t^B, \qquad DM = \frac{\bar{d}}{\sqrt{\hat{\Omega}_d / n}} \approx N(0,1) \text{ under } H_0 : \mathbb{E}[d_t] = 0$$

  - ℓ_t^A, ℓ_t^B: losses (e.g. pinball) of models A and B on day t; $\bar{d}$: mean difference; $\hat{\Omega}_d$: HAC long-run variance of d_t; $DM < -1.96$: A significantly better at 5%

- **Joint VaR–ES loss** Fissler–Ziegel zero-homogeneous (FZ0) Patton, Ziegel & Chen (2019): scores a pair (v_t, e_t) = (return quantile, tail mean), $e_t \le v_t < 0$

$$\mathrm{FZ0} = -\frac{1\{r_t \le v_t\}\,(v_t - r_t)}{\alpha\, e_t} + \frac{v_t}{e_t} + \ln(-e_t) - 1$$

  - $\mathrm{FZ0}(cv_t, ce_t; cr_t) = \mathrm{FZ0}(v_t, e_t; r_t) + \ln c$: differences between the scores of two models, not the score itself, are homogeneous of degree zero

- **FRTB-style capital** (Fundamental Review of the Trading Book; C2, simplified) Basel Committee (2019): ES from a stressed 250-day window, scaled to the liquidity horizon

$$\text{Capital} = m_c \times \mathrm{ES}_{\mathbf{10}}^{stres} \times \sqrt{LH/10} \times \text{amount}$$

  - $m_c = 1.5$: multiplier; $\mathrm{ES}_{\mathbf{10}}^{stres}$: 10-day ES 2.5% in the worst window; LH: liquidity horizon in days (10 or 20)

Exercise Map

Exercise	Question	Type	Prerequisite or example
A1, A2	VaR and ES under the Normal and the Student-t distributions	Solved	–
A3	VaR and ES of a discrete default loss	Solved	–
A4	when VaR punishes diversification	Proposed	A3
B1	historical VaR and ES on BET, with bootstrap intervals	Solved	–
B2	how precise is an ES estimate?	Proposed	B1
B3	tomorrow's filtered VaR for Bitcoin	Proposed	B1
C1	what is wrong in an AI answer?	Proposed	B1
B4	a GPD tail for BET (after the lecture)	Solved	–
A5, B5	risk contributions of two assets and of a four-asset portfolio (after the lecture)	Solved; Proposed	A5
Extension packs	A2, A4, A5, B1, B3, B4, B5 Extended; A6–A8, B6, B7 (optional); team project: one of C2 and C3	Solved, Proposed and Open	named on each task

- **[Solved]**: full solution in the slides and the notebook, a model to copy; **[Proposed]**: you solve it, the solution is discussed in class; **[Open]**: a project question with no single correct answer (C2, C3)
- Part A: derivations on paper; Part B: estimation and inference on real data, each ending with an interpretation question; Part C: AI critique and open projects

Notation and Conventions Used in Every Exercise

- Return X (in %), loss $L = -X$; level α = tail probability; VaR and ES are **positive** numbers when they are losses (definitions: "VaR and ES")
- Empirical convention (all B exercises): `np.quantile(L, 1 - alpha)` and the tail mean in "Historical Simulation"
  - $\widehat{\mathrm{ES}}_\alpha$ is the fractional tail average: weight 1 on the $\lfloor n\alpha \rfloor$ largest losses, $n\alpha - \lfloor n\alpha \rfloor$ on the next; the plain mean of the losses $\geq \widehat{\mathrm{VaR}}_\alpha$ agrees with it only when $n\alpha$ is an integer
- Units: log losses in % of the position for one asset; simple returns for portfolios (they add across assets); percentage points (pp) for standard errors and interval widths
- Symbols that look alike: α = tail probability, α_1 = GARCH coefficient; ξ = GPD shape, $\hat{\xi}_\alpha$ = residual quantile (B3); w = money exposures (A5) or portfolio weights (B5)

A1 [Solved]: Normal VaR and ES — Task

- **Question**: what are VaR and ES of a position whose daily return is Normal?
- **Data**
  - $X \sim N(\mu, \sigma^2)$, daily log return in % of the position, $\mu = 0.04\%$, $\sigma = 1.2\%$; loss $L = -X$ with density f
  - $z_\alpha = \Phi^{-1}(\alpha) < 0$; Φ, φ: distribution and density functions of the standard Normal distribution
- **Tasks**
  1. Derive $\mathrm{ES}_\alpha = -\mu + \sigma\varphi(z_\alpha)/\alpha$ from $\mathrm{ES}_\alpha = -\mathbb{E}[X \mid X \le q_\alpha(X)]$.
  2. Compute z_α and $\varphi(z_\alpha)$ at $\alpha = 1\%$ and $\alpha = 2.5\%$.
  3. Compute VaR and ES at both levels, in % of the position.
- **Report**: a four-cell table (VaR, ES $\times$ 1%, 2.5%)
- **Interpretation**
  - which of the four numbers are almost equal?
  - why does that matter for regulation?
- Notebook: section A1

A1 [Solved]: Solution — Normal VaR and ES

- **1. Derivation.** Write $X = \mu + \sigma Z$; $X \le q_\alpha$ is $Z \le z_\alpha$, of probability α
  - Since $\varphi'(u) = -u\varphi(u)$: $\int_{-\infty}^{z} u\varphi(u)\,du = [-\varphi(u)]_{-\infty}^{z} = -\varphi(z)$
  - $\mathbb{E}[X \mid X \le q_\alpha] = \mu + \dfrac{\sigma}{\alpha}\displaystyle\int_{-\infty}^{z_\alpha} u\varphi(u)\,du = \mu - \sigma\dfrac{\varphi(z_\alpha)}{\alpha}$; change the sign
- **2–3.** $\mathrm{VaR}_\alpha = -(\mu + \sigma z_\alpha)$, e.g. $-(0.04 + 1.2 \cdot (-2.326)) = 2.75$; $\mathrm{ES}_{\mathbf{1\%}} = -0.04 + 1.2 \cdot 2.665 = 3.16$

α	z_α	$\varphi(z_\alpha)$	φ/α	VaR_α	ES_α
1%	−2.326	0.0267	2.665	2.75	3.16
2.5%	−1.960	0.0584	2.338	2.31	2.77

- VaR and ES in % of the position; for a position of 1 million currency units: VaR 1% $\approx$ 27 516, ES 1% $\approx$ 31 583
- Interpretation: ES 2.5% (2.77%) $\approx$ VaR 1% (2.75%): under the Normal distribution the switch to ES 2.5% keeps the capital, which is why FRTB chose this level

A2 [Solved]: Student-t ES — Task

- **Question**: how much do heavy tails change VaR and ES when the standard deviation is the same as in A1?
- **Data**
  - $X = m + sT$, $T \sim t_\nu$ with density g_ν, $t_\alpha = t_\nu^{-1}(\alpha) < 0$; $\nu = 4$, $m = 0.04\%$, standard deviation 1.2%
  - Tail integral (proof in Extension Pack 1): $\int_{-\infty}^{t} x g_\nu(x)\,dx = -g_\nu(t)\,\dfrac{\nu + t^2}{\nu - 1}$ for $\nu > 1$
- **Tasks**
  1. Use the tail integral to write $\mathrm{ES}_\alpha = -m + s\,\dfrac{g_\nu(t_\alpha)}{\alpha}\,\dfrac{\nu + t_\alpha^2}{\nu - 1}$.
  2. Find the scale s that gives a standard deviation of 1.2%, using $\mathrm{Var}(T) = \nu/(\nu - 2)$ (finite only for $\nu > 2$).
  3. Compute VaR and ES at 1% and 2.5% and put them next to A1.
- **Report**: the A1 table with a Student-t column
- **Interpretation**
  - at which level do the two models differ most?
  - why does ES see the difference first?
- Common mistake: using $\sigma = 1.2$ instead of s as the scale

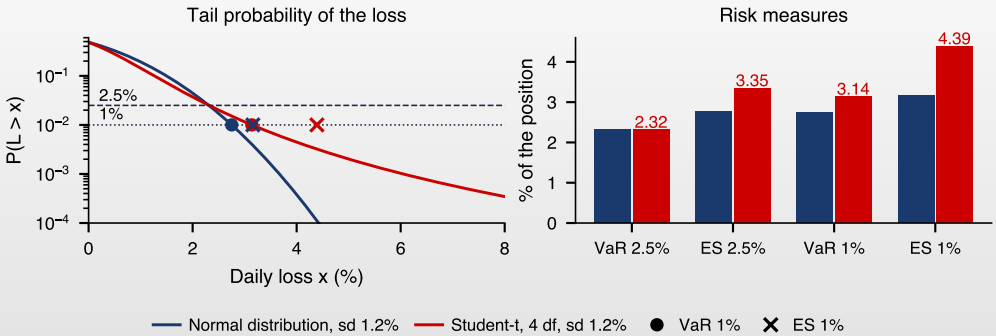
Back Exercise list

A2 [Solved]: Solution — Student-t ES (1/2)

- **1.** Divide the tail integral by α and rescale with $X = m + sT$, as in A1
- **2.** $s^2 \cdot 4/2 = 1.2^2$, so $s = 1.2\sqrt{2/4} = 0.849$
- **3.** At 1%: $t_{1\%} = -3.747$, $g_4(t_{1\%}) = 0.0087$, ES factor $\dfrac{0.0087}{0.01} \cdot \dfrac{4 + 14.04}{3} = 5.221$
  - $\mathrm{VaR}_{1\%} = -(0.04 + 0.849 \cdot (-3.747)) = 3.14$; $\mathrm{ES}_{1\%} = -0.04 + 0.849 \cdot 5.221 = 4.39$

	$\mathrm{VaR}_{2.5\%}$	$\mathrm{ES}_{2.5\%}$	$\mathrm{VaR}_{1\%}$	$\mathrm{ES}_{1\%}$
Normal	2.31	2.77	2.75	3.16
t_4	2.32	3.35	3.14	4.39

A2 [Solved]: Solution — Student-t ES (2/2)



- Interpretation: VaR 2.5% is the same in both models; VaR 1% and both ES are much higher under t_4: heavy tails show only deep in the tail, and ES sees them first

MFM_ch7_seminar

A3 [Solved]: VaR and ES of a Discrete Loss — Task

- **Question**: what are VaR and ES when the loss has an atom (a default)?
- **Data**: a bond loses 60 monetary units (loss given default, LGD) with probability $p = 3\%$ and 0 otherwise
- **Tasks**
  1. Compute VaR 5% and VaR 2.5% from the infimum definition of VaR ("VaR and ES").
  2. Compute ES 5% and ES 2.5% with the formula for a loss with an atom ("Losses with an Atom") Acerbi & Tasche (2002).
  3. Explain why $\mathrm{ES}_{5\%}$ differs from $\mathbb{E}[L \mid L \geq \mathrm{VaR}_{5\%}]$.
- **Report**: four numbers and one sentence

Loss L	0	60
$P(L)$	97%	3%
$P(L \geq \cdot)$	100%	3%

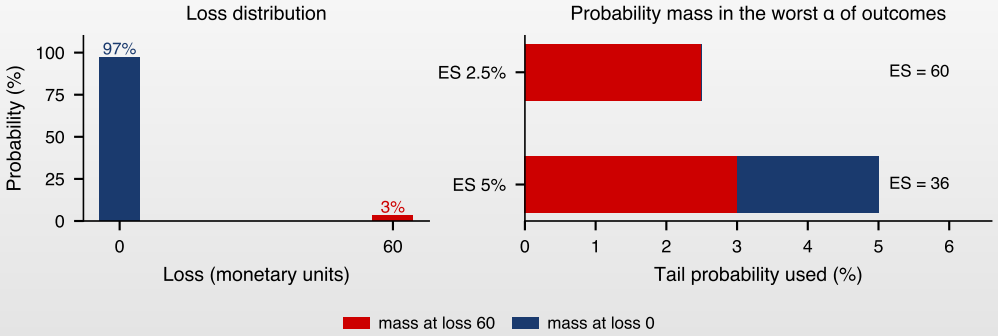
- The worst α of outcomes must be filled with probability mass; at $\alpha = 5\%$ the 3% default mass is not enough

Back Exercise list

A3 [Solved]: Solution — VaR and ES of a Discrete Loss (1/2)

- **1.** $P(X \le -60) = 3\% \le 5\%$, $P(X \le 0) = 100\% > 5\%$: $q_{\mathbf{5\%}} = 0$, $\mathrm{VaR}_{\mathbf{5\%}} = 0$; $P(X \le -60) = 3\% > 2.5\%$: $\mathrm{VaR}_{\mathbf{2.5\%}} = 60$
- **2.** $\mathrm{ES}_{\mathbf{5\%}} = \dfrac{60 \cdot 0.03 + 0 \cdot (0.05 - 0.03)}{0.05} = 36$; $\mathrm{ES}_{\mathbf{2.5\%}} = \dfrac{0 + 60 \cdot (0.025 - 0)}{0.025} = 60$
- **3.** $\mathbb{E}[L \mid L \ge 0] = 1.8$ averages over **all** outcomes; ES 5% uses the 3% default mass plus 2% of the zero-loss atom (chart)
- Interpretation: the same fractional weight gives the historical ES of B1: the last loss in the tail enters with weight $n\alpha - \lfloor n\alpha \rfloor$

A3 [Solved]: Solution — VaR and ES of a Discrete Loss (2/2)



MFM_ch7_seminar

A4 [Proposed]: When VaR Punishes Diversification — Task

- **Question**: can a VaR limit make a diversified portfolio look riskier than its parts?
- **Data**: two independent bonds A and B; each loses 100 with probability $p = 0.7\%$ (the lecture used 0.9%) and 0 otherwise; portfolio A + B Artzner et al. (1999)
- **Tasks**
  1. Write the distribution of the portfolio loss: the values 0, 100, 200 and their probabilities.
  2. Compute VaR 1% and ES 1% of one bond and of the portfolio with the discrete formula of A3.
  3. Show that VaR 1% is not subadditive here while ES 1% is.
  4. Recompute both risk measures and decide whether subadditivity holds at $\alpha = 1.5\%$.
- **Report**: a table of VaR and ES for A, B and A + B at both levels
- **Interpretation**: what would a VaR-based limit reward?
- **Prerequisites and examples**: A3

Back Exercise list

Which Files Define the Sample?

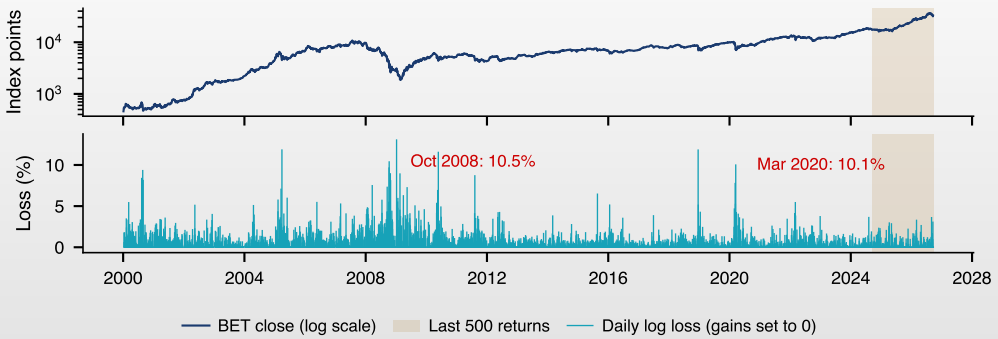
- Course market files use a `date` index. For BET, S&P 500, Bitcoin and gold use `close` in `BET.csv`, `GSPC.INDX.csv`, `BTC-USD.CC.csv`, `XAUUSD.FOREX.csv`; all are in data/market.
- Portfolio ETFs use `adjusted_close`: `SPY.US.csv`, `TLT.US.csv`, `GLD.US.csv`. Bitcoin retains its close; align prices on common observed dates before differencing.
- The BVB basket uses adjusted closes of (BVB ticker in parentheses) Banca Transilvania (TLV), OMV (Österreichische Mineralölverwaltung, Austrian Mineral Oil Administration, OMV group) Petrom (SNP), BRD–Groupe Société Générale (BRD), Transgaz (TGN), Romgaz (SNG), Nuclearelectrica (SNN), Electrica (EL) and Transelectrica (TEL), each in `SYMBOL.RO.csv`. Equal weights apply to simple returns, not their logs.
- EUR/RON comes from the annual BNR XML (eXtensible Markup Language) archive, EUR field: download the annual files `https://curs.bnr.ro/files/xml/years/nbrfxrates<YEAR>.xml`, for each year 2005–2026. Retain the sample endpoints specified in each task.

B1 [Solved]: Historical VaR and ES on BET — Task

- **Questions**
  - how large are the tail losses of the Romanian stock market?
  - how precise are the estimates?
- **Data**
  - BET index, daily closes 2000-01-05 – 2026-09-18, weekdays only; $n = 6\,670$ daily log returns in %, $r_t = 100\,(\ln P_t - \ln P_{t-1})$, loss $L_t = -r_t$
  - Two samples: the full period and the last 500 returns (from 2024-09-12, shaded in the chart)
- **Tasks**
  1. Compute historical VaR and ES at 1% and 2.5% on both samples, and count the losses beyond VaR 1%.
  2. Build a 95% percentile bootstrap interval for VaR 1% with 2 000 i.i.d. resamples, seed 42 Efron (1979).
  3. Repeat step 2 with a moving-block bootstrap: 20-day blocks drawn with replacement, concatenated and cut to length n Künsch (1989).
- **Report**: a table of the eight VaR/ES numbers and the two intervals per sample
- **Interpretation**
  - which estimate would you give a risk committee?
  - why is the block interval wider on the full sample?

Back Exercise list

B1 [Solved]: Historical VaR and ES on BET — Data



- Look first: the largest losses (October 2008: 10.45%; March 2020: 10.08%) come in clusters

MFM_ch7_seminar

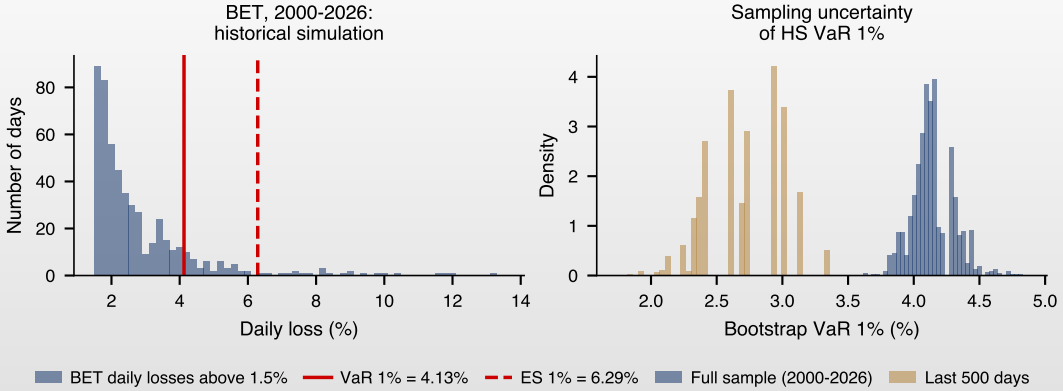
B1 [Solved]: Solution (1/4) — Step by Step

- **1. Last 500 days.** Sort the losses; VaR 1% sits at position $1 + 499 \cdot 0.99 = 495.01$; $L_{(i)}$: the i-th smallest loss
  - $\widehat{\mathrm{VaR}}_{\mathbf{1\%}} = L_{(\mathbf{495})} + 0.01\,(L_{(\mathbf{496})} - L_{(\mathbf{495})}) = 2.710 + 0.01\,(2.925 - 2.710) = 2.712\%$
  - The five losses beyond it: 3.68, 3.35, 3.12, 3.03, 2.93: their mean is $\widehat{\mathrm{ES}}_{\mathbf{1\%}} = 3.22\%$

	n	$\mathrm{VaR}_{\mathbf{1\%}}$	$\mathrm{ES}_{\mathbf{1\%}}$	$\mathrm{VaR}_{\mathbf{2.5\%}}$	$\mathrm{ES}_{\mathbf{2.5\%}}$	beyond VaR 1%
2000–2026	6 670	4.13	6.29	2.80	4.57	67
Last 500 days	500	2.71	3.22	2.05	2.72	5

- The 500-day ES 1% rests on five losses, ES 2.5% on $n\alpha = 12.5$ (the 13th with weight 0.5): report them with these counts
- Both are unconditional summaries: the full sample describes 26 years, the last 500 days a calm period; tomorrow's risk needs a conditional forecast (B3)

B1 [Solved]: Solution (2/4) — The Tail and the Bootstrap



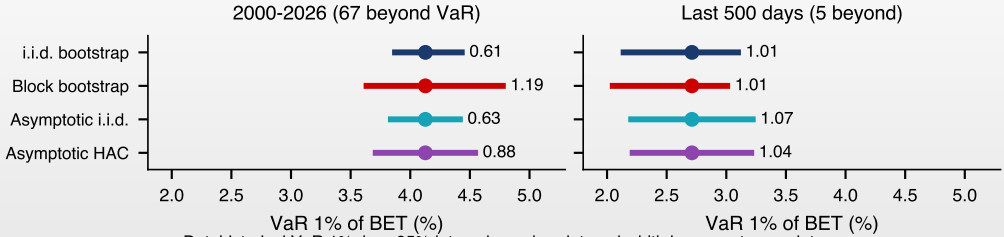
MFM_ch7_seminar

B1 [Solved]: Solution (3/4) — The Tail and the Bootstrap

- Left: BET losses above 1.5% with historical VaR 1% and ES 1%; right: the 2 000 i.i.d. bootstrap values of VaR 1% for each sample
- Each resample recomputes both the quantile and which losses lie beyond it; the 2.5% and 97.5% percentiles of the 2 000 values give the interval
- The two bootstrap distributions do not overlap: the recent market is calmer than the 26-year average

MFM_ch7_seminar

B1 [Solved]: Solution (4/4) — Intervals



Dot: historical VaR 1%; bar: 95% interval; number: interval width in percentage points

- **2.** i.i.d. bootstrap: full $[3.85, 4.46]$, last 500 days $[2.11, 3.12]$
- **3.** Block bootstrap: $[3.61, 4.80]$ and $[2.02, 3.03]$

- Interpretation: exceedances cluster in 2008 and 2020, so blocks that keep a crisis together almost double the width of the full-sample bootstrap interval; the calm last 500 days show no such effect

MFM_ch7_seminar

B2 [Proposed]: How Precise Is an ES Estimate? — Task

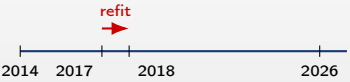
- **Question**: does a 500-day historical ES 2.5% pin the tail risk down to within ± 0.3 pp Chen (2008), Pele & Mazurencu-Marinescu-Pele (2026)?
- **Data**
  - S&P 500 (daily close) and BET, daily log losses in %, 2000–2026; samples: full and last 500 days
  - Use the resampling function of B1 (2 000 resamples, seed 42, 20-day blocks) with ES 2.5% as the statistic
- **Tasks**
  1. Compute ES 2.5%, the tail count $n\alpha$ and the number k of losses with a positive weight in the estimator.
  2. Draw one i.i.d. and one block resample of the S&P 500 last 500 days, and recompute VaR 2.5%, the tail membership and ES on each.
  3. Build the i.i.d. and the block 95% intervals for all four samples and compare their widths with 0.6 pp (± 0.3).
- **Report**: a table ($n\alpha$, k, ES, two intervals) and an interval chart
- **Interpretation**
  - which samples meet the ± 0.3 pp target?
  - what governs the width?
- **Prerequisites and examples**: B1

Back Exercise list

B3 [Proposed]: Next-Day Filtered VaR for Bitcoin — Task

- **Question**: Bitcoin's VaR 1% for 19 September 2026: filtered or historical?
- **Data**
  - Bitcoin (USD), daily close, 17 September 2014 – 18 September 2026, 7 days a week; daily log returns in %
  - AR(1)-GARCH(1,1) by Normal QML, with an AR intercept: $r_t = \mu_t + \sigma_t z_t$ Barone-Adesi et al. (1999); `garch_filter(r)` calls `arch_model(r, mean="AR", lags=1, vol="GARCH", p=1, q=1, dist="normal", rescale=False)`, and `garch_next(r, params)` returns $\hat{\mu}_{t+1}$, $\hat{\sigma}_{t+1}$
- **Tasks**
  1. Estimate the model on the full sample and forecast $\hat{\mu}_{t+1}$, $\hat{\sigma}_{t+1}$.
  2. Compute $\hat{\xi}_{1\%} = -q_{1\%}(\hat{z})$, the fractional tail mean of $-\hat{z}$ beyond it and the filtered VaR 1% and ES 1% ("Filtered Historical Simulation").
  3. Compare with historical VaR and ES: last 500 days and full sample.
  4. With the forecasts of the timeline, count the days in 18 September 2024 – 18 September 2026 with a loss above the FHS VaR and above the rolling 500-day HS VaR.
- **Report**: three VaR/ES pairs, two exceedance counts and the expected count $n\alpha$ for the 731 evaluation days
- **Prerequisites and examples**: B1 (historical quantiles)

Forecast timeline (step 4)



each January: GARCH on all earlier data
each day: $\hat{\sigma}_{t+1}$ from data up to t
residuals used: those before that January

- No forecast uses information from its own day or later: a full-sample threshold is not an out-of-sample forecast
- **Interpretation**
  - why do the three VaRs differ?
  - which method misses more often?

Back Exercise list

C1 [Proposed]: Critique an AI Answer — the Answer

- **Prompt**: "px holds daily BET closes, 2000–2026. Give me the one-day historical VaR 1% and ES 2.5% as positive log losses in %, and the 10-day VaR 1%."
- **Answer**: the code below, which runs without an error message

```
import numpy as np
r = 100 * np.log(px).diff().dropna()                # daily log returns, %
loss = -r                                           # log loss, %
var1 = np.percentile(r, 99)                         # VaR 1%
es25 = loss[loss >= np.percentile(loss, 99)].mean()  # ES 2.5%
var10 = 10 * var1                                   # 10-day VaR 1%
```

- **Conclusion of the AI**: "VaR 1% = 3.95%, ES 2.5% = 6.28%, 10-day VaR 1% = 39.46%."

Back Exercise list

C1 [Proposed]: Critique an AI Answer — Task

- **Questions**
  - which lines of the AI answer are wrong?
  - what are the correct numbers?
- **Data**: the BET data of B1; the answer contains **three** mistakes typical for VaR and ES, everything else is correct
- **Tasks**
  1. For each mistake, name the line number and explain in one sentence why it is wrong, using $\mathrm{VaR}_\alpha = -q_\alpha(r)$ (for a continuous loss, $= q_{\mathbf{1}-\alpha}(L)$).
  2. Say how you would detect it: a count of observations, a number from B1, or a slide of the lecture.
  3. Write the corrected lines, including the 10-day VaR from overlapping 10-day losses, and recompute the three numbers.
  4. Check the corrected ES 2.5%: how many losses does it average? How many should it average?
- **Report**: an error table (line, error, detection, correction) and the corrected numbers; log each error you actually found in `AI_ERRORS.md`
- **Interpretation**: why does the code run without an error message although the answer is wrong?
- **Prerequisites and examples**: B1

B4 [Solved]: A GPD Tail for BET — Task

- **Question**: what does an extreme-value tail add beyond historical simulation, deep in the tail?
- **Data**
  - The BET losses of B1 ($n = 6\,670$)
  - Peaks over threshold McNeil, Frey & Embrechts (2015): the GPD with shape ξ and scale β, its standard errors and the tail VaR and ES (formulas: "Peaks Over Threshold: the GPD"); N_u = number of exceedances
- **Tasks**
  1. Set u = the empirical 95% quantile of L and form the excesses $L - u$ of the losses above it.
  2. Fit the GPD by maximum likelihood with location fixed at 0: `scipy.stats.genpareto.fit(excesses, floc=0)`.
  3. Compute the i.i.d. standard errors $\mathrm{SE}(\hat{\xi})$ and $\mathrm{SE}(\hat{\beta})$.
  4. Compute GPD VaR and ES at 1%, 0.5% and 0.1% and compare with historical simulation.
  5. Refit for thresholds leaving 10% down to 1.5% of the losses above them and plot $\hat{\xi}$ with 95% bands.
- **Report**: the fit with standard errors, a six-row VaR/ES table with tail counts, and the stability chart
- **Interpretation**
  - is the BET tail heavier than the Normal distribution?
  - how sure are we?

Back Exercise list

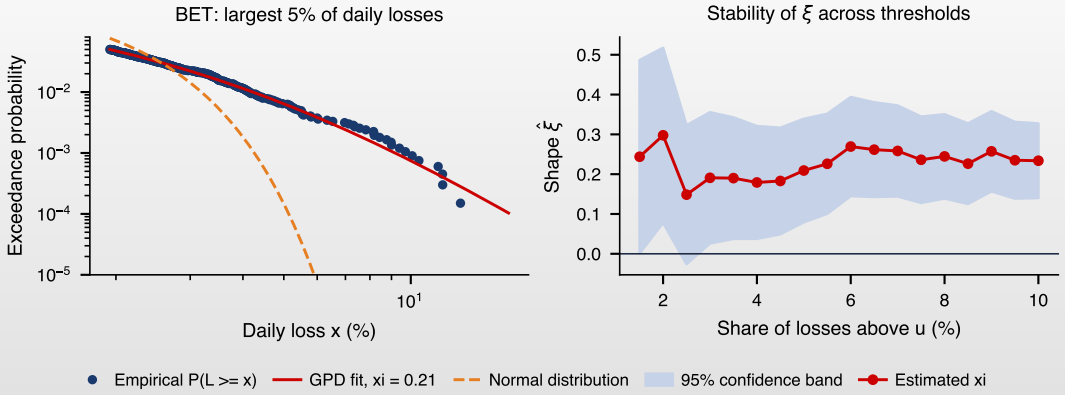
B4 [Solved]: Solution (1/3) — Fit and Table

- **1–2.** $u = 1.93\%$ leaves $N_u = 334$ of $n = 6\,670$ losses above it; $\hat{\xi} = 0.209$, $\hat{\beta} = 1.20$
- **3.** $\mathrm{SE}(\hat{\xi}) = (1 + 0.209)/\sqrt{334} = 0.066$, $\mathrm{SE}(\hat{\beta}) = 0.10$; $\hat{\xi}/\mathrm{SE} = 3.2$
- **4.** VaR 1%: $N_u/n = 0.050$, $(0.01/0.050)^{-\mathbf{0.209}} = 1.401$
  - $\mathrm{VaR}_{\mathbf{1\%}} = 1.93 + \frac{\mathbf{1.20}}{\mathbf{0.209}}(1.401 - 1) = 4.22\%$

	VaR GPD	VaR hist.	ES GPD	ES hist.
1%	4.22	4.13	6.34	6.29
0.5%	5.47	5.38	7.92	7.90
0.1%	9.18	9.50	12.61	11.33

- Losses beyond historical VaR: 67 at 1%, 34 at 0.5%, 7 at 0.1%
- The two methods agree down to 0.5%; at 0.1% historical ES rests on $n\alpha = 6.7$ losses, while the GPD uses all 334 excesses through a fitted tail

B4 [Solved]: Solution (2/3) — Diagnostics



MFM_ch7_seminar

B4 [Solved]: Solution (3/3) — Diagnostics

- Left: the GPD follows the empirical tail, the Normal tail falls far below it; right: $\hat{\xi}$ ranges from 0.15 to 0.30 across thresholds, while VaR 0.1% moves only between 9.04% and 9.51%
- Interpretation: the tail is heavier than the Normal one, but the i.i.d. SE assumes independent exceedances; BET losses cluster, so the honest uncertainty is larger (B4 Extended: extremal index)

MFM_ch7_seminar

A5 [Solved]: Risk Contributions of Two Assets — Task

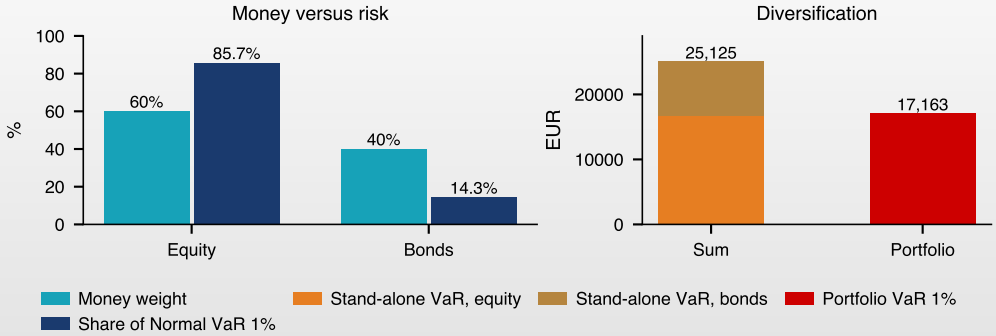
- **Question**: does each asset carry risk in proportion to the money invested in it?
- **Data**
  - Money exposures $w = (600\,000,\ 400\,000)$ EUR in equity and bonds; daily simple returns R in decimals, standard deviations $\sigma_1 = 0.012$, $\sigma_2 = 0.009$, correlation $\rho = -0.2$
  - $\Sigma = \begin{pmatrix} 0.000144 & -0.0000216 \\ -0.0000216 & 0.000081 \end{pmatrix}$; ($\Sigma$: covariance matrix) jointly Normal returns, mean ignored, $\alpha = 1\%$, $z_{99\%} = 2.326$ (the 99% standard Normal quantile)
- **Tasks**
  1. Compute the portfolio standard deviation σ_p in EUR and the portfolio VaR 1% (formulas: "Risk Contributions (Euler Allocation)").
  2. Compute Σw, the marginal VaR per EUR and the component VaR C_i; check $C_1 + C_2 = \mathrm{VaR}$ Tasche (2008).
  3. Compute the diversification benefit: the sum of stand-alone VaRs minus the portfolio VaR.
  4. Under the Normal distribution, rescale each C_i by $\dfrac{\varphi(z_{1\%})/0.01}{z_{99\%}}$ to get the ES 1% contributions.
- **Report**: a table of C_i, C_i^{ES} and shares, and the diversification benefit in EUR
- **Interpretation**: do the bonds carry risk in proportion to their 40% money weight?

Back Exercise list

A5 [Solved]: Solution — Risk Contributions of Two Assets (1/2)

- **1.** $\sigma_p^2 = 7\,200^2 + 3\,600^2 + 2(-0.2)(7\,200)(3\,600) = 54\,432\,000$, $\sigma_p = 7\,378$ EUR, $\mathrm{VaR}_{1\%} = 2.326 \cdot 7\,378 = 17\,163$ EUR
- **2.** $\Sigma w = (77.76, 19.44)$; marginal VaR per EUR $(0.0245, 0.0061)$; $C_1 = 14\,711$ (85.7%), $C_2 = 2\,452$ (14.3%), sum 17 163 EUR
- **3.** Stand-alone $16\,750 + 8\,375 = 25\,125$ EUR vs 17 163 EUR
- **4.** Factor 1.146: $C_1^{ES} = 16\,854$, $C_2^{ES} = 2\,809$, total 19 663 EUR = ES 1%; same shares as VaR under the Normal distribution

A5 [Solved]: Solution — Risk Contributions of Two Assets (2/2)



⊡ Interpretation: the bonds hold 40% of the money but 14.3% of the risk: the negative correlation makes them a partial hedge

MFM_ch7_seminar

B5 [Proposed]: Risk Contributions of a Four-Asset Portfolio — Task

- **Question**: in an equally weighted portfolio of SPY, TLT, GLD and Bitcoin, who carries the risk?
- **Data**
  - SPY, TLT, GLD (adjusted closes) and Bitcoin (USD closes), 18 September 2014 – 18 September 2026; align prices first, then retain 3 017 common daily simple returns R in %
  - Fractional weights $w_i = 25\%$, rebalanced daily; Σ = sample covariance of R (in $\%^2$); Normal VaR 1% with the mean ignored; historical ES 2.5%
- **Tasks**
  1. Compute the Normal VaR 1%, the marginal VaR and the component VaR of each asset, as in A5.
  2. Compute the historical ES 2.5% of the portfolio loss $L_p = -w^\top R$ and the contributions $\mathbb{E}[-w_i R_i \mid L_p \geq \mathrm{VaR}_{\mathbf{2.5\%}}]$ as averages over the tail days, with the weights of historical ES.
  3. Compute the undiversified VaR $\sum_i z_{\mathbf{99\%}} w_i \sigma_i$.
- **Report**: a table of weights, contributions and shares, and a chart of money weights vs risk shares
- **Interpretation**: is equal money equal risk? (comparing Normal VaR 1% with historical ES 2.5% changes both the model and the level)
- **Prerequisites and examples**: A5

Back Exercise list

Exit Check and Team Project

Exit check: can you say why

- a 500-day ES 1% should be reported with the number 5?
- $\sqrt{10}$ scaling is too low for BET?
- a 25% position can carry most of the risk (B5, after the lecture)?
- portfolio losses need simple, not log returns?

Team project

- Choose one of C2 or C3 (Extension Pack 4), with the common assessment criteria
- Log AI use in `AI_USE.md` and every error you actually caught in `AI_ERRORS.md` (no quota)

Extension Packs: What Each One Needs

Pack	Exercises	Prerequisites
1. Tail algebra	A2 Extended [Solved], A4 Extended, A6, A7, A5 Extended [Solved], A8	A1–A5
2. Inference	B1 Extended [Solved], B3 Extended, B6, B5 Extended	B1, B3, B5
3. Extremes and dynamics	B4 Extended [Solved], B7	B3, B4
4. Projects	C2 (ES capital, BVB vs S&P 500), C3 (deep quantile regression); assessment criteria	the lecture case study, B1–B3

A2 Extended [Solved]: The Student-t Tail Integral — Task and Solution

- **Question**: where does the Student-t ES factor of A2 come from?
- **Data**: $g_\nu(x) = c_\nu(1 + x^2/\nu)^{-(\nu+1)/2}$, the t_ν density; c_ν: the constant that makes it integrate to 1
- **Tasks**
  1. Differentiate $(1 + x^2/\nu)^{-(\nu-1)/2}$.
  2. Deduce $\int_{-\infty}^{t} x g_\nu(x)\,dx = -g_\nu(t)\,\frac{\nu+t^2}{\nu-1}$ and state where $\nu > 1$ is needed.
- **Report**: the two-line proof

Solution

- **1.** $\frac{d}{dx}(1 + x^2/\nu)^{-(\nu-1)/2} = -\frac{\nu-1}{\nu}\,x\,(1 + x^2/\nu)^{-(\nu+1)/2}$, so $x g_\nu(x) = -\frac{\nu c_\nu}{\nu-1}\,\frac{d}{dx}(1 + x^2/\nu)^{-(\nu-1)/2}$
- **2.** Integrating from $-\infty$ to t: $-\frac{\nu}{\nu-1} g_\nu(t)(1 + t^2/\nu) = -g_\nu(t)\,\frac{\nu+t^2}{\nu-1}$; the term at $-\infty$ vanishes only if $\nu > 1$
- ES needs $\nu > 1$; matching a variance needs $\nu > 2$

Back Exercise list

A4 Extended [Proposed]: For Which Levels Does VaR Fail? — Task

- **Question**: for which tail probabilities α is VaR not subadditive in A4?
- **Data**: the A4 bonds, $p = 0.7\%$, and the VaR definition in "VaR and ES"
- **Tasks**
  1. Write $\mathrm{VaR}_\alpha(A)$ and $\mathrm{VaR}_\alpha(A+B)$ as step functions of α.
  2. Find the interval of α where $\mathrm{VaR}_\alpha(A+B) > 2\mathrm{VaR}_\alpha(A)$ and explain both endpoints.
- **Report**: the interval and one sentence per endpoint
- **Prerequisites and examples**: A4

A6 [Proposed]: When Is Cornish–Fisher a Valid Quantile? — Task

- **Question**: for which excess kurtosis is the Cornish–Fisher map still a quantile function?
- **Data**
  - $\mu = 0$, $\sigma = 1.2\%$, skewness $S = -0.5$, excess kurtosis K Cornish & Fisher (1938)
  - The Cornish–Fisher map $\tilde{z}(z)$ in "The Cornish–Fisher (CF) Expansion"; $\mathrm{VaR}_\alpha = -\sigma\tilde{z}(z_\alpha)$
- **Tasks**
  1. With $K = 3$, compute the three correction terms at $z_{1\%}$, $\tilde{z}_{1\%}$ and the Cornish–Fisher VaR 1%.
  2. Write $d\tilde{z}/dz = Az^2 + Bz + C$ and locate the vertex of this parabola.
  3. Find the largest K for which $\tilde{z}$ is increasing on $[z_{0.1\%}, 0]$, and decide whether $K = 10$ is admissible.
- **Report**: the VaR, the bound on K and a plot of $\tilde{z}$ and $d\tilde{z}/dz$ for $K = 3$ and $K = 10$
- **Interpretation**: does a valid quantile map also give an accurate VaR?
- **Prerequisites and examples**: A1

Back Exercise list

A7 [Proposed]: ES as a Minimum, and Subadditivity — Task

- **Question**: why can ES never reward splitting a portfolio?
- **Data**
  - Integrable loss L; the objective $F_\alpha(v)$ Rockafellar & Uryasev (2000) in "ES as a Minimum"
  - Numbers: the A4 portfolio (two bonds, $p = 0.7\%$) at $\alpha = 1\%$
- **Tasks**
  1. For the A4 portfolio, compute the slopes of F on $(0, 100)$ and $(100, 200)$, the minimiser and the minimum.
  2. Show that the right derivative of F is $1 - P(L > v)/\alpha$ and deduce that VaR_α minimises F, with minimum ES_α.
  3. Deduce $\mathrm{ES}_\alpha(L_1 + L_2) \leq \mathrm{ES}_\alpha(L_1) + \mathrm{ES}_\alpha(L_2)$ from $(a + b)_+ \leq a_+ + b_+$.
- **Report**: the three numbers of step 1 and two short proofs
- **Prerequisites and examples**: A3, A4

A5 Extended [Solved]: ES Contributions from the Minimum — Task and Solution

- **Question**: why are the historical ES contributions of B5 averages over the tail days?
- **Data**: $\mathrm{ES}_\alpha(w) = \min_v F(w, v)$, $F(w, v) = v + \mathbb{E}[(L_p - v)_+]/\alpha$, $L_p = -w^\top R$ (A7)
- **Tasks**
  1. Use the envelope argument at $v^* = \mathrm{VaR}_\alpha$ to compute $\partial \mathrm{ES}_\alpha / \partial w_i$.
  2. Show that $w_i \, \partial \mathrm{ES}_\alpha / \partial w_i = \mathbb{E}[-w_i R_i \mid L_p \geq \mathrm{VaR}_\alpha]$ for a continuous L_p.
- **Report**: a short proof
- **Prerequisites and examples**: A7

Solution

- $\partial F / \partial v = 0$ at v^*, so $\partial \mathrm{ES}_\alpha / \partial w_i = \partial F / \partial w_i (w, v^*) = \mathbb{E}[-R_i 1\{L_p > v^*\}]/\alpha$
- For continuous L_p, $P(L_p > v^*) = \alpha$: this is $\mathbb{E}[-R_i \mid L_p \geq \mathrm{VaR}_\alpha]$; summing $w_i \times$ it gives ES_α (Euler)
- Historical estimate: average $-w_i R_i$ over the tail days with the weights of historical ES, so the contributions add up to it (B5); under the Normal distribution it equals the VaR rescaling of A5

Back Exercise list

A8 [Proposed]: The Sampling Variance of Historical ES — Task

- **Question**: how imprecise is a 500-day historical ES 2.5%, even for Normal losses?
- **Data**
  - n i.i.d. losses, continuous, positive density at VaR_α, $\mathbb{E}[L^2] < \infty$; $Y = (L - \mathrm{VaR}_\alpha)_+$, the excess of the loss over VaR (0 below it)
  - The influence function IF Chen (2008) and $\mathrm{Var}(\widehat{\mathrm{ES}}_\alpha) \approx \mathrm{Var}(\mathrm{IF}(L))/n$ in "Precision of ES"
- **Tasks**
  1. Show that $\mathrm{Var}(\mathrm{IF}(L)) = \mathrm{Var}(Y)/\alpha^2$ and explain, with A7, why estimating VaR adds no first-order term.
  2. For $L \sim N(0,1)$, $\alpha = 2.5\%$, integrate, with $v = \mathrm{VaR}_\alpha = -z_\alpha$, $\mathbb{E}[Y] = \int_v^\infty (x - v)\varphi(x)\,dx$ and $\mathbb{E}[Y^2]$ in closed form.
  3. Compute the SE of $\widehat{\mathrm{ES}}_{2.5\%}$ for $n = 500$ and $\sigma = 1.2\%$, and check it with 4,000 independent samples (NumPy default_rng(42)); use the fractional historical ES of "Historical Simulation" ($n\alpha = 12.5$) and ddof=1 for the simulated standard deviation.
- **Report**: the two moments, the SE in pp and the simulated standard deviation
- **Interpretation**: why is the S&P 500 value (B2) much larger?
- **Prerequisites and examples**: A1, A7

Back Exercise list

B1 Extended [Solved]: Asymptotic and HAC Intervals for VaR — Task

- **Question**: do the asymptotic formulas agree with the bootstrap of B1?
- **Data**
  - The BET losses of B1; the i.i.d. and HAC standard errors in "Precision of VaR" (kernel density $\hat{f}$, Newey–West lag $\lfloor 4(n/100)^{2/9} \rfloor$)
- **Tasks**
  1. Normal benchmark (from A1): with $f(\mathrm{VaR}_{1\%}) = \varphi(z_{1\%})/1.2$, compute the SE for $n = 500$ and $n = 6\,670$.
  2. For BET, compute $\hat{f}$, the i.i.d. and HAC SEs and the two 95% intervals on both samples.
- **Report**: the four intervals of VaR 1% next to the two bootstrap intervals of B1
- **Interpretation**: why are the BET intervals wider than the Normal benchmark?

B1 Extended [Solved]: Solution — Asymptotic and HAC Intervals for VaR

- **1.** $f = 0.0267/1.2 = 0.0222$: SE 0.20 pp ($n = 500$, width 0.79), 0.055 pp ($n = 6\,670$, width 0.22)
- **2.** BET full: $\hat{f} = 0.0076$, SE 0.16 (i.i.d.), 0.23 (HAC); intervals $[3.81, 4.44]$, $[3.69, 4.57]$
- Last 500 days: $\hat{f} = 0.0163$, SE 0.27 and 0.27
- Interpretation: the BET tail is heavier than the Normal one, so the density at VaR is lower and the quantile less precise; HAC and blocks agree that clustering widens the full-sample interval

MFM_ch7_seminar

B3 Extended [Proposed]: A Bootstrap Interval for the Next-Day FHS VaR — Task

- **Question**: how uncertain is the filtered VaR of B3?
- **Data**: the B3 model and data; residual bootstrap Christoffersen & Gonçalves (2005) with 999 paths, seed 42, function `fhs_bootstrap` of the seminar Quantlet
- **Tasks**
  1. For each path, simulate an AR(1)-GARCH(1,1) series of the sample length with resampled $\hat{z}$ and re-estimate $\theta^* = (c^*, \phi^*, \omega^*, \alpha_1^*, \beta_1^*)$.
  2. Compute μ_{t+1}^* and σ_{t+1}^* by running the **original** Bitcoin returns through θ^*, and $\hat{\xi}_{1\%}^*$ from the path's own residuals.
  3. Report the 90% percentile interval of $\mathrm{VaR}^* = -\mu_{t+1}^* + \sigma_{t+1}^* \hat{\xi}_{1\%}^*$ and the intervals of its two sources.
- **Report**: the histogram of the 999 forecasts with the interval
- **Interpretation**
  - does the interval include the 500-day historical VaR?
  - is that a test?
- **Prerequisites and examples**: B3

B6 [Proposed]: Cornish–Fisher vs Empirical Quantiles — Task

- **Question**: does the Cornish–Fisher correction fix the Normal VaR on real data?
- **Data**
  - BET (compulsory): the B1 data; extension: S&P 500, Bitcoin and gold (daily closes, weekdays for gold) and the BNR EUR/RON reference rate (2005–2026), all to 18 September 2026
  - Moments of the returns X: mean, standard deviation with `ddof=1`, `scipy.stats.skew` and `scipy.stats.kurtosis` (excess, biased)
- **Tasks**
  1. Compute S and K for BET.
  2. Compute VaR 1% and VaR 2.5% by the Normal distribution, Cornish–Fisher (A6) and historical simulation.
  3. Build 95% i.i.d. bootstrap intervals (1 000 resamples) for the two historical VaRs and check whether the Cornish–Fisher values fall inside.
  4. Separately, check whether K exceeds the monotonicity bound of A6.
- **Report**: a small table per asset and a chart of the three methods with the bootstrap intervals
- **Interpretation**: at which tail level does Cornish–Fisher agree better with the historical benchmark, and does the pattern hold for every asset?
- **Prerequisites and examples**: A1, A6, B1

Back Exercise list

B5 Extended [Proposed]: Risk-Parity Weights — Task

- **Question**: which long-only, fully invested weights give each asset of B5 the same share of Normal VaR?
- **Data**: the B5 returns and covariance Σ
- **Tasks**
  1. Solve $w_i(\Sigma w)_i = w_j(\Sigma w)_j$ for all i, j, with $w \geq 0$ and $\sum_i w_i = 1$, numerically.
  2. Compare with the inverse-volatility weights $w_i \propto 1/\sigma_i$.
  3. At the risk-parity weights, compute the Normal VaR 1%, the historical ES 2.5% and the ES shares.
- **Report**: the weights, VaR, ES and shares
- **Interpretation**: does equal VaR risk mean equal ES risk?
- **Prerequisites and examples**: B5

Back Exercise list

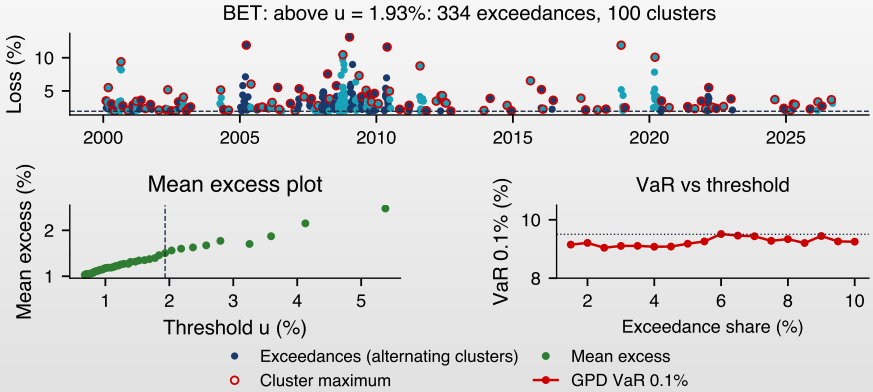
B4 Extended [Solved]: Clustered Exceedances and the Extremal Index — Task

- **Question**: how many independent pieces of tail information do the 334 BET exceedances of B4 contain?
- **Data**
  - Exceedance times $S_1 < \cdots < S_N$ of u (B4), gaps $T_i = S_{i+1} - S_i$
  - The intervals estimator $\hat{\theta}$ Ferro & Segers (2003) in "Clustered Exceedances: the Extremal Index"
- **Tasks**
  1. Compute $\hat{\theta}$ at full precision (round only its printed value) and the target number of clusters $C = \lfloor \hat{\theta} N \rfloor + 1$.
  2. Split the exceedances at the $C - 1$ largest gaps; if the $(C-1)$-th and C-th largest gaps are equal, reduce C until they differ.
  3. Fit the GPD to the cluster maxima and compare $\hat{\xi}/\mathrm{SE}$ with B4.
- **Report**: $\hat{\theta}$, the number of clusters, the declustered fit and a timeline chart
- **Interpretation**: is the i.i.d. SE of B4 credible?

B4 Extended [Solved]: Solution — Clustered Exceedances and the Extremal Index (1/2)

- **1.** $\hat{\theta} = 0.32$, $C = 108$
- **2.** Tie at the cut ($T = 12$ days): 100 clusters of 3.3 exceedances on average; largest cluster 33, 47 isolated exceedances
- **3.** $\hat{\xi} = 0.269$ (SE 0.127) vs 0.209 (SE 0.066); $\hat{\xi}/\mathrm{SE}$ falls from 3.2 to 2.1; at the VaR 1% threshold $\hat{\theta} = 0.40$
- Interpretation: the tail stays heavy, but the i.i.d. SE is not credible; report the declustered fit or filter with GARCH first (conditional EVT, McNeil & Frey (2000))

B4 Extended [Solved]: Solution — Clustered Exceedances and the Extremal Index (2/2)



MFM_ch7_seminar

B7 [Proposed]: CAViaR on BET — Task

- **Question**: does modelling the BET quantile directly beat filtered and historical simulation out of sample?
- **Data**
  - The BET returns of B1 ($n = 6\,670$)
  - SAV CAViaR Engle & Manganelli (2004) and the pinball loss in "Quantile Regression, the Pinball Loss and CAViaR"; $\mathrm{VaR}_{\mathbf{1}} = -\hat{q}_{\mathbf{1}\%}$ of the first 300 returns; routines `caviar_path`, `caviar_fit` (seminar Quantlet): $10^{\mathbf{4}}$ starts in $(0, 1)^{\mathbf{3}}$, seed 42, the best 10 refined
- **Tasks**
  1. Estimate β on 2000–2015 by minimising the summed pinball loss at $\alpha = 1\%$.
  2. Compute one pinball-loss term step by step for the first out-of-sample day and check that the forecast for day t uses only $r_{t-\mathbf{1}}$.
  3. With β fixed, forecast VaR for 2016–2026 and compare exceedances and mean pinball loss with FHS (refitted each January) and rolling 500-day HS.
  4. For 20 February – 30 April 2020 count exceedances and report the peak VaR of each method.
- **Report**: the estimates, a comparison table and a chart with a Coronavirus Disease 2019 (COVID-19) zoom
- **Interpretation**: can the difference be attributed to the model class when CAViaR is fitted once and FHS every year?
- **Prerequisites and examples**: B3; lecture, CAViaR on the S&P 500

Back Exercise list

C2 [Open]: ES Capital, BVB vs S&P 500 — Task (1/2)

- **Question**: which portfolio needs more ES capital per unit invested: Bucharest blue chips or the S&P 500?
- **Data**
  - BVB basket: Banca Transilvania, OMV Petrom (OMV: Österreichische Mineralölverwaltung, the Austrian OMV group), BRD–Groupe Société Générale, Transgaz, Romgaz, Nuclearelectrica, Electrica, Transelectrica (adjusted closes), equal weights, rebalanced daily; keep common observed price dates, retain individual zero returns, and remove days when all eight returns are zero
  - S&P 500 portfolio: SPY, the S&P 500 ETF (adjusted closes); both from 2016-09-20 to 18 September 2026; daily simple returns R_t in % in local currency
- **Tasks**
  1. Construction: compute the one-day historical ES 2.5% of each portfolio (loss $-R_t$) and its tail count.
  2. Horizon: compute the 10-day ES 2.5% by $\sqrt{10}$ scaling, from overlapping compounded losses $100[1 - \prod_{i=1}^{10}(1 + R_{t+i}/100)]$, and by filtered Monte Carlo: first turn the basket simple return into $r_t = 100 \ln(1 + R_t/100)$, fit AR(1)-GARCH(1,1) to this series, simulate 100 000 paths from today's state (seed 42) and convert each simulated log loss to a simple loss before reporting money.
- **Report**: continued on the next slide
- **Prerequisites and examples**: B1, B3; lecture, horizon and Basel sections

Back Exercise list

C2 [Open]: ES Capital, BVB vs S&P 500 — Task (2/2)

- **Tasks (continued)**

3. Stress: find the 250 consecutive days with the largest one-day historical ES 2.5%, plot the rolling ES with the window shaded, and report the stressed ES at 1 day and, by $\sqrt{10}$, at 10 days.
4. Capital: capital per 1 000 000 invested $= 1\,000\,000 \times 1.5 \times \mathrm{ES}_{\mathbf{10}}^{\mathit{stress}}$ (%)/100 (1.5: the FRTB multiplier m_c; $\mathrm{ES}_{\mathbf{10}}^{\mathit{stress}}$: the 10-day stressed ES), for a 10-day and a 20-day liquidity horizon (a further factor $\sqrt{2}$); amounts in RON for BVB and in USD for SPY, ratios comparable.
5. Uncertainty: 95% moving-block bootstrap intervals (20-day blocks, 2 000 resamples) for the one-day ES and for the stressed ES with the window held fixed.

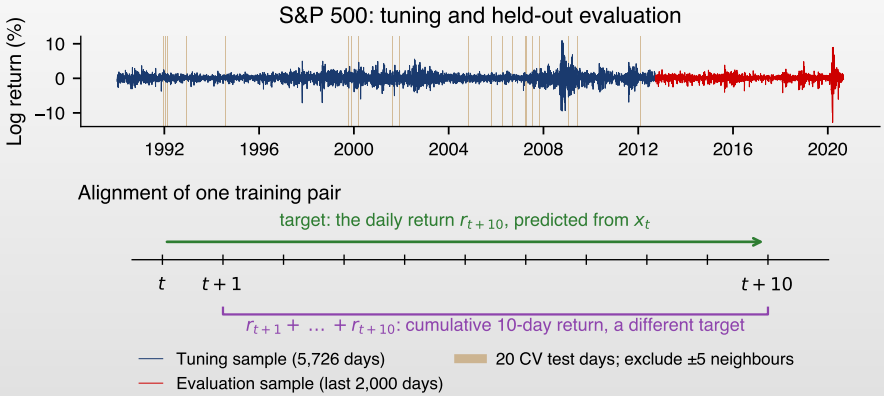
- **Report**: a table of all numbers with intervals and a one-page argument
  - which portfolio needs more capital?
  - how do thin trading and currency change the answer?
- A simplified illustration inspired by FRTB Basel Committee (2019), not the full capital calculation of Basel Committee, MAR33

C3 [Open]: Deep Quantile Regression for VaR — Task (1/2), Minimum Version

- **Question**: does a small neural network beat linear quantile regression for the S&P 500 VaR 1%, as in the lecture case study Chronopoulos, Raftapostolos & Kapetanios (2024)?
- **Data**
  - S&P 500 daily closes, 1990-01-03 – 31 August 2020 (the paper starts in 1985); unscaled daily log returns r_t; tuning: first 5 726 days, evaluation: last 2 000 (from 2012-09-19)
  - One predictor set: GARCH(1,1) σ_t; $\alpha = 1\%$
- **Tasks**
  1. Build training pairs $(x_t, r_{t+\mathbf{10}})$ whose targets all lie in the tuning sample, as in the chart; the target is the daily return $r_{t+\mathbf{10}}$, not the cumulative 10-day return.
  2. Fit a linear quantile regression and one network (one hidden ReLU (Rectified Linear Unit) layer of 5 units, ridge penalty 0.01, Adam optimiser (stochastic gradient descent with adaptive step sizes), early stopping) by the pinball loss. In each chronological cross-validation split, re-estimate GARCH and standardise the predictors on the training part only, and drop the pairs whose targets $r_{t+\mathbf{10}}$ cross the end of the training part; once the configuration is chosen, re-estimate GARCH on the whole tuning sample.
  3. Before evaluating, freeze the learning rate, batch size, epoch cap, chronological validation split, early-stopping patience and five seeds in the submitted protocol. Compare mean pinball losses on the evaluation sample with a Diebold–Mariano test (HAC, 10 lags) Diebold & Mariano (1995), over 5 random seeds.
- **Report**: a table of mean pinball losses, DM statistics and exceedance counts, for each seed

Back Exercise list

C3 [Open]: Deep Quantile Regression for VaR — Data



MFM_ch7_seminar

C3 [Open]: Deep Quantile Regression for VaR — Task (2/2), Extensions

- **Replication extension**: the four predictor sets and the cross-validation grid of the paper (Appendix of the lecture); relative Root Mean Squared Forecast Error (RMSFE) only to reproduce Table 6, since RMSFE is minimised by the conditional mean, not the quantile
- **Rolling extension**: re-estimate once a year on a rolling window and forecast September 2020 – 2026 for the S&P 500 and the BET
- **Joint extension**: two outputs (v_t, e_t) trained on the FZ0 loss Patton, Ziegel & Chen (2019), (v_t: return quantile, e_t: tail mean, as in the lecture) with $e_t = -\exp(a_t)$ and $v_t = e_t + \exp(b_t)$, a_t, b_t unconstrained network outputs, so that $e_t < 0$ and $e_t \leq v_t$
- **Interpretation** (one page): can a gain be attributed to non-linearity when architecture, penalty and tuning change together?

Project Choice and Assessment Criteria

Choose one

- C2: compare ES capital for the specified BVB basket and SPY, following all five tasks; an additional portfolio of at least three assets is an optional extension
- C3: the minimum version, then one extension
- Deliver a Quantlet folder (code, `Metainfo.txt`, charts) and three slides: numbers, interpretation, limitations

Assessment criteria

- Reproducibility: data loaded by code, fixed seeds, every number from a script (30%)
- Correct conventions: sign, level, horizon, return type (20%)
- Uncertainty: tail counts and intervals for every tail estimate (25%)
- Interpretation and limits, including logged AI use and the AI errors actually found (25%)

Acronyms Used in This Material

- **AI**: Artificial Intelligence
- **AR**: AutoRegressive (model)
- **BET**: Bucharest Exchange Trading (index)
- **BNR**: Banca Națională a României — *the National Bank of Romania*
- **BVB**: Bursa de Valori București — *the Bucharest Stock Exchange*
- **CAViaR**: Conditional Autoregressive Value at Risk
- **CF**: Cornish–Fisher (expansion)
- **CI**: Confidence Interval
- **COVID-19**: COronaVIrus Disease 2019
- **DM**: Diebold–Mariano (test)
- **ES**: Expected Shortfall
- **ETF**: Exchange-Traded Fund
- **EVT**: Extreme Value Theory
- **FHS**: Filtered Historical Simulation
- **FRTB**: Fundamental Review of the Trading Book
- **FZ0**: Fissler–Ziegel loss, zero-homogeneous version
- **GARCH**: Generalised AutoRegressive Conditional Heteroskedasticity
- **GPD**: Generalised Pareto Distribution
- **HAC**: Heteroskedasticity and Autocorrelation Consistent (standard errors)
- **HS**: Historical Simulation
- **LGD**: Loss Given Default
- **MAR33**: Basel Framework, chapter MAR33 (internal models approach: capital requirements calculation)
- **OMV**: Österreichische Mineralölverwaltung — *Austrian Mineral Oil Administration (OMV group)*
- **QML**: Quasi-Maximum Likelihood
- **RMSFE**: Root Mean Squared Forecast Error
- **SAV**: Symmetric Absolute Value (CAViaR specification)
- **SE**: Standard Error
- **US**: United States

References (1/2)

- Acerbi, C., Tasche, D. (2002). On the coherence of expected shortfall. *Journal of Banking & Finance*, 26(7), 1487–1503.
- Artzner, P., Delbaen, F., Eber, J.-M., Heath, D. (1999). Coherent measures of risk. *Mathematical Finance*, 9(3), 203–228.
- Barone-Adesi, G., Giannopoulos, K., Vosper, L. (1999). VaR without correlations for portfolios of derivative securities. *Journal of Futures Markets*, 19(5), 583–602.
- Basel Committee on Banking Supervision (2019). Minimum capital requirements for market risk. Bank for International Settlements.
- Chen, S.X. (2008). Nonparametric estimation of expected shortfall. *Journal of Financial Econometrics*, 6(1), 87–107.
- Christoffersen, P., Gonçalves, S. (2005). Estimation risk in financial risk management. *Journal of Risk*, 7(3), 1–28.
- Chronopoulos, I., Raftapostolos, A., Kapetanios, G. (2024). Forecasting Value-at-Risk using deep neural network quantile regression. *Journal of Financial Econometrics*, 22(3), 636–669.
- Cornish, E.A., Fisher, R.A. (1938). Moments and cumulants in the specification of distributions. *Revue de l'Institut International de Statistique*, 5(4), 307–320.
- Diebold, F.X., Mariano, R.S. (1995). Comparing predictive accuracy. *Journal of Business & Economic Statistics*, 13(3), 253–263.
- Efron, B. (1979). Bootstrap methods: another look at the jackknife. *Annals of Statistics*, 7(1), 1–26.
- Engle, R.F., Manganelli, S. (2004). CAViaR: conditional autoregressive value at risk by regression quantiles. *Journal of Business & Economic Statistics*, 22(4), 367–381.
- Ferro, C.A.T., Segers, J. (2003). Inference for clusters of extreme values. *Journal of the Royal Statistical Society: Series B*, 65(2), 545–556.

References (2/2)

- Künsch, H.R. (1989). The jackknife and the bootstrap for general stationary observations. *Annals of Statistics*, 17(3), 1217–1241.
- McNeil, A.J., Frey, R. (2000). Estimation of tail-related risk measures for heteroscedastic financial time series: an extreme value approach. *Journal of Empirical Finance*, 7(3–4), 271–300.
- McNeil, A.J., Frey, R., Embrechts, P. (2015). Quantitative Risk Management: Concepts, Techniques and Tools, revised ed. Princeton University Press.
- Patton, A.J., Ziegel, J.F., Chen, R. (2019). Dynamic semiparametric models for expected shortfall (and Value-at-Risk). *Journal of Econometrics*, 211(2), 388–413.
- Pele, D.T., Mazurencu-Marinescu-Pele, M. (2026). Finite-sample precision limits for expected shortfall forecast comparisons. *Mathematics*, 14(13), 2316.
- Rockafellar, R.T., Uryasev, S. (2000). Optimization of conditional value-at-risk. *Journal of Risk*, 2(3), 21–41.
- Tasche, D. (2008). Capital allocation to business units and sub-portfolios: the Euler principle. arXiv:0708.2542.