








Modelling Financial Markets

Seminar 4: Portfolio Optimisation — Hands-on



Daniel Traian PELE
Antoaneta Amza

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 Do We Choose Weights and Compare Portfolios?

- Question of the seminar
  - **can an optimised portfolio beat equal weights (1/N) out of sample, and how would we know that it does?**
- 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: chapter4_seminar_notebook.ipynb (Colab); prices: data/market
- Seminar: Prof. Daniel Traian Pele, danpele@ase.ro; Antoaneta Amza, PhD candidate, antoaneta.amza@csie.ase.ro

Which Exposures Do the Funds Buy?

ETF (exchange-traded fund): a fund traded on an exchange.

SPY	S&P 500 (Standard & Poor's 500)
EFA	developed markets, excluding US/Canada
EEM	emerging markets
TLT / IEF	US Treasuries: long / 7–10 years
LQD / HYG	corporate: investment-grade / high-yield
GLD	gold

XLB	materials
XLE	energy
XLF	financials
XLI	industrials
XLK	technology
XLP	consumer staples
XLU	utilities
XLV	health care
XLY	consumer discretionary

USD (US dollar): US (United States) dollar; RON: Romanian leu; BTC: Bitcoin. OOS: out-of-sample, i.e. outside the estimation sample.

How Do We Read the Bucharest Stock Exchange Symbols?

BVB (Bursa de Valori București): Bucharest Stock Exchange; BET (Bucharest Exchange Trading): the main market index; BET-TR (total return): its dividend-reinvestment version.

TLV	Banca Transilvania	SNP	OMV Petrom
BRD	BRD Groupe Société Générale	TGN	Transgaz
SNG	Romgaz	SNN	Nuclearelectrica
EL	Electrica	TEL	Transelectrica

Data Used

Universe	Assets (price used)	Currency	Monthly returns	Months
Sectors	XLB, XLE, XLF, XLI, XLK, XLP, XLU, XLV, XLY (adjusted close)	USD	Jan 1999 – Jul 2026	331
Multi-asset	SPY, EFA, EEM, TLT, IEF, LQD, HYG, GLD (adjusted close)	USD	May 2007 – Jul 2026	231
Combined	9 sectors + the 7 multi-asset ETFs other than SPY	USD	May 2007 – Jul 2026	231
Bitcoin	Bitcoin (BTC/USD, close)	USD	Aug 2010 – Jul 2026	B8 only
BVB	TLV, SNP, BRD, TGN, SNG, SNN, EL, TEL; the BET-TR index (close)	RON	Oct 2014 – Aug 2026	C2 only

- Daily prices of the course data, one series per asset (field used: close / adjusted close)
- Align the *prices* of a universe on common trading days, keep the last price of each month, then take simple returns in decimals ($0.01 = 1\%$)
- Risk-free rate: one-month Treasury bill, Kenneth French Data Library, monthly decimals aligned on the month-end date; excess return $r_t^e = r_t - r_{f,t}$; USD/RON: National Bank of Romania (BNR) reference rate (C2)

Setup: Run and Verify

```
# notebook section "Setup" defines the helpers
r, rex, rf = us_monthly(SECTORS)
# r: monthly simple returns (T x 9)
# rex = r - rf: excess returns; rf: T-bill
print(r.shape, r.index[0], r.index[-1])
print(r.iloc[0][['XLB', 'XLK', 'XLU']], rf.iloc
    [0])
ret, to, W = backtest(r, rex)
# ret: OOS excess returns of each rule
# to: monthly turnover; W: weights per rule
```

- Expected output: 331 months, first return 1999-01-31, last 2026-07-31
- First row (January 1999): XLB (materials) −3.70%, XLK (technology) 15.90%, XLU (utilities) −2.48%; T-bill 0.35% a month
- `backtest` uses months $t-60,\ldots,t-1$ for the weights of month t: 271 out-of-sample months, 2004-01 to 2026-07
- Other helpers: `sr_diff_hac(r1, r2)` returns the annualised difference, its heteroskedasticity and autocorrelation consistent (HAC) standard error and p; `sr_diff_boot(r1, r2, block)` the studentized bootstrap test; `lw_cc(X)` the shrunk covariance and δ
- If a count or a first value differs, stop and check the symbols, the price type (adjusted or close) and the date alignment

Portfolio Return and Variance

- Details: Chapter 4
- **Portfolio return**: the weighted sum of the asset returns Markowitz (1952)

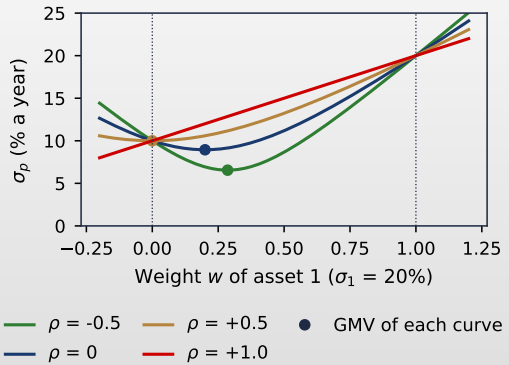
$$r_p = \mathrm{w}^\top \mathrm{r} = \sum_{i=\mathbf{1}}^{N} w_i r_i, \qquad \mathrm{w}^\top 1 = 1$$

  - N: number of assets; w_i: share of wealth in asset i ($w_i < 0$: a short position)
  - 1: vector of ones, so $\mathrm{w}^\top 1 = 1$ means fully invested
- **Mean and variance** of the portfolio

$$\mu_p = \mathrm{w}^\top \boldsymbol{\mu}, \qquad \sigma_p^{\mathbf{2}} = \mathrm{w}^\top \boldsymbol{\Sigma} \mathrm{w} = \sum_i \sum_j w_i w_j \sigma_{ij}$$

  - $\boldsymbol{\mu}$: vector of expected excess returns; $\boldsymbol{\Sigma}$: covariance matrix, $\sigma_{ij} = \rho_{ij}\sigma_i\sigma_j$

GMV (global minimum variance): the lowest portfolio variance.



Diversification with Two Assets

- With two assets and $w = w_1$, $1 - w = w_2$, the variance has three terms

$$\sigma_p^2 = w^2\sigma_1^2 + (1-w)^2\sigma_2^2 + 2w(1-w)\rho\sigma_1\sigma_2$$

  - the third term is the only one that depends on the correlation ρ
  - Long-only weights and $\rho = 1$: $\sigma_p = w\sigma_1 + (1-w)\sigma_2$, a straight line, no diversification
  - Long-only weights and $\rho < 1$: σ_p lies below that line; the lower ρ, the deeper the curve
- Chart on the previous slide: $\sigma_1 = 20\%$, $\sigma_2 = 10\%$, four correlations
  - $\rho = 0$: the minimum is at $w = 0.2$, $\sigma_p = 8.94\%$, below both volatilities
  - $\rho = -0.5$: the minimum falls to 6.55% at $w = 0.29$
  - $\rho = +0.5 = \sigma_2/\sigma_1$: the minimum is the less volatile asset alone ($w = 0$)
- Worked example: $\rho = 0$, w $= (0.5, 0.5)$: $\sigma_p^2 = 0.25(0.04) + 0.25(0.01) = 0.0125$, $\sigma_p = 11.2\%$
- **What do you think**: with $\sigma_1 = 20\%$, $\sigma_2 = 10\%$ and $\rho = 0.8$, what weight does asset 1 get in the unconstrained minimum-variance portfolio?
  - **Answer**: $w = (\sigma_2^2 - \rho\sigma_1\sigma_2)/(\sigma_1^2 + \sigma_2^2 - 2\rho\sigma_1\sigma_2) = -0.006/0.018 = -1/3$, a short position, because $\rho > \sigma_2/\sigma_1$

The Global Minimum-Variance Portfolio

- **GMV** (global minimum variance): the fully invested weights with the smallest variance

$$\min_{w} \; w^\top \Sigma w \quad \text{subject to} \quad w^\top 1 = 1$$

- Solution (derived in A1 with a Lagrange multiplier)

$$w_{GMV} = \frac{\Sigma^{-1}1}{1^\top \Sigma^{-1}1}, \qquad \sigma^2_{GMV} = \frac{1}{1^\top \Sigma^{-1}1}$$

  - Σ^{-1}: inverse covariance matrix; the denominator rescales the weights so that they sum to 1
  - needs no expected returns: only Σ is estimated

- Example: the same volatilities, $\sigma_1 = 20\%$ and $\sigma_2 = 10\%$, with correlation $\rho = -0.5$, so $\sigma_{12} = -0.01$
  - $\Sigma^{-1}1 = (66.7, 166.7)$, sum 233.3
  - $w_{GMV} = (0.29, 0.71)$; $\sigma_{GMV} = 1/\sqrt{233.3} = 6.55\%$

Sharpe Ratio, Tangency Portfolio and 1/N

- **Sharpe ratio**: mean excess return per unit of volatility

$$SR_p = \frac{\mu_p}{\sigma_p} = \frac{\mathrm{w}^\top \boldsymbol{\mu}}{\sqrt{\mathrm{w}^\top \boldsymbol{\Sigma} \mathrm{w}}}$$

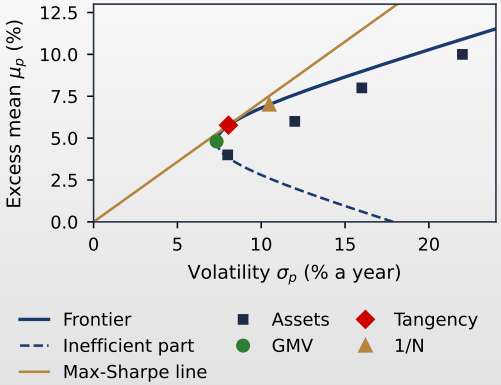
- **MV** (mean–variance tangency portfolio): the weights with the highest SR_p, provided the denominator below is positive (Chapter 3)

$$\mathrm{w}_{\mathbf{MV}} = \frac{\boldsymbol{\Sigma}^{-1}\boldsymbol{\mu}}{1^\top \boldsymbol{\Sigma}^{-1}\boldsymbol{\mu}}$$

  - **MV-LO**: the same maximum with $0 \leq w_i \leq 1$ (long-only), solved numerically

- **1/N**: $w_i = 1/N$, nothing estimated DeMiguel, Garlappi & Uppal (2009)

- Chart: four simulated assets; the line from the origin through the tangency point has slope $SR = 0.72$; GMV 0.65, 1/N 0.67

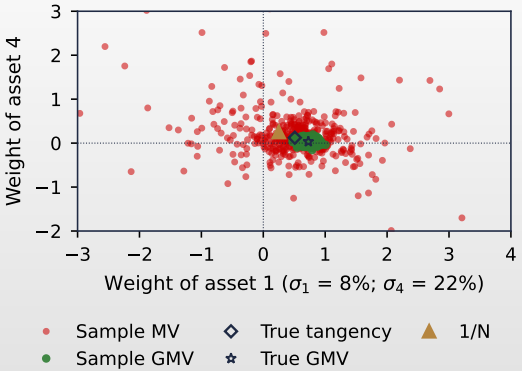


Estimation Error in the Weights

- In practice $\boldsymbol{\mu}$ and $\boldsymbol{\Sigma}$ are replaced by sample estimates from T months

$$\hat{\boldsymbol{\mu}} = \frac{1}{T}\sum_{t=1}^{T} \mathrm{r}_t, \qquad \hat{\boldsymbol{\Sigma}} = \frac{1}{T-1}\sum_{t=1}^{T} (\mathrm{r}_t - \hat{\boldsymbol{\mu}})(\mathrm{r}_t - \hat{\boldsymbol{\mu}})^{\top}$$

- Chart: 500 samples of $T = 60$ months from the same four assets
  - sample MV weights scatter widely: interquartile range 0.66 for asset 1; 38% of the samples have some $|w_i| > 1$
  - sample GMV weights stay close to the truth (range 0.12): no $\hat{\boldsymbol{\mu}}$ in the formula
- Why 1/N is a hard benchmark
  - $\hat{\boldsymbol{\mu}}$ has large errors and $\hat{\boldsymbol{\Sigma}}^{-1}$ amplifies them into extreme weights



Risk Contributions

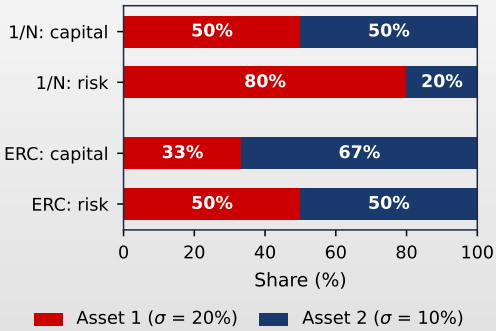
- **Euler decomposition**: the portfolio volatility is the sum of the contributions of the assets

$$\sigma_p = \sum_i RC_i, \qquad RC_i = w_i \frac{\partial \sigma_p}{\partial w_i} = \frac{w_i (\Sigma \mathrm{w})_i}{\sigma_p}$$

  - $(\Sigma \mathrm{w})_i = \sum_j \sigma_{ij} w_j$: covariance of asset i with the portfolio
  - RC_i: contribution of asset i to the volatility; risk share RC_i/σ_p, the shares add up to 100%

- Example: $\sigma_1 = 20\%$, $\sigma_2 = 10\%$, $\rho = 0$, $\mathrm{w} = (0.5, 0.5)$
  - $\Sigma \mathrm{w} = (0.020, 0.005)$, $\sigma_p^2 = 0.0125$
  - shares $0.5 \times 0.020/0.0125 = 80\%$ and 20%: half of the money, four fifths of the risk

ERC (equal risk contribution): equal contributions to risk.



Equal Risk Contribution

- **ERC** (equal risk contribution, risk parity) Maillard, Roncalli & Teïletche (2010): weights with equal contributions

$$RC_1 = RC_2 = \cdots = RC_N = \sigma_p / N$$

  - uses only $\boldsymbol{\Sigma}$; in general solved numerically (helper `w_erc`)
  - with two uncorrelated assets: $w_i \propto 1/\sigma_i$
- Example ($\sigma_1 = 20\%$, $\sigma_2 = 10\%$, $\rho = 0$)
  - $w \propto (1/0.2,\ 1/0.1) = (5, 10)$, so w $= (1/3, 2/3)$
  - risk shares 50% and 50%; $\sigma_p = 9.43\%$ (bottom bars of the previous chart)
- Capital weights and risk shares are different things
  - a portfolio balanced in money can be concentrated in risk (A3)
  - ERC sits between 1/N (equal money) and GMV (lowest risk)

The Rolling Backtest

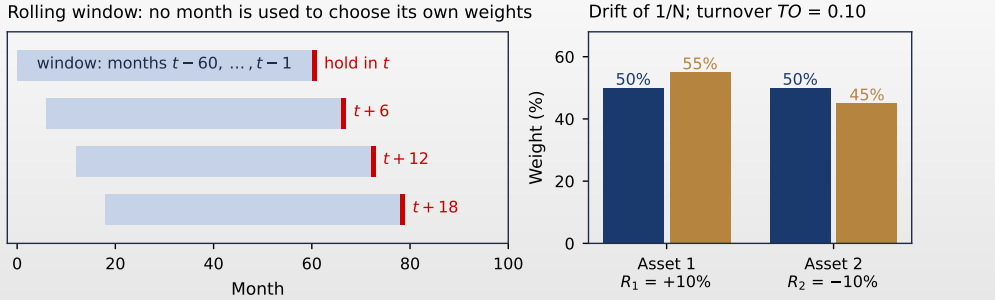
- **Out-of-sample** (OOS) backtest: each month, choose weights with past data only, then record the next return
  - rolling window: the weights held in month t use only months $t-60, \ldots, t-1$
  - any estimate that uses month t is **look-ahead**: it uses information not yet available
- **Drift**: during the month the weights move with the returns

$$w_i^+ = \frac{w_i(1+R_i)}{\sum_j w_j(1+R_j)}$$

  - R_i: simple return of asset i in the month; w_i^+: its weight at the end of the month
- **Turnover**: the share of wealth traded to return to the new target weights $w_{i,\text{new}}$

$$TO = \sum_i |w_{i,\text{new}} - w_i^+|$$

Windows, Drift and Turnover in a Chart



- Left: every 60-month window ends one month before the month it is used for
- Right: 1/N with returns $+10\%$ and -10% drifts to $\mathrm{w}^+ = (0.55, 0.45)$
  - back to 1/N: $TO = |0.5 - 0.55| + |0.5 - 0.45| = 0.10$, a tenth of wealth traded

Performance Measures

- From the monthly excess returns $r_1^e, \ldots, r_T^e$ of a strategy (sample mean $\bar{r}^e$, standard deviation s)

$$\hat{\mu}_{\text{ann}} = 12\,\bar{r}^e, \qquad \hat{\sigma}_{\text{ann}} = \sqrt{12}\,s, \qquad \widehat{SR}_{\text{ann}} = \sqrt{12}\,\frac{\bar{r}^e}{s}$$

  - $\sqrt{12}$ for monthly data ($\sqrt{252}$ is for daily data): the mean grows with 12, the volatility with $\sqrt{12}$
- **Maximum drawdown**: the largest fall of wealth W_t from a previous peak

$$MDD = \min_t \Big(\frac{W_t}{\max_{s \le t} W_s} - 1 \Big), \qquad W_t = \prod_{s \le t} (1 + r_s)$$

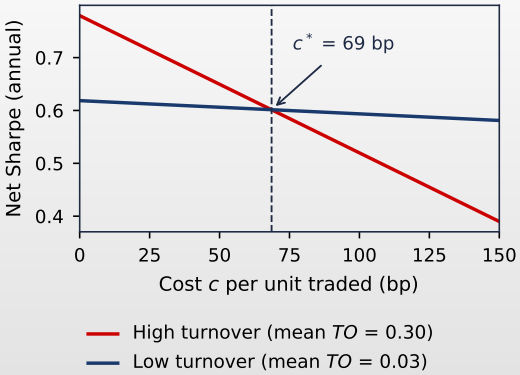
- **Gross exposure** $\sum_i |w_i|$: equals 1 for long-only weights, larger with short positions

Trading Costs and the Break-Even Cost

- **Net return**: the gross return minus the cost of the trades

$$r_t^{\text{net}} = r_t - c\, TO_t$$

  - c: cost per unit traded; 1 basis point (bp) $= 10^{-4}$, so 20 bp is $c = 0.002$
  - $TO = 0.10$ at 20 bp costs $0.002 \times 0.10 = 2$ bp of wealth
- **Break-even cost** c^*: the cost at which two net Sharpe ratios are equal
- Chart: two simulated strategies, gross Sharpe 0.78 and 0.62
  - the one with 10 times more turnover loses its edge at $c^* = 69$ bp

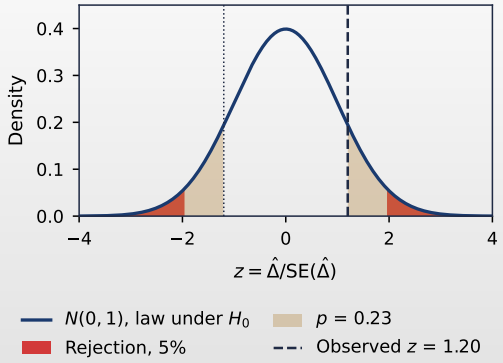


A Two-Sided z-Test

- $\hat{\Delta} = \widehat{SR}_1 - \widehat{SR}_2$: difference of two Sharpe ratios on the same months
- H_0: $SR_1 = SR_2$; statistic and p-value

$$z = \frac{\hat{\Delta}}{\mathrm{SE}(\hat{\Delta})}, \qquad p = 2\,[1 - \Phi(|z|)]$$

  - $\mathrm{SE}(\hat{\Delta})$: standard error of $\hat{\Delta}$ (next slides)
  - Φ: distribution function of $N(0,1)$, the law of z under H_0 in large samples
- Decision at level $\alpha = 5\%$: reject H_0 if $|z| > 1.96$, equivalently if $p < 0.05$
- Chart: $z = 1.20$ gives $p = 0.23$, not rejected
  - not rejecting does not prove that the two Sharpe ratios are equal



The Delta Method

- **Delta method**: the standard error of a smooth function of estimates, by a linear approximation

$$\sqrt{T}(\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}) \to N(0, \boldsymbol{\Psi}) \;\Rightarrow\; \sqrt{T}\big(f(\hat{\boldsymbol{\theta}}) - f(\boldsymbol{\theta})\big) \to N(0,\; \nabla f^{\top} \boldsymbol{\Psi} \nabla f)$$

  - $\boldsymbol{\theta}$: the means and variances of the two strategies; $\hat{\boldsymbol{\theta}}$: their sample estimates
  - $\boldsymbol{\Psi}$: covariance matrix of $\sqrt{T}(\hat{\boldsymbol{\theta}} - \boldsymbol{\theta})$; T: number of months
  - $f(\boldsymbol{\theta}) = \mu_1/\sigma_1 - \mu_2/\sigma_2$, the Sharpe difference; ∇f: its gradient (vector of partial derivatives)
- Result: $\mathrm{SE}(\hat{\Delta}) = \sqrt{\nabla f^{\top} \hat{\boldsymbol{\Psi}} \nabla f / T}$, with ∇f evaluated at $\hat{\boldsymbol{\theta}}$
- The two estimates $\widehat{SR}_1$, $\widehat{SR}_2$ are correlated: the test must use the pair, not two separate standard errors

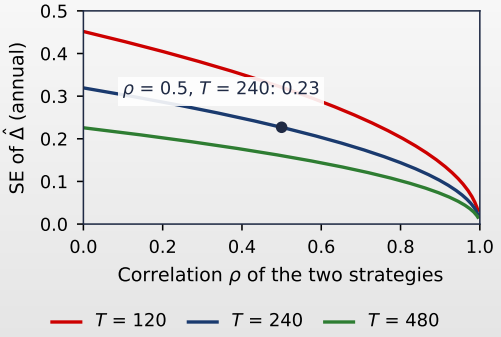
The Variance of a Sharpe Difference

- i.i.d. (independent and identically distributed) returns with the Normal distribution Jobson & Korkie (1981), corrected in Ledoit & Wolf (2008)

$$T\,\mathrm{Var}(\hat{\Delta}) \to 2 - 2\rho + \tfrac{1}{2}\left(SR_1^2 + SR_2^2 - 2SR_1 SR_2 \rho^2\right)$$

  - ρ: correlation of the two strategies' returns
  - SR_1, SR_2: monthly Sharpe ratios
  - a high ρ makes the difference precise

- Example: $SR_1 = SR_2 = 0.2$ a month, $\rho = 0.5$, $T = 240$
  - bracket 1.03; standard error (SE) $\sqrt{1.03/240} = 0.066$ a month, 0.23 a year
  - an estimated gap must exceed about $1.96 \times 0.227 \approx 0.45$ a year for two-sided significance at 5%
  - for 80% power, the true gap must be about $(1.96 + 0.84) \times 0.227 \approx 0.64$ a year



Which Moments Enter the Sharpe Gradient?

- The HAC test (next slide) describes each strategy by its mean μ_i and its uncentred second moment $\gamma_i = \mathbb{E}[r_i^2] = \sigma_i^2 + \mu_i^2$, not by its variance.
- Write one Sharpe ratio as $g(\mu, \gamma) = \mu / \sqrt{\gamma - \mu^2}$. The chain rule gives

$$g_\mu = \gamma / (\gamma - \mu^2)^{3/2}, \qquad g_\gamma = -\mu / [2(\gamma - \mu^2)^{3/2}]$$

- For the difference, put the derivatives in the order $(\mu_1, \mu_2, \gamma_1, \gamma_2)$ and reverse the signs for strategy 2.
- Toy check: $\mu = 0.01$, $\gamma = 0.0005$, hence $\sigma = 0.02$ and $SR = 0.5$. The derivatives are 62.5 and -625, whereas holding variance fixed gives $1/\sigma = 50$.
- Use the gradient and covariance for the same parameterisation; mixing them changes the standard error.

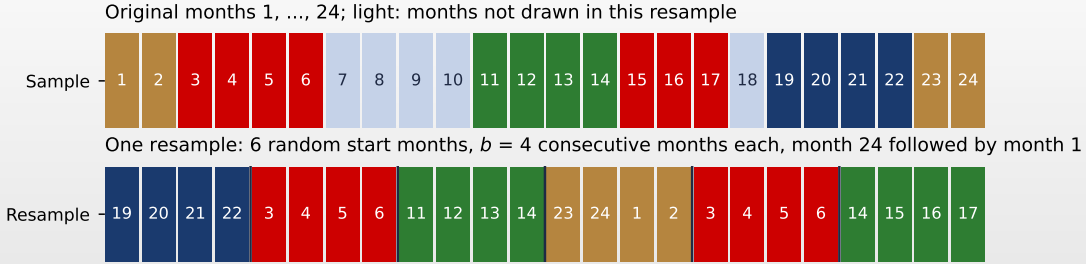
The HAC Test

- Monthly returns have heavy tails and volatility clustering: the i.i.d. formula can understate or overstate the uncertainty
  - GMV − 1/N on sectors: i.i.d. SE 0.167, HAC SE 0.159 a year
- Conditions: a finite fourth moment; stationarity and weak dependence, so that (r_t, r_t^2) obeys a central limit theorem
- **HAC test** Ledoit & Wolf (2008): the delta method with $\boldsymbol{\Psi}$ estimated by a HAC (heteroskedasticity and autocorrelation consistent) covariance of $\mathrm{y}_t = (r_{1t} - \hat{\mu}_1,\ r_{2t} - \hat{\mu}_2,\ r_{1t}^2 - \hat{\gamma}_1,\ r_{2t}^2 - \hat{\gamma}_2)^\top$

$$\hat{\boldsymbol{\Psi}} = \hat{\boldsymbol{\Gamma}}_0 + \sum_{j=1}^{L} \Big(1 - \frac{j}{L+1}\Big)\big(\hat{\boldsymbol{\Gamma}}_j + \hat{\boldsymbol{\Gamma}}_j^\top\big), \qquad \hat{\boldsymbol{\Gamma}}_j = \frac{1}{T}\sum_{t=j+1}^{T} \mathrm{y}_t \mathrm{y}_{t-j}^\top$$

  - $\hat{\gamma}_i$: mean of r_{it}^2; $\hat{\boldsymbol{\Gamma}}_j$: autocovariance matrix at lag j
  - Bartlett weights $1 - j/(L+1)$ decline linearly; $L = \lfloor 4(T/100)^{2/9} \rfloor$ lags, 4 for $T = 271$
- Then $z = \hat{\Delta}/\mathrm{SE}$ and $p = 2[1 - \Phi(|z|)]$, as in A Two-Sided z-Test; helper `sr_diff_hac`

The Circular Block Bootstrap



- A resample of T months is built from blocks of b consecutive months with random starts Künsch (1989)
  - blocks keep the dependence between neighbouring months; "circular": month T is followed by month 1
  - some months appear twice, others not at all; the **pair** $(r_{\mathbf{1}t}, r_{\mathbf{2}t})$ is always drawn together

The Studentized Bootstrap Test

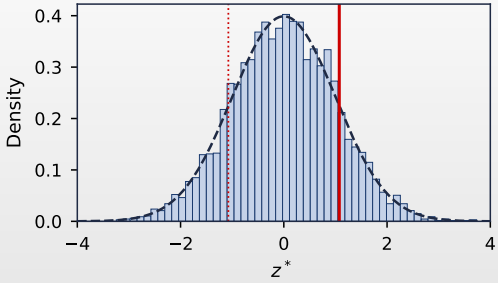
- On each of M resamples compute $\hat{\Delta}^*_m$, its HAC standard error SE^*_m and the studentized statistic Ledoit & Wolf (2008)

$$z^*_m = \frac{\hat{\Delta}^*_m - \hat{\Delta}}{\mathrm{SE}^*_m}$$

  - centred at $\hat{\Delta}$, so z^*_m mimics the law of z under H_0
- p-value: the share of $|z^*_m|$ at least as large as the observed $|z|$

$$p = \frac{1 + \#\{m : |z^*_m| \geq |z|\}}{M + 1}$$

- 95% interval $\hat{\Delta} \pm q^* \, \mathrm{SE}$, q^*: 95th percentile of the $|z^*_m|$
- Chart: simulated pair, block 1, $M = 4\,999$: $z = 1.07$, bootstrap $p = 0.28$; helper `sr_diff_boot`
- It can improve the finite-sample approximation; it does not repair a missing fourth moment (nor does it for the variance ratio of B3)

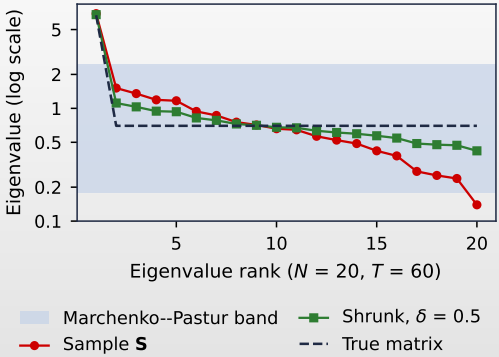


Covariance Shrinkage (B3, after the Lecture)

- **Ledoit–Wolf** Ledoit & Wolf (2004): a weighted average of the sample matrix and a structured target

$$\hat{\boldsymbol{\Sigma}}_{\mathbf{LW}} = \delta \mathrm{F} + (1-\delta)\mathrm{S}, \qquad f_{ij} = \bar{r}\sqrt{s_{ii}s_{jj}},\ f_{ii} = s_{ii}$$

  - S: sample covariance, elements s_{ij}; $\bar{r}$: average of the $N(N-1)$ sample correlations
  - $\delta \in [0,1]$: estimated weight of the target, large when S is noisy relative to $\|\mathrm{S}-\mathrm{F}\|$
- **Condition number** $\lambda_{\max}/\lambda_{\min}$ (largest over smallest eigenvalue): large means S^{-1} is unstable
- Chart, $N = 20$, $T = 60$: the sample spreads the eigenvalues (condition number 49.6, true 9.6); $\delta = 0.5$ gives 16.1



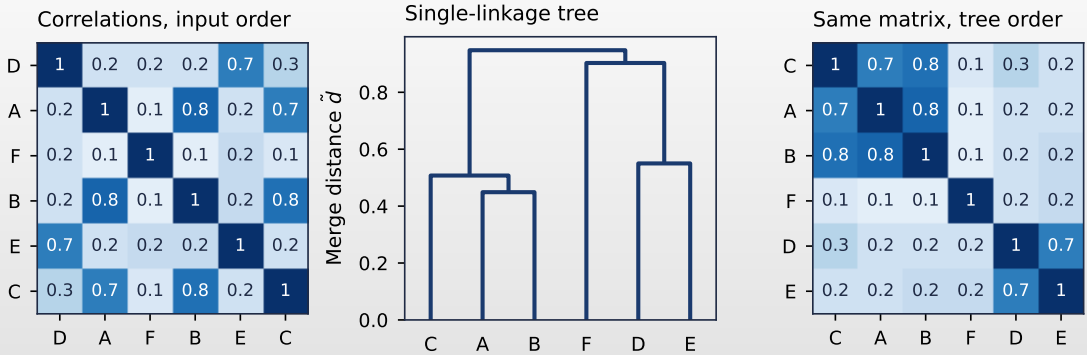
Hierarchical Risk Parity (B4, after the Lecture)

- **HRP** (Hierarchical Risk Parity) López de Prado (2016): allocate along a tree of similar assets, without inverting $\boldsymbol{\Sigma}$
- Step 1, distances: $d_{ij} = \sqrt{(1-\rho_{ij})/2} \in [0,1]$, then $\tilde{d}_{ij} = \sqrt{\sum_k (d_{ki} - d_{kj})^2}$ between the columns of $\mathrm{D} = (d_{ij})$
- Step 2, tree: single linkage merges the two closest clusters; the distance between clusters is the smallest $\tilde{d}$ between their members
- Step 3, order: list the assets as the leaves of the tree (similar assets next to each other)
- Step 4, recursive bisection: split the ordered list in two halves C_1, C_2

$$V_k = \mathrm{v}_k^\top \boldsymbol{\Sigma}_k \mathrm{v}_k, \quad v_{k,i} \propto 1/\sigma_i^2, \qquad \alpha = 1 - \frac{V_1}{V_1 + V_2}$$

  - v_k: inverse-variance weights inside half k; V_k: variance of that half; α to C_1, $1-\alpha$ to C_2
  - repeat inside each half until single assets remain

The Tree in a Chart



- Six simulated assets: two groups of related assets (A, B, C and D, E) and one loosely related asset (F)
- The tree merges A–B first ($\tilde{d} = 0.45$); in tree order the matrix shows the two blocks on the diagonal

Notation Used in Every Exercise

- N assets, T months; w: weights, $\mathrm{w}^\top 1 = 1$; $\boldsymbol{\mu}$, $\boldsymbol{\Sigma}$: mean vector and covariance matrix of monthly excess returns
- $R_{i,t}$: simple return of asset i in month t; $r^e_t = r_t - r_{f,t}$: excess return (Data Used)
- In this seminar, r and R denote the same simple return; r^e and R^e are their excess-return versions. Lower case does not indicate a log return here.
- Rules and the slides where they are defined
  - 1/N: $w_i = 1/N$; GMV: The Global Minimum-Variance Portfolio; MV (tangency) and MV-LO (long-only): Sharpe Ratio, Tangency Portfolio and 1/N
  - ERC: Equal Risk Contribution (risk contributions RC_i: Risk Contributions); HRP: Hierarchical Risk Parity
- Performance and trading
  - Sharpe ratio SR: Sharpe Ratio, Tangency Portfolio and 1/N, annualised in Performance Measures; difference $\hat{\Delta} = \widehat{SR}_1 - \widehat{SR}_2$: A Two-Sided z-Test
  - drifted weights $w^+_{i,t}$ and turnover TO: The Rolling Backtest; net return, cost c and basis points (bp): Trading Costs and the Break-Even Cost

Exercise Map

Exercise	Question	Type	Model
A1	where is the minimum-variance mix of SPY and TLT?	Solved	–
A2, A3	who carries the risk: ERC versus 60/40	Solved, Proposed	A3: A2
B1	do Markowitz rules beat 1/N on sectors?	Solved	–
A4, A5	is a Sharpe difference significant?	Solved, Proposed	A5: A4
B2	at which cost does MV-LO stop paying?	Proposed	B1
C1	what is wrong in an AI (artificial intelligence) backtest?	Proposed	B1, A4
After the lecture	B3 [Solved], B4 [Solved]: covariance shrinkage, ERC and HRP	Solved: reproduce	B1
C2	a BVB portfolio in RON (take-home project)	Proposed	B1, B4, A4
Extension packs	A3 Extended, A6, A7, A8, A9, B3 Extended, B5–B9, C3 (optional)	Solved or Proposed	marked on each task

- **[Solved]**: full solution in the slides and in the notebook, a model to copy; **[Proposed]**: you solve it, the solution is discussed in class
- Part A: derivations on paper; Part B: backtests and inference on real data, each ending with an interpretation question; Part C: open questions and the critique of an AI (Artificial Intelligence) answer

A1 [Solved]: The GMV Portfolio — Task

- **Question**: which mix of SPY (S&P 500 ETF) and TLT (20+ year US Treasury ETF) has the lowest variance, and what does it earn?
- **Assumptions**: N risky assets, known positive definite $\mathbf{\Sigma}$, short sales allowed, $\mathrm{w}^\top 1 = 1$
- **Inputs** (monthly total returns August 2002 – July 2026, $T = 288$, annualised; the numbers below suffice)

	mean	volatility	variance
SPY	11.73%	14.72%	0.02166
TLT	4.32%	13.49%	0.01820
correlation $\rho = -0.100$, covariance $\sigma_{\mathbf{12}} = -0.00199$			

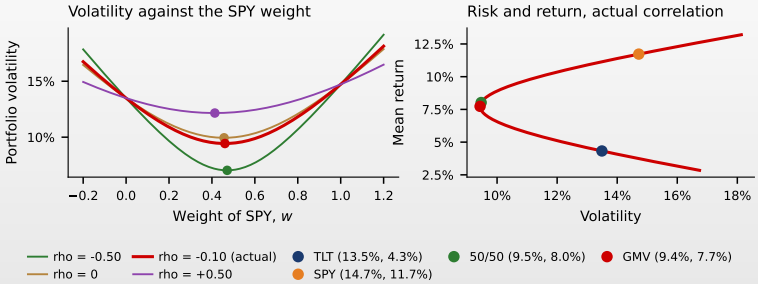
1. Write the Lagrangian of $\min_{\mathrm{w}} \mathrm{w}^\top \mathbf{\Sigma} \mathrm{w}$ s.t. $\mathrm{w}^\top 1 = 1$ and derive $\mathrm{w}_{\mathrm{GMV}} = \mathbf{\Sigma}^{-1} 1 / A$, $A = 1^\top \mathbf{\Sigma}^{-1} 1$.
2. Show that the minimum variance is $1/A$.
3. Invert the 2×2 matrix and obtain $w^* = (\sigma_2^2 - \sigma_{12}) / (\sigma_1^2 + \sigma_2^2 - 2\sigma_{12})$.
4. Compute A, w^*, σ_{GMV} and the mean of the mix; **report** the four numbers and one sentence on why σ_{GMV} is below both volatilities.

Back Exercise list

A1 [Solved]: Derivation

- Step 1: $\mathcal{L} = \mathrm{w}^\top \boldsymbol{\Sigma} \mathrm{w} - \lambda(\mathrm{w}^\top 1 - 1)$; first-order condition $2\boldsymbol{\Sigma}\mathrm{w} - \lambda 1 = 0$, so $\mathrm{w} = \frac{\lambda}{2}\boldsymbol{\Sigma}^{-1}1$
- Imposing $1^\top \mathrm{w} = 1$: $\frac{\lambda}{2}A = 1$, hence $\mathrm{w}_{\mathrm{GMV}} = \boldsymbol{\Sigma}^{-1}1/A$; the objective is strictly convex, so the minimum is unique
- Step 2: $\sigma^2_{\mathrm{GMV}} = 1^\top \boldsymbol{\Sigma}^{-1}\boldsymbol{\Sigma}\boldsymbol{\Sigma}^{-1}1/A^2 = A/A^2 = 1/A$
- Step 3: $\boldsymbol{\Sigma}^{-1} = \frac{1}{\det \boldsymbol{\Sigma}} \begin{pmatrix} \sigma_2^2 & -\sigma_{12} \\ -\sigma_{12} & \sigma_1^2 \end{pmatrix}$, so $\boldsymbol{\Sigma}^{-1}1 = \frac{1}{\det \boldsymbol{\Sigma}}(\sigma_2^2 - \sigma_{12},\ \sigma_1^2 - \sigma_{12})^\top$
- Dividing the first element by the sum of both: $w^* = (\sigma_2^2 - \sigma_{12})/(\sigma_1^2 + \sigma_2^2 - 2\sigma_{12})$; $\det \boldsymbol{\Sigma}$ cancels

A1 [Solved]: Numbers and the Minimum



- Step 4, numbers:
  - $\det \boldsymbol{\Sigma} = 0.02166 \times 0.01820 - 0.00199^2 = 0.000390$; $\boldsymbol{\Sigma}^{-1}\mathbf{1} = (0.02019,\ 0.02366)/0.000390 = (51.7,\ 60.6)$
  - $A = 112.4$; $w^* = 51.7/112.4 = 0.461$ in SPY, 0.539 in TLT; $\sigma_{\mathbf{GMV}} = 1/\sqrt{A} = 9.43\%$; mean $0.461 \times 11.73\% + 0.539 \times 4.32\% = 7.73\%$
- Interpretation: with $\rho < 0$ the two risks partly cancel; the mix is 36% less volatile than SPY and keeps two-thirds of its mean; at $\rho = +0.5$ the minimum is 12.2%

MFM_ch4_seminar

A2 [Solved]: ERC with Two Assets — Task

- **Question**: which SPY–IEF mix gives both assets the same risk contribution?
- **Inputs**: SPY and IEF (7–10 year Treasuries), monthly August 2002 – July 2026, annualised

$$\Sigma = \begin{pmatrix} 0.02166 & -0.00098 \\ -0.00098 & 0.00441 \end{pmatrix}, \quad \sigma_1 = 14.72\%,\ \sigma_2 = 6.64\%,\ \rho = -0.101$$

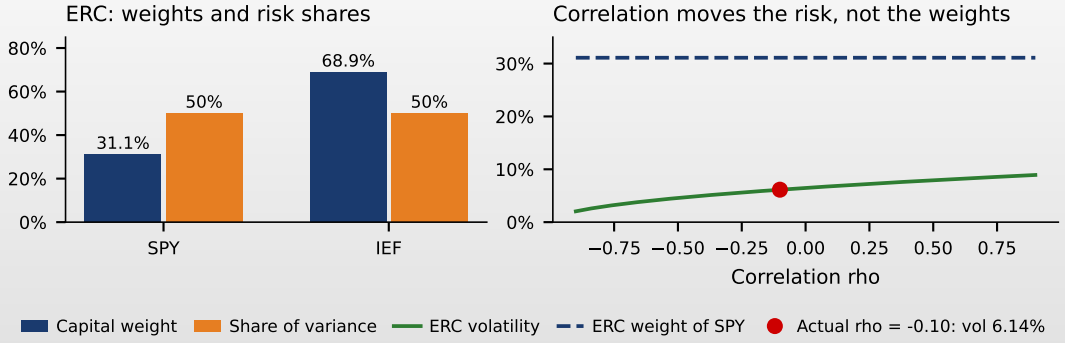
- Three different quantities: $RC_i\sigma_p = w_i(\Sigma \mathrm{w})_i$ (contribution to the variance), RC_i (to the volatility), RC_i/σ_p (the share, in %)
- **Tasks**
  1. Write $RC_1\sigma_p$ and $RC_2\sigma_p$ explicitly for two assets, long-only weights $w_1 + w_2 = 1$.
  2. Prove that $RC_1 = RC_2$ if and only if $w_1\sigma_1 = w_2\sigma_2$, for any $\rho \in (-1, 1)$.
  3. Compute the ERC weight of SPY, then σ_{ERC}, then the two risk shares.
- **Report**: the proof, w_1, σ_{ERC} and the two shares

Back Exercise list

A2 [Solved]: Proof and Numbers

- (1) $RC_1\sigma_p = w_1^2\sigma_1^2 + w_1w_2\sigma_{12}$, $RC_2\sigma_p = w_2^2\sigma_2^2 + w_1w_2\sigma_{12}$
- (2) The shaded cross terms are identical and cancel:
  $RC_1 = RC_2 \iff w_1^2\sigma_1^2 = w_2^2\sigma_2^2 \iff w_1\sigma_1 = w_2\sigma_2$; ρ drops out of the weights, not out of σ_p
- (3) $w_1 = \sigma_1^{-1}/(\sigma_1^{-1} + \sigma_2^{-1}) = 6.79/(6.79 + 15.06) = 0.311$;
  $\sigma_p^2 = 0.311^2(0.02166) + 0.689^2(0.00441) + 2(0.311)(0.689)(-0.00098) = 0.00377$, $\sigma_{\text{ERC}} = 6.14\%$;
  shares 50%, 50%
- Interpretation: equal risk needs more than twice as much of the bond ETF, which is less than half as volatile

A2 [Solved]: Weights, Risk Shares and Correlation



⊡ Left: capital weights against shares of variance; right: the ERC weight of SPY is the same for every ρ, the ERC volatility is not (6.14% at the actual ρ)

MFM_ch4_seminar

A3 [Proposed]: Who Carries the Risk in 60/40? — Task

- **Model**: A2 [Solved]
- **Question**: a "balanced" 60/40 portfolio holds 40% in bonds; how much of its risk do the bonds carry?
- **Inputs**: $\boldsymbol{\Sigma}$ of A2; $\mathrm{w} = (0.6, 0.4)$ in SPY and IEF
- **Tasks**
  1. Compute $\boldsymbol{\Sigma}\mathrm{w}$ row by row: $(\boldsymbol{\Sigma}\mathrm{w})_1 = 0.6\sigma_1^2 + 0.4\sigma_{12}$ and $(\boldsymbol{\Sigma}\mathrm{w})_2 = 0.6\sigma_{12} + 0.4\sigma_2^2$.
  2. Compute $\sigma_p^2 = \mathrm{w}^\top(\boldsymbol{\Sigma}\mathrm{w})$, σ_p and the shares $w_i(\boldsymbol{\Sigma}\mathrm{w})_i/\sigma_p^2$.
  3. Compare σ_p with the ERC volatility of A2 and draw the capital and risk shares of both portfolios side by side.
- **Report**: $(\boldsymbol{\Sigma}\mathrm{w})_1$, $(\boldsymbol{\Sigma}\mathrm{w})_2$, σ_p, the two risk shares and the comparison chart with ERC
- **Interpretation**: is 60/40 a balanced portfolio in risk terms?

B1 [Solved]: 1/N versus Markowitz on Sectors — Task

- **Question**: do estimated Markowitz rules beat 1/N out of sample on nine sector ETFs?
- **Data**: the sector panel of the setup; monthly excess returns January 1999 – July 2026
- **Tasks**
  1. Each month t estimate 1/N, MV (tangency), MV-LO and GMV on months $t-60, \ldots, t-1$ only; hold them in month t: January 2004 – July 2026 ($T = 271$).
  2. Apply the liquidation rule of Chapter 4: a rule whose wealth reaches zero (a monthly loss of 100% or more) is liquidated; its later returns are not an investment outcome, so it is not ranked or tested.
  3. Compute the first month step by step: the window, the weights, the January 2004 return, the drifted weights and the February 2004 turnover.
  4. Compute the annualised mean, volatility, Sharpe ratio, mean turnover and maximum drawdown of each rule that is not liquidated.
  5. Test H_0: $SR_{\mathbf{rule}} = SR_{\mathbf{1}/N}$ with `sr_diff_hac` (Ledoit & Wolf (2008), Section 3.1; Bartlett kernel, 4 lags) and `sr_diff_boot` (studentized circular block bootstrap, 4999 draws; calibrated block lengths (given): MV-LO 1, GMV 1), $\alpha = 5\%$.
- **Report**: the table below and a wealth chart

	Sharpe	Volatility	Turnover	Max. drawdown	$\hat{\Delta}$	HAC p / bootstrap p
1/N, MV-LO, GMV	...	...	...	...	...	...
MV	liquidation month, or the full set of statistics if it survives					

- **Interpretation**: MV-LO beats 1/N by a few hundredths: is that evidence of skill?

Back Exercise list

B1 [Solved]: The First Out-of-Sample Month



- Window 1999-01 – 2003-12 (60 months) gives the weights held in 2004-01; MV-LO holds XLB 81.8%, XLE (energy) 14.3%, XLF (financials) 3.9%, the others 0; MV has gross exposure 8.3
- January 2004 excess returns: 1/N $\sum_i w_i r_i^e = 1.11\%$, MV-LO -3.24%, GMV 1.36%, MV 1.12% (T-bill 0.07%)
- Drift: XLB returned -4.4%, XLK 4.3%; by the end of January the 1/N weights of 11.11% become 10.50% (XLB) and 11.45% (XLK)
- February 2004 turnover $\sum_i |w_{i,\mathbf{Feb}} - w^+_{i,\mathbf{Jan}}|$: 1/N 0.0191, MV-LO 0.748 (the new window also moves the MV-LO weights)

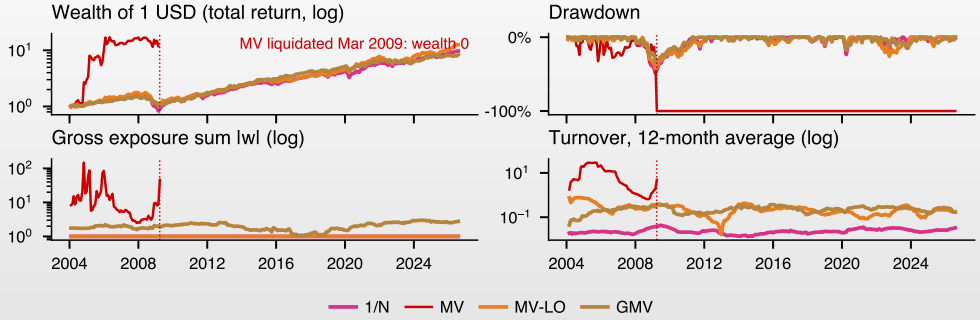
MFM_ch4_seminar

B1 [Solved]: Performance and Tests

	1/N	MV-LO	GMV
Sharpe	0.66	0.72	0.71
Volatility	14.3%	14.7%	11.9%
Turnover	0.02	0.23	0.23
Max. drawdown	−49%	−42%	−33%
$\hat{\Delta}$ vs 1/N	–	+0.06	+0.05
HAC p	–	0.60	0.75
Bootstrap p	–	0.63	0.76

- MV is liquidated in March 2009 (rule of Chapter 4): it is not ranked or tested
  - modelled total return of that month: −167.6%, a theoretical loss of the levered positions; by convention wealth stays at zero and the drawdown is −100%
  - gross exposure: 46.2 in March 2009; median over the 63 months up to liquidation, 8.1
  - mean turnover up to liquidation: 9.8 times wealth a month, in the gross backtest without costs
- MV-LO and GMV land within 0.06 of 1/N; no p-value is below 0.05

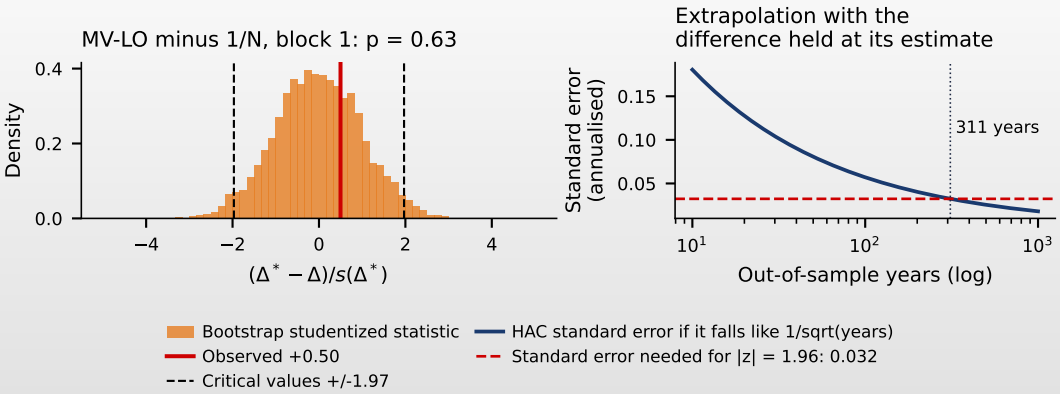
B1 [Solved]: Wealth, Drawdown, Exposure and Turnover



- Top: wealth of 1 USD in total returns and its drawdown; bottom: gross exposure $\sum_i |w_i|$ and turnover, both on log scales; MV-LO and 1/N have gross exposure 1 by construction
- MV: wealth stays at zero after March 2009 (drawdown -100%); exposure and turnover are shown only up to liquidation

MFM_ch4_seminar

B1 [Solved]: Bootstrap Distribution and Extrapolation



MFM_ch4_seminar

B1 [Solved]: Decision and Interpretation

- MV-LO: $\hat{\Delta} = 0.063$, HAC standard error 0.120, HAC $p = 0.60$; studentized bootstrap (block 1) $p = 0.63$, 95% interval $[-0.19, 0.31]$ covers zero: do not reject
- Bootstrap uses a prewhitened quadratic-spectral (QS) standard error: $0.063/0.127 \approx 0.50$. The Bartlett HAC test uses 0.120, hence $z \approx 0.53$.
- Extrapolation (right): if the true gap equalled the estimate and the standard error fell like $1/\sqrt{\text{years}}$, $|z| = 1.96$ would need 311 years of data; a conditional significance calculation, not a power calculation
- Interpretation: the interval admits losses and gains up to 0.31; failure to reject does not establish equality

MFM_ch4_seminar

A4 [Solved]: The Variance of a Sharpe Difference — Task

- **Question**: is the GMV–1/N gap of B1 larger than sampling noise, under the simplest assumptions?
- **Inputs**: the B1 monthly **excess** returns of GMV and 1/N, $T = 271$; monthly $SR_{\text{GMV}} = 0.2051$, $SR_{1/N} = 0.1908$, correlation $\rho = 0.694$
- **Assumption**: i.i.d. bivariate Normal returns; Φ is the standard Normal cumulative distribution function
- **Tasks**
  1. Write the 4×4 asymptotic covariance matrix of $\sqrt{T}(\hat{\mu}_1 - \mu_1, \hat{\mu}_2 - \mu_2, \hat{\sigma}_1^2 - \sigma_1^2, \hat{\sigma}_2^2 - \sigma_2^2)$.
  2. Write the gradient of $f = \mu_1/\sigma_1 - \mu_2/\sigma_2$ and apply the delta method to obtain the formula of The Variance of a Sharpe Difference.
  3. Test H_0: $SR_{\text{GMV}} = SR_{1/N}$ with a two-sided z-test at 5%; compare with the HAC p-value 0.75.
- **Report**: the matrix, the formula, z, p and the decision

Back Exercise list

A4 [Solved]: The Delta Method

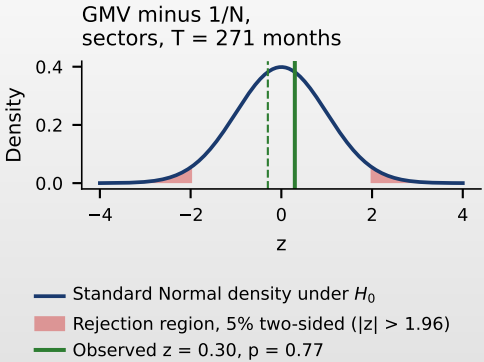
- (1) Under Normality, sample means and variances are asymptotically independent, $\sigma_{12} = \rho\sigma_1\sigma_2$:

$$\Psi = \begin{pmatrix} \sigma_1^2 & \sigma_{12} & 0 & 0 \\ \sigma_{12} & \sigma_2^2 & 0 & 0 \\ 0 & 0 & 2\sigma_1^4 & 2\sigma_{12}^2 \\ 0 & 0 & 2\sigma_{12}^2 & 2\sigma_2^4 \end{pmatrix}, \qquad \nabla f = \begin{pmatrix} 1/\sigma_1 \\ -1/\sigma_2 \\ -SR_1/(2\sigma_1^2) \\ +SR_2/(2\sigma_2^2) \end{pmatrix}$$

- (2) $\nabla f^\top \Psi \nabla f$ splits into two blocks
  - Means: $\sigma_1^2/\sigma_1^2 + \sigma_2^2/\sigma_2^2 - 2\sigma_{12}/(\sigma_1\sigma_2) = 2 - 2\rho$
  - Variances: $\frac{SR_1^2}{4\sigma_1^4} 2\sigma_1^4 + \frac{SR_2^2}{4\sigma_2^4} 2\sigma_2^4 - 2\frac{SR_1 SR_2}{4\sigma_1^2\sigma_2^2} 2\sigma_{12}^2 = \frac{1}{2}(SR_1^2 + SR_2^2 - 2SR_1 SR_2\rho^2)$
- Adding the blocks gives the formula (Jobson & Korkie (1981); corrected form in Ledoit & Wolf (2008))

A4 [Solved]: The Test

- (3) Bracket: $2 - 2(0.694) + 0.5(0.2051^2 + 0.1908^2 - 2 \times 0.2051 \times 0.1908 \times 0.694^2)$
  $= 0.612 + 0.5(0.0421 + 0.0364 - 0.0377) = 0.6324$
- Standard error $\sqrt{0.6324/271} = 0.0483$;
  $z = (0.2051 - 0.1908)/0.0483 = 0.30$;
  $p = 2[1 - \Phi(0.30)] = 0.77$
- Decision: do not reject H_0 at 5%; the HAC test agrees ($p = 0.75$)
- Interpretation: the annualised gap, 0.71 against 0.66, is well inside sampling noise; 22 years cannot rank GMV and 1/N



MFM_ch4_seminar

A5 [Proposed]: ERC versus 1/N — Task

- **Model**: A4 [Solved]
- **Question**: ERC beats 1/N by less than GMV does; can the smaller gap be closer to significance?
- **Inputs**: B1 backtest with ERC added, $T = 271$; monthly $SR_{\text{ERC}} = 0.2018$, $SR_{1/N} = 0.1908$, correlation 0.9927; keep four decimals until the end (the bracket is small when ρ is near 1)
- Reference: the formula for $T\,\text{Var}(\hat{\Delta})$ in The Variance of a Sharpe Difference
- **Tasks**
  1. Compute the bracket, the standard error, z and the two-sided p; decide at 5%.
  2. Show that the bracket tends to $\frac{1}{2}(SR_1 - SR_2)^2$ as $\rho \to 1$.
  3. Fill a comparison table (gap, ρ, standard error, z) for A4 and A5 and plot both 95% intervals.
- **Report**: the bracket, the standard error, z, the p-value and the decision at 5%; the comparison table and the interval chart
- **Interpretation**
  - why is a smaller difference here closer to significance?
  - what happens if the two return series are identical?

Back Exercise list

B2 [Proposed]: The Break-Even Cost — Task

- **Model**: B1 [Solved]
- **Question**: MV-LO beats 1/N gross on 16 ETFs; at which trading cost does the edge vanish?
- **Data**: the combined universe (16 ETFs), OOS May 2012 – July 2026 ($T = 171$); strategy 1 = MV-LO, strategy 2 = 1/N; turnover as in B1
- Net return $r_t - c\,TO_t$, cost c per unit traded: Trading Costs and the Break-Even Cost
- **Tasks**
  1. Compute the gross Sharpe ratios, the HAC p-value of their difference and the mean turnovers.
  2. On a grid of 0–200 bp (step 0.1 bp), recompute the mean **and** the standard deviation of both net return series and find the cost c^* where the net Sharpe ratios cross; plot both curves.
  3. Derive the fixed-volatility approximation $c^* \approx \Delta SR / [12(TO_1/\sigma_1 - TO_2/\sigma_2)]$ (annual σ) and compare it with the grid value.
- **Report**: the two gross Sharpe ratios, the HAC p-value, the mean turnovers, c^* from the grid and from the approximation, the chart
- **Interpretation**: retail costs on BVB or US ETFs: is MV-LO worth implementing?

Back Exercise list

C1 [Proposed]: Critique an AI Answer — Task

- **Model**: B1 and A4 [Solved]. **Prompt**: "`prices` holds daily adjusted closes of the nine sector ETFs of B1 (December 1998 – July 2026). Write Python code for an out-of-sample GMV backtest with monthly rebalancing and give me its annualised Sharpe ratio."

```
import numpy as np
r = prices.resample('ME').last().pct_change().dropna()  # monthly simple returns
S = r.cov().values                                      # covariance matrix
w = np.linalg.solve(S, np.ones(len(S)))                 # Sigma^{-1} 1
w = w / w.sum()                                         # GMV weights, sum to one
port = r.dot(w)                                         # portfolio return
sr = port.mean() / port.std() * np.sqrt(252)
print(f"Annualised Sharpe ratio: {sr:.2f}")
```

- **Tasks** (`rf`: the monthly T-bill of the setup, decimals, aligned on month-end dates)
  1. Find the three intended errors; name the line, explain why it is wrong and how you would detect it (a date check, a number from B1 or A4, a definition). Further qualified remarks are welcome.
  2. Use the B1 design: 60-month rolling estimation, weights chosen before each month, monthly rebalancing with drift. Write the corrected code and report its Sharpe ratio; check the date lag, missing values, weight sums and the annualisation.
  3. Log each error in the format of `AI_ERRORS.md`: request, answer, error, how found, correction.
- **Interpretation**
  - how far is the AI's number from a defensible one?
  - which error matters most?

Back Exercise list

B3 [Solved]: Does Shrinkage Lower GMV Risk? — Task

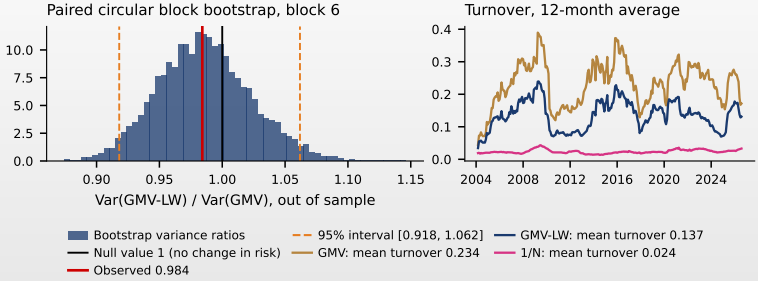
- **Question**: does the Ledoit–Wolf estimator lower the out-of-sample risk of GMV on sectors?
- **Data**: the B1 backtest; GMV with the sample covariance versus GMV-LW (GMV with the Ledoit–Wolf covariance), the constant-correlation target of Ledoit & Wolf (2004) (helper `lw_cc(X)`: input a 60×9 window of excess returns, output $\hat{\Sigma}_{\text{LW}}$ and δ); S, F, $\bar{r}$, δ and the condition number as in Covariance Shrinkage (B3, after the Lecture)
- **Tasks**
  1. On the last window, follow the pipeline: S $\rightarrow$ target F $\rightarrow \delta \rightarrow \hat{\Sigma}_{\text{LW}} \rightarrow$ GMV weights; report $\bar{r}$, δ, both condition numbers and the gross exposures.
  2. Compute the out-of-sample volatilities of GMV, GMV-LW and 1/N.
  3. Compute Var(GMV-LW)/ Var(GMV) with a 95% percentile interval from a circular block bootstrap (Künsch (1989); blocks of 6 months, 5000 draws); both series use the **same** resampled months.
  4. Compare the mean monthly turnover of the two rules.
- **Interpretation**
  - why is the gain small with 9 assets and 60 months?
  - when would it be larger?

Back Exercise list

B3 [Solved]: From S to Weights

- Last window 2021-08 – 2026-07, $T = 60$, $N = 9$
- Step 1: average sample correlation $\bar{r} = 0.511$; e.g. XLK–XLU 0.31 is replaced by $\bar{r}$ in F, the variances are kept
- Step 2: $\delta = 0.24$: $\hat{\Sigma}_{\text{LW}}$ is 24% target, 76% sample
- Step 3: condition number 51.6 $\rightarrow$ 32.5: the smallest eigenvalues are pulled up
- Step 4: GMV weights; gross exposure $\sum_i |w_i|$ 2.64 (sample) $\rightarrow$ 1.91 (shrunk); XLP (consumer staples) 66.8% $\rightarrow$ 57.8%, XLK 10.1% $\rightarrow$ 6.4%
- One paired resample: draw 46 random start months, take 6 consecutive months from each (wrapping around), cut at $T = 271$, and compute the variance of both strategies on those same months

B3 [Solved]: Risk and Turnover



- Volatility: GMV 11.90%, GMV-LW 11.80%, 1/N 14.28%: both GMV rules cut risk by about a sixth
- Variance ratio 0.984, interval $[0.918, 1.062]$ contains 1: no significant reduction; turnover 0.234 $\rightarrow$ 0.137 (-41%): the benefit here is stability, not risk
- Interpretation: with $N/T = 0.15$ the sample matrix is already decent; the gain grows with N/T and with a target close to the truth

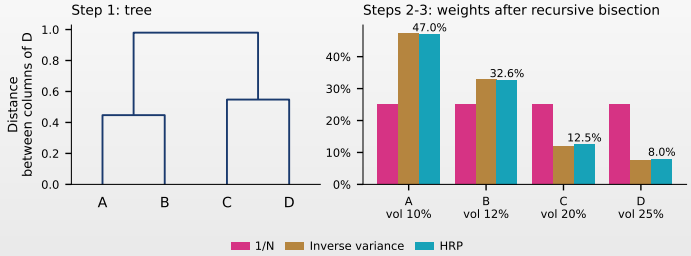
MFM_ch4_seminar

B4 [Solved]: ERC and HRP on Sectors — Task

- **Questions**
  - how do the two risk-based rules allocate?
  - do they beat 1/N?
- **Part 1, computing the weights**: first the four-asset example, then the last window August 2021 – July 2026, with the sample covariance
  1. Apply HRP step by step on the four-asset example of the next slide: distances, tree, order, bisection.
  2. Compute the nine ERC weights (`w_erc(S)`, the convex problem of Maillard, Roncalli & Teïletche (2010)) and HRP weights (`w_hrp(S)`, the four steps of Hierarchical Risk Parity (B4, after the Lecture)).
  3. Compute the risk shares of 1/N, ERC and HRP.
- **Part 2, backtest** (January 2004 – July 2026, rolling 60 months, as in B1)
  4. Test the Sharpe ratios of ERC and HRP against 1/N with the HAC test at 5%; report turnover and the correlation with 1/N.
- **Interpretation**: ERC and 1/N have correlation 0.993: is ERC a different strategy or a small tilt?

Back Exercise list

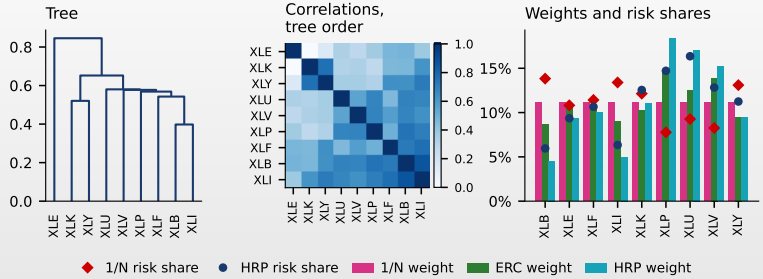
B4 [Solved]: HRP on Four Assets



- Volatilities $10\%, 12\%, 20\%, 25\%$; correlations A–B 0.8, C–D 0.7, all cross pairs 0.2
- Step 1: $d_{AB} = \sqrt{(1-0.8)/2} = 0.316$, $d_{CD} = 0.387$, cross pairs 0.632; distances between columns of D: $\tilde{d}_{AB} = 0.447$, $\tilde{d}_{CD} = 0.548$; single linkage merges A–B, then C–D, then the two clusters
- Step 2: order A, B, C, D; split into $\{A, B\}$ and $\{C, D\}$
- Step 3: inverse-variance weights inside each half, $(0.59, 0.41)$ and $(0.61, 0.39)$; cluster variances $V_{\mathbf{1}} = 0.01055$, $V_{\mathbf{2}} = 0.04105$; $\alpha = 1 - V_{\mathbf{1}}/(V_{\mathbf{1}} + V_{\mathbf{2}}) = 0.796$ to $\{A, B\}$
- Then split each pair again by inverse variance: A 47.0%, B 32.6%, C 12.5%, D 8.0%

MFM_ch4_seminar

B4 [Solved]: Nine Sectors, Last Window



- Tree order: XLE, XLK, XLY (consumer discretionary), XLU, XLV (health care), XLP, XLF, XLB, XLI (industrials)
- ERC: XLP 14.8%, XLV 13.9% ... XLB 8.6%; HRP: XLP 18.4%, XLU 17.1% ... XLB 4.5%
- 1/N risk shares: from 7.8% (XLP) to 13.8% (XLB)

MFM_ch4_seminar

B4 [Solved]: Backtest

- OOS January 2004 – July 2026, rolling 60 months, HAC test against 1/N
- Backtest: ERC Sharpe 0.70 vs 0.66, $p = 0.09$; HRP 0.70, $p = 0.32$; turnover 1/N 0.024, ERC 0.028, HRP 0.082; correlation with 1/N: ERC 0.993, HRP 0.979
- Interpretation: ERC is a small, cheap tilt towards low-volatility sectors; HRP a larger tilt that trades three times as much; neither rejects H_0

C2 [Proposed]: A BVB Portfolio in RON — the Take-Home Project

- **Model**: B1, B4 and A4 [Solved]
- **Question**: does diversification beyond the BVB pay for a Romanian investor?
- **Data and currency**: the eight BVB stocks, TLV (Banca Transilvania), SNP (OMV Petrom), BRD (BRD–Groupe Société Générale), TGN (Transgaz), SNG (Romgaz), SNN (Nuclearelectrica), EL (Electrica), TEL (Transelectrica) and the BET-TR index (October 2014 – August 2026); SPY, EFA (developed markets ex US and Canada), IEF, GLD (gold) converted to RON on common days, $P_{\mathbf{RON}} = P_{\mathbf{USD}} \times$ USD/RON (BNR); e.g. 2017-10-02: SPY 219.84 USD $\times$ 3.9086 = 859.27 RON
- **Base case**: monthly returns in RON, 36-month window, first portfolio October 2017, $R_f = 0$
- **Tasks**
  1. Backtest 1/N, GMV-LW, ERC and HRP on the eight stocks, then on the stocks plus the four foreign ETFs.
  2. Compare each with BET-TR (never BET): Sharpe ratio, volatility, maximum drawdown; plot wealth and drawdown.
  3. Test each global rule (stocks plus ETFs) against BET-TR and against its BVB-only version (HAC), with a Holm adjustment (Noise Eigenvalues and Multiple Tests) over this family of eight tests.
  4. Sensitivity: a RON rate of 5% a year, 0.05/12 a month, subtracted from all returns.
- **Deliverables**: a Quantlet folder in your repository, a table of the eight tests (raw and Holm p), one slide with the conclusion and its main limitation (currency risk, costs, taxes)
- **Assessment criteria**: a backtest without look-ahead; HAC tests and the Holm adjustment; the interpretation of the limits (currency risk, costs, taxes); reproducible code

Back Exercise list

Extension Packs: What Each One Needs

Pack	Exercises	Prerequisites
1. Risk parity theory	A3 Extended	A2, A3
2. Black–Litterman	A6 [Solved], A7, B5	lecture, Black–Litterman
3. Bias and regression inference	A8, A9, B6	A1, A4, B1
4. Covariance estimation	B3 Extended [Solved], B7	B3
5. Beyond the sectors	B8, B9	B1, B4
6. Research projects	C3 (C2 is the base homework)	B1, A4, the lecture case study

- Every pack has its own notebook section; do them in order within a pack
- Further study: Hierarchical Clustering, Minimum Spanning Tree

Black–Litterman: Views and Uncertainty

- **Black–Litterman** Black & Litterman (1992): combine a prior for the expected returns with the investor's views
- Prior from reverse optimisation: the means that make the benchmark w_b optimal

$$\boldsymbol{\pi} = \delta_{\mathrm{RA}} \boldsymbol{\Sigma} \mathrm{w}_b, \qquad \boldsymbol{\mu} \sim N(\boldsymbol{\pi},\ \tau \boldsymbol{\Sigma})$$

  - δ_{RA}: risk aversion; τ: small scalar, the uncertainty of the prior relative to $\boldsymbol{\Sigma}$
- Views: K linear statements about $\boldsymbol{\mu}$, each with its own uncertainty

$$\mathrm{q} = \mathrm{P}\boldsymbol{\mu} + \mathrm{e}, \qquad \mathrm{e} \sim N(0, \boldsymbol{\Omega})$$

  - P: $K \times N$ matrix (a row $(0, \dots, +1, \dots, -1, \dots)$ is a spread view); $\boldsymbol{\Omega}$: covariance of the view errors
- One asset, one view: the posterior mean is the precision-weighted average

$$\mu_{\mathrm{BL}} = \frac{\pi/(\tau\sigma^2) + q/\omega}{1/(\tau\sigma^2) + 1/\omega}, \qquad \text{precision} = 1/\text{variance}$$

  - $\tau\sigma^2$: variance of the prior; ω: variance of the view; the more precise source gets the larger weight

Noise Eigenvalues and Multiple Tests

- **Marchenko–Pastur band** Laloux et al. (1999): asymptotic noise benchmark for eigenvalues of a correlation matrix of N independent series, T observations each:

$$\left[(1-\sqrt{c})^2,\ (1+\sqrt{c})^2\right], \qquad c = N/T$$

  - departures from the band can indicate structure; finite samples and dependence also affect them, so this is not an exact classification

- **Holm adjustment** for m tests (B8, C2): order the p-values $p_{(1)} \leq \cdots \leq p_{(m)}$

$$\tilde{p}_{(k)} = \max_{j \leq k} \ \min\left\{1,\ (m-j+1)\, p_{(j)}\right\}$$

  - reject the hypotheses with $\tilde{p}_{(k)} < \alpha$; the chance of any false rejection in the family stays below α

- Exact F and HAC Wald tests of linear restrictions $\mathrm{R}b = 0$: the formulas are on the B6 slide

A3 Extended [Proposed]: Euler and the Uniqueness of ERC — Task

- **Model**: A2 and A3
- **Tasks**
  1. Show that $\sigma_p(cw) = c\,\sigma_p(w)$ for $c > 0$ and deduce $\sum_i RC_i = \sigma_p$ from Euler's theorem for functions homogeneous of degree one.
  2. For $\boldsymbol{\Sigma}$ positive definite, let $f(y) = \frac{1}{2} y^\top \boldsymbol{\Sigma} y - \frac{1}{N} \sum_i \ln y_i$ on $y > 0$. Compute its Hessian and show that f is strictly convex.
  3. Show that $f \to +\infty$ as any $y_i \to 0$ and as $\|y\| \to \infty$; conclude that a minimiser exists and is unique.
  4. Write the first-order condition and show that it equalises the risk contributions, before and after the normalisation $w = y/1^\top y$.
- **Report**: the two proofs in at most ten lines

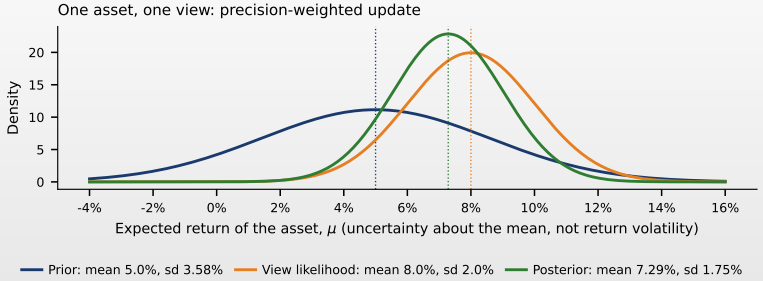
Back Exercise list

A6 [Solved]: Black–Litterman with One Asset — Task

- **Question**: how much should a precise view move an uncertain prior mean?
- **Inputs**: one asset, return volatility $\sigma = 16\%$, $\tau = 0.05$; prior for the **expected** return $\mu \sim N(\pi, \tau\sigma^2)$ with $\pi = 5\%$; view $q = 8\%$ with standard deviation 2%
  - Both variances measure uncertainty about the mean μ, not the volatility of returns; Normality is assumed for the posterior density
- **Tasks**
  1. Compute the prior variance $\tau\sigma^2$ and the two precisions (inverse variances).
  2. Compute μ_{BL} as the precision-weighted mean, the weight of the view and the posterior standard deviation; plot the three densities.
  3. Generalise: stack $\boldsymbol{\pi} = \boldsymbol{\mu} + \mathrm{e}_1$ and $\mathrm{q} = \mathrm{P}\boldsymbol{\mu} + \mathrm{e}_2$ as $\mathrm{y} = \mathrm{Z}\boldsymbol{\mu} + \mathrm{e}$ (Z: I stacked over P; Var(e) = V, block-diagonal with $\tau\boldsymbol{\Sigma}$ and $\boldsymbol{\Omega}$) and derive the generalised least squares (GLS) estimator (Theil & Goldberger (1961)); show that $\mathrm{P} = \mathrm{I}$, $\boldsymbol{\Omega} = \tau\boldsymbol{\Sigma}$ gives $(\boldsymbol{\pi} + \mathrm{q})/2$.
- **Report**: the prior variance, the two precisions, μ_{BL}, the weight on the view, the posterior standard deviation and the chart of the three densities

Back Exercise list

A6 [Solved]: One Asset



- (1) Prior variance $0.05 \times 0.16^2 = 0.00128$, precision 781; view variance $0.02^2 = 0.0004$, precision 2500
- (2) $\mu_{\mathbf{BL}} = (781 \times 5\% + 2500 \times 8\%)/3281 = 7.29\%$; view weight $2500/3281 = 0.76$; posterior standard deviation $1/\sqrt{3281} = 1.75\%$
- The view is three times as precise as the prior, so it gets three-quarters of the weight; the posterior is more precise than either source

MFM_ch4_seminar

A6 [Solved]: Black–Litterman as GLS

- (3) Dimensions: $\boldsymbol{\mu}$ is $N \times 1$, q is $K \times 1$, P is $K \times N$

$$\underbrace{\begin{pmatrix} \boldsymbol{\pi} \\ \mathrm{q} \end{pmatrix}}_{\mathrm{y}\ ((N+K)\times 1)} = \underbrace{\begin{pmatrix} \mathrm{I}_N \\ \mathrm{P} \end{pmatrix}}_{\mathrm{Z}} \boldsymbol{\mu} + \mathrm{e}, \qquad \mathrm{V} = \mathrm{Var}(\mathrm{e}) = \begin{pmatrix} \tau\boldsymbol{\Sigma} & 0 \\ 0 & \boldsymbol{\Omega} \end{pmatrix}$$

- GLS: $\hat{\boldsymbol{\mu}} = (\mathrm{Z}^\top \mathrm{V}^{-1} \mathrm{Z})^{-1} \mathrm{Z}^\top \mathrm{V}^{-1} \mathrm{y}$; one block multiplication: $\mathrm{Z}^\top \mathrm{V}^{-1} = \left((\tau\boldsymbol{\Sigma})^{-1},\ \mathrm{P}^\top \boldsymbol{\Omega}^{-1}\right)$, so

$$\boldsymbol{\mu}_{\mathrm{BL}} = \left[(\tau\boldsymbol{\Sigma})^{-1} + \mathrm{P}^\top \boldsymbol{\Omega}^{-1} \mathrm{P}\right]^{-1} \left[(\tau\boldsymbol{\Sigma})^{-1} \boldsymbol{\pi} + \mathrm{P}^\top \boldsymbol{\Omega}^{-1} \mathrm{q}\right]$$

- Covariance $[(\tau\boldsymbol{\Sigma})^{-1} + \mathrm{P}^\top \boldsymbol{\Omega}^{-1} \mathrm{P}]^{-1}$: the Black–Litterman posterior
- $\mathrm{P} = \mathrm{I}$, $\boldsymbol{\Omega} = \tau\boldsymbol{\Sigma}$: $[2(\tau\boldsymbol{\Sigma})^{-1}]^{-1} (\tau\boldsymbol{\Sigma})^{-1} (\boldsymbol{\pi} + \mathrm{q}) = (\boldsymbol{\pi} + \mathrm{q})/2$: equal precisions, equal weights

A7 [Proposed]: A Sector View — Task

- **Model**: A6 [Solved]
- **Inputs**: nine sectors in the order XLB, XLE, XLF, XLI, XLK, XLP, XLU, XLV, XLY; monthly excess returns August 2016 – July 2026, $\boldsymbol{\Sigma}$ annualised; benchmark 1/N, risk aversion $\delta_{\mathrm{RA}} = 2.5$, $\tau = 0.05$; Δw: change of the weights relative to the benchmark
- View XLK minus XLU $= 3\%$ a year: $\mathrm{P} = (0, 0, 0, 0, +1, 0, -1, 0, 0)$; implied $\mathrm{P}\boldsymbol{\pi} = 2.22\%$; spread variance $\mathrm{P}\boldsymbol{\Sigma}\mathrm{P}^{\top} = 0.0477$ (enough: the full $\boldsymbol{\Sigma}$ is only needed for a check)
- **Tasks**
  1. Write the prior distribution of $\mathrm{P}\boldsymbol{\mu}$ implied by $\boldsymbol{\mu} \sim N(\boldsymbol{\pi}, \tau\boldsymbol{\Sigma})$.
  2. With $\Omega = \tau\,\mathrm{P}\boldsymbol{\Sigma}\mathrm{P}^{\top}$ (He & Litterman (2002)), show that the posterior view is the simple average of prior and view, and compute it.
  3. Show that $\Delta\mathrm{w} = \mathrm{P}^{\top}(q - \mathrm{P}\boldsymbol{\pi})/(2\delta_{\mathrm{RA}}\mathrm{P}\boldsymbol{\Sigma}\mathrm{P}^{\top})$; compute the scalar, the new XLK and XLU weights, and check that the tilt sums to zero.
- **Report**: the prior of $\mathrm{P}\boldsymbol{\mu}$, the posterior view, the scalar and the new XLK and XLU weights
- **Interpretation**: how does the tilt change if the view variance Ω doubles (its precision halves)?

Back Exercise list

B5 [Proposed]: Black–Litterman on a Multi-Asset Set — Task

- **Model**: A6 and A7
- **Inputs**: SPY, EFA, EEM (iShares MSCI Emerging Markets ETF), TLT, IEF, LQD (iShares iBoxx Investment Grade Corporate Bond ETF), HYG (iShares iBoxx High Yield Corporate Bond ETF), GLD (this order), monthly excess returns August 2016 – July 2026 ($T = 120$), $\boldsymbol{\Sigma}$ annualised; benchmark 1/N, $\delta_{\mathrm{RA}} = 2.5$, $\tau = 0.05$; view GLD minus TLT $= 4\%$ a year, $\mathrm{P} = (0, 0, 0, -1, 0, 0, 0, +1)$, $\Omega = \tau \mathrm{P} \boldsymbol{\Sigma} \mathrm{P}^{\top}$
- Unconstrained static allocation: a negative weight is a short position
- **Tasks**
  1. Compute $\boldsymbol{\pi} = \delta_{\mathrm{RA}} \boldsymbol{\Sigma} \mathrm{w}_b$ and the implied GLD minus TLT.
  2. Compute the spread variance, the surprise $q - \mathrm{P}\boldsymbol{\pi}$, the scalar tilt of A7 and the new weights; plot them against the benchmark.
  3. Compare with the sample mean of GLD minus TLT, with a ten times lower confidence ($\Omega \times 10$) and with a long-only allocation maximizing $\boldsymbol{\mu}_{\mathrm{BL}}^{\top} \mathrm{w} - \delta_{\mathrm{RA}} \mathrm{w}^{\top} \boldsymbol{\Sigma} \mathrm{w}/2$ subject to nonnegative weights summing to one. Compare the sample mean with the stated view; do not substitute it for the view.
- **Report**: $\boldsymbol{\pi}$, the implied spread, the surprise, the scalar, the new weights and the three comparisons, with the chart
- **Interpretation**: the weights move far although the view is modest: why?

Back Exercise list

A8 [Proposed]: The In-Sample Sharpe Bias — Task

- **Model**: A4 [Solved] (for the moments), estimation risk in Chapter 4
- **Assumptions**: i.i.d. $N(\boldsymbol{\mu}, \boldsymbol{\Sigma})$ returns; maximum-likelihood $\hat{\boldsymbol{\mu}}$, $\hat{\boldsymbol{\Sigma}}$ (divisor T); given: $\hat{\boldsymbol{\mu}} \sim N(\boldsymbol{\mu}, \boldsymbol{\Sigma}/T)$, $\mathbb{E}[\hat{\boldsymbol{\Sigma}}^{-1}] = \frac{T}{T-N-2}\boldsymbol{\Sigma}^{-1}$ for $T > N+2$, $\hat{\boldsymbol{\mu}}$ independent of $\hat{\boldsymbol{\Sigma}}$, and $\mathbb{E}[\mathrm{x}^\top \mathrm{A}\mathrm{x}] = \mathrm{tr}(\mathrm{A}\,\mathrm{Var}(\mathrm{x})) + \mathbb{E}[\mathrm{x}]^\top \mathrm{A}\,\mathbb{E}[\mathrm{x}]$
- **Tasks**
  1. Compute $\mathbb{E}[\hat{\theta}^2 \mid \hat{\boldsymbol{\mu}}]$ for $\hat{\theta}^2 = \hat{\boldsymbol{\mu}}^\top \hat{\boldsymbol{\Sigma}}^{-1} \hat{\boldsymbol{\mu}}$.
  2. Take the expectation over $\hat{\boldsymbol{\mu}}$ and conclude $\mathbb{E}[\hat{\theta}^2] = (T\theta^2 + N)/(T - N - 2)$ (Kan & Zhou (2007)).
  3. Evaluate for $N = 9$, $T = 60$, $\theta = 0.66/\sqrt{12}$ monthly; report $\sqrt{12\,\mathbb{E}[\hat{\theta}^2]}$, the annualised root mean square of the estimate (not its mean), and compare with a simulation.
  4. Give the unbiased estimator of θ^2 implied by the formula.

Back Exercise list

A9 [Proposed]: Tangency Weights by Regression — Task

- **Model**: A1 [Solved] (matrix algebra); the Chapter 4 Appendix on Britten-Jones (1999)
- **Setting**: X is the $T \times N$ matrix of excess returns, rows $r_t^\top$; 1_T the response (a column of ones), 1_N the normalisation vector; $\hat{\boldsymbol{\mu}} = \mathrm{X}^\top 1_T / T$, $\hat{\boldsymbol{\Sigma}} = \mathrm{X}^\top \mathrm{X}/T - \hat{\boldsymbol{\mu}}\hat{\boldsymbol{\mu}}^\top$, invertible
- **Tasks**
  1. Show $\mathrm{X}^\top \mathrm{X}/T = \hat{\boldsymbol{\Sigma}} + \hat{\boldsymbol{\mu}}\hat{\boldsymbol{\mu}}^\top$ and write the ordinary least squares (OLS) estimator of $1 = \mathrm{b}^\top \mathrm{r}_t + u_t$ (no constant) in terms of $\hat{\boldsymbol{\mu}}$, $\hat{\boldsymbol{\Sigma}}$.
  2. Apply Sherman–Morrison, $(\mathrm{A} + \mathrm{uu}^\top)^{-1} = \mathrm{A}^{-1} - \mathrm{A}^{-1}\mathrm{uu}^\top \mathrm{A}^{-1}/(1 + \mathrm{u}^\top \mathrm{A}^{-1}\mathrm{u})$, to show $\hat{\mathrm{b}} = \hat{\boldsymbol{\Sigma}}^{-1}\hat{\boldsymbol{\mu}}/(1 + \hat{\theta}^2)$.
  3. Show that $\hat{\mathrm{b}}/1_N^\top \hat{\mathrm{b}}$ are the sample tangency weights when $1_N^\top \hat{\boldsymbol{\Sigma}}^{-1}\hat{\boldsymbol{\mu}} > 0$.
  4. Optional: for $N = 3$ write the restriction matrix of "tangency = 1/N", then generalise.
- **Report**: the three proofs (and the restriction matrix, optional), in at most one page

Back Exercise list

B6 [Proposed]: Is the Tangency Portfolio 1/N? — Task

- **Model**: A9 and B1
- **Data**: nine sectors, monthly excess returns January 1999 – July 2026 ($T = 331$) and the halves January 1999 – September 2012, October 2012 – July 2026
- **Tests** of H_0: $\mathrm{Rb} = 0$, $\mathrm{R} = [\mathrm{I}_8,\ -1_8]$, 8 restrictions
  - Exact: $F = (\mathrm{R}\hat{\mathrm{b}})^\top[\mathrm{R}(\mathrm{X}^\top\mathrm{X})^{-1}\mathrm{R}^\top]^{-1}(\mathrm{R}\hat{\mathrm{b}})/(8s^2) \sim F(8, T-9)$, $s^2 = \hat{\mathrm{u}}^\top\hat{\mathrm{u}}/(T-9)$
  - HAC Wald: $W = (\mathrm{R}\hat{\mathrm{b}})^\top(\mathrm{R}\hat{\mathrm{V}}\mathrm{R}^\top)^{-1}\mathrm{R}\hat{\mathrm{b}} \to \chi^2_8$, $\hat{\mathrm{V}} = T(\mathrm{X}^\top\mathrm{X})^{-1}\hat{\mathbf{\Omega}}(\mathrm{X}^\top\mathrm{X})^{-1}$, $\hat{\mathbf{\Omega}}$ Newey–West of $\mathrm{r}_t\hat{u}_t$, Bartlett, $\lfloor 4(T/100)^{2/9} \rfloor$ lags
- **Tasks**
  1. Fill a 3×2 table of statistics and p-values (three samples, two tests); plot the coefficients with OLS and HAC intervals.
  2. Simulate 2000 samples of $T = 166$ from a Normal model with $\boldsymbol{\mu} \propto \boldsymbol{\Sigma}1$ (so H_0 holds, $\boldsymbol{\Sigma}$ of the second half) and report the rejection rate of each test at 5%. Scale the mean vector so 1/N has its observed second-half monthly Sharpe ratio.
- **Report**: the 3×2 table, the coefficient chart with OLS and HAC intervals, the empirical size of both tests from the simulation
- **Interpretation**
  - when the tests disagree, which do you trust?
  - does non-rejection mean the tangency portfolio is 1/N?

Back Exercise list

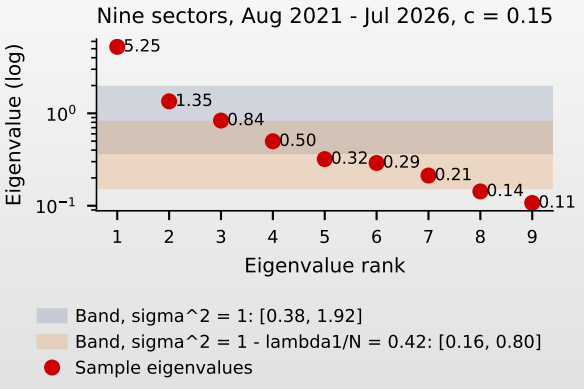
B3 Extended [Solved]: Eigenvalues of the Last Window — Task

- **Question**: how many eigenvalues of the sector correlation matrix are distinguishable from noise?
- **Data**: last window of B3, August 2021 – July 2026, $N = 9$, $T = 60$
- **Tasks**
  1. Compute the nine eigenvalues of the sample correlation matrix.
  2. Compute the Marchenko–Pastur band (Noise Eigenvalues and Multiple Tests) and count the eigenvalues above and below it.
  3. Repeat with the noise variance rescaled to $\sigma^2 = 1 - \lambda_1/N$, λ_1 the largest eigenvalue (Laloux et al. (1999)).
- **Report**: the nine eigenvalues, the Marchenko–Pastur band and the number of eigenvalues above and below it, in both variants
- **Interpretation**: why does GMV care most about the eigenvalues *below* the band?

Back Exercise list

B3 Extended [Solved]: Eigenvalues Against the Noise Band

- (1) eigenvalues: 5.25, 1.35, 0.84, 0.50, 0.32, 0.29, 0.21, 0.14, 0.11
- (2) $c = 0.15$, band $[0.38, 1.92]$: 1 above (the market mode), 5 below
- (3) $\sigma^2 = 0.42$, band $[0.16, 0.80]$: 3 above, 2 below
- Interpretation: GMV loads on the smallest eigenvalues, the directions of low standardised in-sample variance; if they are underestimated, so is GMV risk; shrinkage pulls them up



MFM_ch4_seminar

B7 [Proposed]: Shrinkage on 16 ETFs — Task

- **Model**: B3 [Solved]
- **Data**: combined universe (nine sectors plus EFA, EEM, TLT, IEF, LQD, HYG, GLD), common months May 2007 – July 2026; 60-month window, OOS May 2012 – July 2026 ($T = 171$)
- **Tasks** (output: one table with rows GMV, GMV-LW, GMV-NL (GMV with nonlinear shrinkage, task 3), 1/N and columns volatility, turnover, variance ratio with interval)
  1. Out-of-sample volatility of GMV, GMV-LW and 1/N; the variance ratio GMV-LW/GMV with the paired block bootstrap of B3; mean turnovers.
  2. Mean δ and the condition numbers of the sample and shrunk matrices on the last window.
  3. Extension: nonlinear shrinkage (Ledoit & Wolf (2020); helper `nl_shrink(X)`), the GMV-NL backtest and its variance ratios to GMV and GMV-LW.
- **Interpretation**: is constant correlation a sensible target for stocks and Treasuries together? Distinguish a significant change from a point estimate

Back Exercise list

B8 [Proposed]: Risk Parity with Bitcoin — Task

- **Model**: B4 [Solved]; literature: Petukhina, Trimborn, Härdle & Elendner (2021), Pele et al. (2023)
- **Data**: the eight multi-asset ETFs with and without Bitcoin (BTC/USD), prices aligned on common trading days (weekends dropped), month-end returns from August 2010; both universes on the **same** months
- **Tasks**
  1. Descriptive, August 2010 – July 2026: Bitcoin's volatility and its risk share in 1/N of nine assets.
  2. Out of sample, 60-month window, August 2015 – July 2026 ($T = 132$): backtest 1/N, ERC and HRP with and without Bitcoin; Sharpe ratios, volatilities, drawdowns; HAC test of each with-minus-without difference.
  3. Plot Bitcoin's weight in ERC and HRP and its risk share in 1/N over time.
  4. Treat the three tests as one family: report raw and Holm-adjusted p-values.
- **Report**: Bitcoin's volatility and risk share, the backtest table, the weight chart and the raw and Holm p-values
- **Interpretation**
  - would you expect the past Bitcoin premium to persist?
  - what about the ex-post choice of Bitcoin?

Back Exercise list

B9 [Proposed]: HRP versus GMV on 16 ETFs — Task

- **Model**: B1 [Solved]
- **Data**: the combined universe and backtest of B7 (OOS May 2012 – July 2026, $T = 171$)
- **Tasks**
  1. Report Sharpe ratio, volatility and turnover of HRP and GMV, gross and net of 50 bp per unit of turnover; plot the gross and net points.
  2. Primary test on **gross** excess returns: H_0: $SR_{\text{HRP}} = SR_{\text{GMV}}$ with the HAC test and the studentized bootstrap (block 1), at 5% and 10%; report the bootstrap interval.
  3. Fill a decision table (test $\times$ level).
- **Interpretation**
  - GMV wins on volatility, HRP on Sharpe ratio: which criterion fits which investor?
  - what would a comparison at equal volatility require?

Which Firms Are Compared in Extension C3?

JPM / BAC	JPMorgan Chase / Bank of America
GS / MS	Goldman Sachs / Morgan Stanley
WFC / C	Wells Fargo / Citigroup
NVDA / AAPL	Nvidia / Apple
MSFT / AMZN	Microsoft / Amazon
GOOGL / GME	Alphabet / GameStop
CSCO / MSTR	Cisco Systems / MicroStrategy

C3 [Proposed]: The Implementable Frontier on Our Data — Question

- **Question**: does the cost-aware ranking of Jensen, Kelly, Malamud & Pedersen (2026) (Table 2; the case study of Chapter 4) hold on a small universe?
- **A research project**, an alternative to C2; starter code: `jkmp_replication` in the Quantlet `MFM_ch4_implementable_frontier`
- **Notation**
  - π_t: risky weights; trade $\pi_t - g_t\pi_{t-1}$, with g_t the drift of last month's weights; cost $\frac{W_t}{2}(\pi_t - g_t\pi_{t-1})^\top \mathbf{\Lambda}(\pi_t - g_t\pi_{t-1})$, $\Lambda_i = 0.2/V_i$, V_i the 6-month average daily dollar volume
  - Utility per month: net excess return $-\frac{\gamma}{2}\times$ its variance, $\gamma = 10$; the paper's rule trades part of the way (m) towards an aim portfolio learned from signals
- **Data**: JPM (JPMorgan Chase), BAC (Bank of America), C (Citigroup), GS (Goldman Sachs), MS (Morgan Stanley), WFC (Wells Fargo), NVDA (NVIDIA), AAPL (Apple), MSFT (Microsoft), AMZN (Amazon), GOOGL (Alphabet, Class A), GME (GameStop), CSCO (Cisco Systems), MSTR (Strategy, formerly MicroStrategy) and the nine sector ETFs, prices aligned on common days (from 19 August 2004); T-bill of the setup
- Signals at each month end, ranked into $[0, 1]$: $\ln(P_t/P_{t-1})$, $\ln(P_{t-1}/P_{t-6})$, $\ln(P_{t-1}/P_{t-12})$, 1-year volatility of daily returns, log 6-month dollar volume

Back Exercise list

C3 [Proposed]: The Implementable Frontier on Our Data — Tasks

- **Timeline**: signals and covariances use only past data; covariance windows expand from August 2004 until they reach 10 years (August 2014); hyperparameters chosen each year on validation years from 2011 only; first test weights end of January 2014, test returns February 2014 – August 2026
- **Tasks**
  1. Baseline first: build Portfolio-ML with the fixed warm-up choice $\lambda = 1$, $p = 64$, $\eta = e^{-3}$ (the next task tunes the grid); check one month's cost step by step (W_t, trade vector, $\mathbf{\Lambda}$).
  2. Tune on the grid of the paper (Table 1) and report, for Portfolio-ML, 1/N, minimum variance, Markowitz-ML and Static-ML: return, volatility, gross and net Sharpe, cost, utility, turnover, leverage (Table 2 layout).
  3. Probability of a higher utility, $\Phi(\bar{d}/(s_d/\sqrt{T}))$ for the monthly utility differences d_t: a Normal i.i.d. approximation, conditional on the estimated procedure (not an exact posterior probability); add a 95% HAC interval for $\bar{d}$.
  4. Extension: wealth \$1, 10, 100 billion $\times\ \gamma \in \{1, 5, 10, 20, 100\}$; plot the implementable frontiers with 1/N's trading costs computed at each wealth level.
- Monthly utility: $u_t = r_t^{net} - \frac{\gamma}{2}\frac{T}{T-1}(r_t^{net} - \bar{r}^{net})^2$, using the full evaluation-period mean, so the mean of u_t is the reported utility (variance with $T-1$); $d_t = u_{1t} - u_{2t}$. Multiplying every flow by 12 leaves the displayed probability unchanged.
- **Report**: two tables, the frontier figure, a Quantlet folder; separate what follows the paper from any simplification you made
- **Interpretation**
  - if Portfolio-ML has a lower net Sharpe ratio than 1/N but a higher utility, which criterion matches the investor with $\gamma = 10$?
  - for which γ would the ranking reverse?

Back Exercise list

Acronyms used in this material (1/2)

- **AAPL**: Apple Inc. (Nasdaq ticker)
- **AI**: Artificial Intelligence
- **AMZN**: Amazon.com, Inc. (Nasdaq ticker)
- **BAC**: Bank of America Corporation (NYSE ticker)
- **BET**: Bucharest Exchange Trading (index)
- **BET-TR**: BET Total Return (index)
- **BNR**: Banca Națională a României — *the National Bank of Romania*
- **BVB**: Bursa de Valori București — *the Bucharest Stock Exchange*
- **CSCO**: Cisco Systems, Inc. (Nasdaq ticker)
- **ERC**: Equal Risk Contribution portfolio
- **ETF**: Exchange-Traded Fund
- **GLS**: Generalised Least Squares
- **GME**: GameStop Corp. (NYSE ticker)
- **GMV**: Global Minimum-Variance portfolio
- **GMV-LW**: Global Minimum-Variance portfolio with Ledoit–Wolf covariance
- **GMV-NL**: Global Minimum-Variance portfolio with Nonlinear shrinkage covariance
- **GOOGL**: Alphabet Inc., Class A shares (Nasdaq ticker)

Acronyms used in this material (2/2)

- **GS**: The Goldman Sachs Group, Inc. (NYSE ticker)
- **HAC**: Heteroskedasticity and Autocorrelation Consistent (standard errors)
- **HRP**: Hierarchical Risk Parity
- **HYG**: iShares iBoxx High Yield Corporate Bond ETF
- **IEF**: iShares 7–10 Year Treasury Bond ETF
- **JPM**: JPMorgan Chase & Co. (NYSE ticker)
- **LQD**: iShares iBoxx Investment Grade Corporate Bond ETF
- **MS**: Morgan Stanley (NYSE ticker)
- **MSFT**: Microsoft Corporation (Nasdaq ticker)
- **MSTR**: Strategy Inc., formerly MicroStrategy (Nasdaq ticker)
- **MV**: Mean–Variance tangency portfolio estimated from sample moments
- **MV-LO**: Mean–Variance tangency portfolio, Long-Only
- **NVDA**: NVIDIA Corporation (Nasdaq ticker)
- **OLS**: Ordinary Least Squares
- **OOS**: Out-Of-Sample
- **SE**: Standard Error
- **US**: United States
- **WFC**: Wells Fargo & Company (NYSE ticker)

References (1/3)

- Black, F., Litterman, R. (1992). Global portfolio optimization. *Financial Analysts Journal*, 48(5), 28–43.
- Britten-Jones, M. (1999). The sampling error in estimates of mean-variance efficient portfolio weights. *Journal of Finance*, 54(2), 655–671.
- DeMiguel, V., Garlappi, L., Uppal, R. (2009). Optimal versus naive diversification: how inefficient is the 1/N portfolio strategy?. *Review of Financial Studies*, 22(5), 1915–1953.
- Franke, J., Härdle, W.K., Hafner, C.M. (2019). Statistics of Financial Markets: An Introduction, 5th ed. Springer.
- He, G., Litterman, R. (2002). The intuition behind Black-Litterman model portfolios. SSRN Working Paper 334304.
- Jensen, T.I., Kelly, B., Malamud, S., Pedersen, L.H. (2026). Machine learning and the implementable efficient frontier. *Review of Financial Studies*, 39(10), 3035–3078.
- Jobson, J.D., Korkie, B.M. (1981). Performance hypothesis testing with the Sharpe and Treynor measures. *Journal of Finance*, 36(4), 889–908.

References (2/3)

- Kan, R., Zhou, G. (2007). Optimal portfolio choice with parameter uncertainty. *Journal of Financial and Quantitative Analysis*, 42(3), 621–656.
- Künsch, H.R. (1989). The jackknife and the bootstrap for general stationary observations. *Annals of Statistics*, 17(3), 1217–1241.
- Laloux, L., Cizeau, P., Bouchaud, J.-P., Potters, M. (1999). Noise dressing of financial correlation matrices. *Physical Review Letters*, 83(7), 1467–1470.
- Ledoit, O., Wolf, M. (2004). Honey, I shrunk the sample covariance matrix. *Journal of Portfolio Management*, 30(4), 110–119.
- Ledoit, O., Wolf, M. (2008). Robust performance hypothesis testing with the Sharpe ratio. *Journal of Empirical Finance*, 15(5), 850–859.
- Ledoit, O., Wolf, M. (2017). Nonlinear shrinkage of the covariance matrix for portfolio selection: Markowitz meets Goldilocks. *Review of Financial Studies*, 30(12), 4349–4388.
- Ledoit, O., Wolf, M. (2020). Analytical nonlinear shrinkage of large-dimensional covariance matrices. *Annals of Statistics*, 48(5), 3043–3065.

References (3/3)

- López de Prado, M. (2016). Building diversified portfolios that outperform out of sample. *Journal of Portfolio Management*, 42(4), 59–69.
- Maillard, S., Roncalli, T., Teïletche, J. (2010). The properties of equally weighted risk contribution portfolios. *Journal of Portfolio Management*, 36(4), 60–70.
- Markowitz, H. (1952). Portfolio selection. *Journal of Finance*, 7(1), 77–91.
- Pele, D.T., Wesselhöfft, N., Härdle, W.K., Kolossiatis, M., Yatracos, Y.G. (2023). Are cryptos becoming alternative assets?. *The European Journal of Finance*, 29(10), 1064–1105.
- Petukhina, A., Trimborn, S., Härdle, W.K., Elendner, H. (2021). Investing with cryptocurrencies – evaluating their potential for portfolio allocation strategies. *Quantitative Finance*, 21(11), 1825–1853.
- Theil, H., Goldberger, A.S. (1961). On pure and mixed statistical estimation in economics. *International Economic Review*, 2(1), 65–78.
- Data: Kenneth R. French Data Library; BNR reference exchange rates, www.bnr.ro.