








Modelling Financial Markets

Seminar 3: Factor Models and Asset Pricing — 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

From the Question to the Applications: What We Aim For

- Questions of the seminar
  - **does beta explain the differences in average returns across assets?**
  - **how precise are the estimates?**
- Deliverables
  - reproduce the solved A/B exercises and interpret the result
  - Proposed exercises: submit the requested derivation, table or figure
  - the corrected C1 code
  - one interpretation sentence per exercise
- Seminar notebook: chapter3_seminar_notebook.ipynb (Colab); prices: data/market
  - each task names its notebook section
- For exercises A1–C1 you need regression, CAPM (Capital Asset Pricing Model), inference and multiple testing
- Principal component analysis (PCA), the Dimson beta, Fama–MacBeth regressions and generalised least squares (GLS) serve the after-lecture exercises (B3, B4) and the extension packs in the map
- Seminar: Prof. Daniel Traian Pele, danpele@ase.ro; Antoaneta Amza, doctoral (PhD) candidate, antoaneta.amza@csie.ase.ro

What Do the Factor Series Represent?

ETF	Exchange-traded fund: a fund traded on an exchange; SPY tracks Standard & Poor's 500 (S&P 500).
RF	Risk-free: the risk-free rate.
Mkt-RF	Market minus risk-free: the US stock market return minus RF.
B/M	Book-to-market: book equity / market equity.
SMB	Small minus big: small firms minus large firms.
HML	High minus low: high minus low B/M.
RMW	Robust minus weak: strong minus weak profitability.
CMA	Conservative minus aggressive: conservative minus aggressive investment.
MOM	Momentum: past winners minus past losers.
CAPM	Capital Asset Pricing Model: the single-factor model, Mkt-RF.
FF3, FF5	Fama–French: models with three and five factors, respectively.

How Do We Read the Bucharest Stock Exchange Symbols?

BVB (Bursa de Valori București): the Bucharest Stock Exchange; BET (Bucharest Exchange Trading): the 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
H2O	Hidroelectrica	EL	Electrica
TEL	Transelectrica	FP	Fondul Proprietatea
DIGI	Digi Communications		

Data Used

Series	Source, price used	Frequency	Sample
9 sector ETFs (XLB, XLE, XLF, XLI, XLK, XLP, XLU, XLV, XLY) and SPY	adjusted close	month-end	Jan 1999 – Jul 2026
11 sector ETFs (the 9 + XLRE, XLC), SPY	adjusted close	daily	20 Jun 2018 – 18 Sep 2026
Mkt-RF, SMB, HML, RMW, CMA, MOM, RF	Kenneth French Data Library, percent	monthly, daily	Jul 1963 – Jul 2026
25 size × B/M portfolios	Kenneth French Data Library, value-weighted	monthly	Jul 1963 – Jul 2026
30 industry portfolios, 100 size × B/M portfolios, single-sort portfolios	Kenneth French Data Library, value-weighted	monthly	Jul 1963 – Jul 2026
7 factor ETFs (MTUM, VLUE, QUAL, USMV, IWM, IWD, IWF)	adjusted close	month-end; daily (MTUM, QUAL)	2013 – Jul 2026
11 BVB blue chips, BET	stocks: adjusted close; BET: close	daily	6 Jan 2015 (or listing) – 18 Sep 2026

- Daily prices of the course data, one series per asset (field used: adjusted close for ETFs and stocks, close for BET)
- Multi-asset analyses: align the *prices* on common days first, then take returns; monthly returns from month-end prices
- Sectors (the Select Sector SPDR funds): XLB (materials), XLE (energy), XLF (financials), XLI (industrials), XLK (technology), XLP (consumer staples), XLU (utilities), XLV (health care), XLY (consumer discretionary), XLRE (real estate), XLC (communication services)

Setup: From Prices to Excess Returns

```
# notebook section "Setup" defines
# prices(), factors(), monthly_excess()
ex, F = monthly_excess(['XLK.US',
                        'SPY.US'])
# month-end adjusted prices -> simple
# returns -> minus RF -> align with factors
print(len(ex), ex.index[0].date(),
      ex.index[-1].date())
print(ex['XLK.US'].head(3))
print(F[['Mkt-RF', 'RF']].head(3))
```

- Expected: $T = 331$ months, 1999-01-31 to 2026-07-31
- If the count differs, check the field, the month-end rule and the division by 100

	Jan 1999	Feb 1999	Mar 1999
P_{t-1}, XLK	12.1103	14.0359	12.6439
P_t, XLK	14.0359	12.6439	13.5835
$R_t = P_t/P_{t-1} - 1$	15.90%	−9.92%	7.43%
RF (/100)	0.35%	0.35%	0.43%
$R^e_t = R_t - \mathrm{RF}_t$	15.55%	−10.27%	7.00%
Mkt-RF	3.47%	−4.08%	3.48%

- Worked line: $14.0359/12.1103 - 1 = 15.90\%$; subtracting 0.35% gives 15.55%
- RF (the risk-free rate) is already a rate, and Mkt-RF is already an excess return

Returns and Portfolios

- **Simple return** and **excess return** of asset i in period t

$$R_{i,t} = \frac{P_{i,t}}{P_{i,t-\mathbf{1}}} - 1, \qquad R^e_{i,t} = R_{i,t} - R_{f,t}$$

  - $P_{i,t}$: the adjusted price (dividends reinvested); $R_{f,t}$: the risk-free rate of the period (RF)
  - f_t: a factor, an excess return itself (Mkt-RF unless stated); R^e_t: the $N \times 1$ vector of the N assets

- **Portfolio return**: weights w_i applied to *simple* returns

$$R_{p,t} = \sum_{i=\mathbf{1}}^{N} w_i R_{i,t}$$

  - log returns $r_t = \ln(1 + R_t)$ add over *time*, simple returns add over *assets* (Chapter 0)
  - long–short: $+1/n$ on n assets bought, $-1/n$ on n assets sold short

- Annualising: $\times 12$ (means), $\times\sqrt{12}$ (volatilities and Sharpe ratios); daily 252

Moments, Matrices, Cross-Sections

- Expected excess returns and covariance of the N assets

$$\boldsymbol{\mu} = \mathbb{E}[\mathrm{R}_t^e] \ (N \times 1), \qquad \boldsymbol{\Sigma} = \mathrm{Cov}(\mathrm{R}_t^e) \ (N \times N)$$

  - $\boldsymbol{\Sigma}$ positive definite: every portfolio has positive variance, so $\boldsymbol{\Sigma}^{-1}$ exists
  - 1: the vector of ones; I: the identity matrix; $^\top$: transpose; $\mathrm{a}^\top \boldsymbol{\Sigma}^{-1}\mathrm{a}$: a quadratic form, a squared distance that accounts for correlations

- Maximum-likelihood (ML) moments divide by T

$$\hat{\boldsymbol{\Sigma}} = \frac{1}{T}\hat{\mathrm{E}}^\top\hat{\mathrm{E}}, \qquad \hat{\sigma}_f^2 = \frac{1}{T}\sum_{t=1}^{T}(f_t - \bar{f})^2$$

  - $\hat{\mathrm{E}}$: the $T \times N$ matrix of regression residuals; $\bar{f}$: the time-series mean of the factor; exception: the factor covariance of B8 divides by $T-1$

- **Cross-sectional** moments: across the N assets at a given time, not over time
  - $\mathrm{Var}_{cs}(\hat{\beta}) = \frac{1}{N-1}\sum_i(\hat{\beta}_i - \bar{\beta})^2$ (the unbiased estimator), $\bar{\beta} = \frac{1}{N}\sum_i \hat{\beta}_i$; Cov_{cs} likewise
  - $\otimes$ (Kronecker product, B7): A $\otimes$ B replaces each element a_{jk} of A by the block a_{jk}B

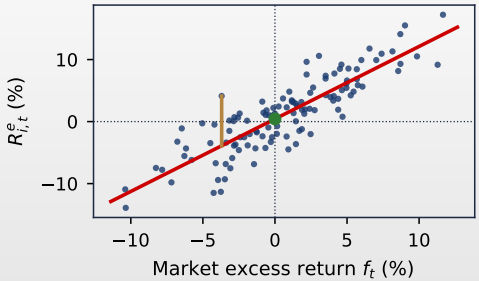
Factor Models

- Details: Chapter 3; textbooks: Campbell, Lo & MacKinlay (1997) and, for principal components, Härdle & Simar (2019)
- **Factor model**: the excess return of an asset is explained by a few common factors

$$R^e_{i,t} = \alpha_i + \beta_{i,\mathbf{1}} f_{\mathbf{1},t} + \cdots + \beta_{i,K} f_{K,t} + \varepsilon_{i,t}$$

  - $f_{k,t}$: factor k, an excess return (a traded portfolio); K: the number of factors
  - $\beta_{i,k}$: the loading (exposure) of asset i on factor k; $\varepsilon_{i,t}$: the idiosyncratic shock, mean 0
  - α_i: the part of the mean return the factors do not explain; a correct model has $\alpha_i = 0$ for every asset

The Market Model and Ordinary Least Squares



- Months $(f_t, R^e_{i,t})$
- OLS: $\hat{\alpha} = 0.43$, $\hat{\beta} = 1.17$, $R^2 = 0.75$
- $\hat{\alpha}$: the line at $f_t = 0$
- one residual $\hat{\varepsilon}_t$

- **Market model**: one factor, f_t = Mkt-RF

$$R^e_{i,t} = \alpha_i + \beta_i f_t + \varepsilon_{i,t}$$

- OLS (ordinary least squares) estimates, minimising $\sum_t \hat{\varepsilon}_t^2$

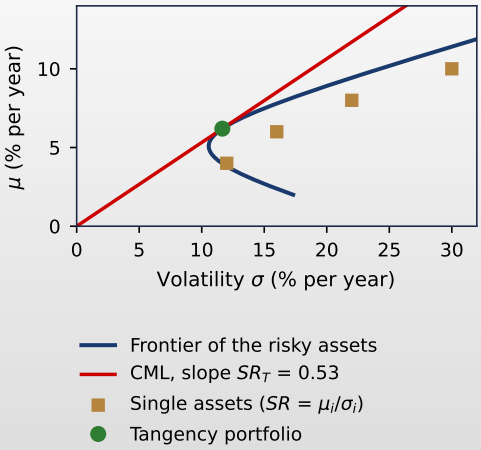
$$\hat{\beta}_i = \frac{\sum_t (f_t - \bar{f})(R^e_{i,t} - \bar{R}^e_i)}{\sum_t (f_t - \bar{f})^2}, \qquad \hat{\alpha}_i = \bar{R}^e_i - \hat{\beta}_i \bar{f}$$

- Reading
  - $\hat{\beta}_i$: the slope; $\beta > 1$: the asset amplifies market moves
  - $\hat{\alpha}_i$: the intercept, the mean excess return in a month with $f_t = 0$
  - $R^2 = 1 - \sum_t \hat{\varepsilon}_t^2 / \sum_t (R^e_{i,t} - \bar{R}^e_i)^2$: the share of the variance explained by the market

The Fama–French Factors

- Factors are long–short portfolios sorted on a characteristic (Fama & French (1993); Fama & French (2015); Carhart (1997)); data: Kenneth French Data Library
  - Mkt-RF is value-weighted; HML (High Minus Low, value factor) sets value against growth stocks; RMW (Robust Minus Weak, profitability factor) sorts on operating profitability, CMA (Conservative Minus Aggressive, investment factor) on asset growth
  - MOM: momentum, the 12–1 signal (slide "Momentum and Look-Ahead Bias")
- The models
  - CAPM: Mkt-RF; FF3 (Fama–French three-factor model): Mkt-RF, SMB (Small Minus Big, size factor), HML; Carhart: FF3 + MOM; FF5 (Fama–French five-factor model): Mkt-RF, SMB, HML, RMW, CMA
  - FF3 regression: $R^e_{i,t} = \alpha_i + b_i\,\text{Mkt-RF}_t + s_i\,\text{SMB}_t + h_i\,\text{HML}_t + \varepsilon_{i,t}$
- Test assets
  - 25 size × B/M portfolios: stocks sorted into 5 size groups × 5 book-to-market groups
    - the sorts are made every June, independently, with NYSE (New York Stock Exchange) quintile breakpoints
    - each intersection is a value-weighted portfolio, held from July to June
    - mean returns differ widely across portfolios: 3.9% a year for small-growth, 13.3% for small-value
  - Sector ETFs hold one sector, factor ETFs follow a factor (MTUM: momentum, QUAL: quality)

Mean–Variance and the Tangency Portfolio

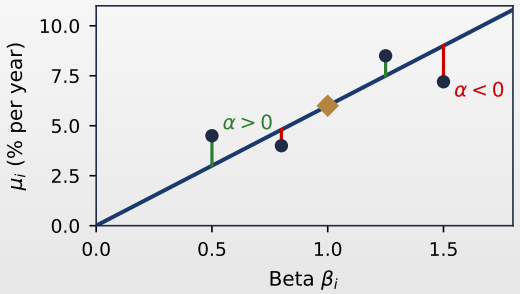


- **Sharpe ratio**: mean excess return per unit of volatility, $SR = \mu/\sigma$
  - mean–variance investors (Markowitz (1952)) with a risk-free asset all hold the risky portfolio with the highest SR: the **tangency portfolio**
- For positive definite $\boldsymbol{\Sigma}$ and $1^\top \boldsymbol{\Sigma}^{-1}\boldsymbol{\mu} > 0$, tangency weights and squared Sharpe ratio

$$w_T = \frac{\boldsymbol{\Sigma}^{-1}\boldsymbol{\mu}}{1^\top \boldsymbol{\Sigma}^{-1}\boldsymbol{\mu}}, \qquad SR_T^2 = \boldsymbol{\mu}^\top \boldsymbol{\Sigma}^{-1}\boldsymbol{\mu}$$

- **CML** (capital market line): the risk-free asset mixed with the tangency portfolio
  - chart: four assets, SR from 0.33 to 0.38; $SR_T = 0.53$
  - example: two uncorrelated assets, $\boldsymbol{\mu} = (4\%, 6\%)$, variances 0.04 and 0.09: $w_T = (0.60, 0.40)$, $SR_T^2 = 0.04 + 0.04$, $SR_T = 0.283$

The CAPM and the Security Market Line



SML: $\mu_i = \beta_i \mu_m$, $\mu_m = 6\%$

Assets: average excess return

Market: $\beta = 1$

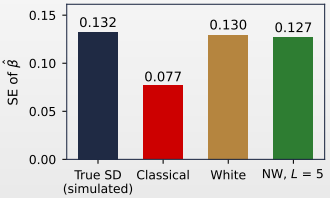
- **CAPM** (Sharpe (1964)): if everyone holds the tangency portfolio, it is the market, and for every asset

$$\mu_i = \beta_i \mu_m, \qquad \beta_i = \frac{\mathrm{Cov}(R_i^e, R_m^e)}{\mathrm{Var}(R_m^e)}$$

  - $\mu_i = \mathbb{E}[R_i^e]$, $\mu_m = \mathbb{E}[R_m^e]$: expected excess returns of the asset and of the market
  - SML (security market line): only market risk is priced; $\varepsilon_{i,t}$ is diversified away
- **Alpha**: the vertical distance to the SML
  - testable consequence: $\alpha_i = 0$ for every asset
  - example: $\beta = 1.2$, premium 6% a year: CAPM gives 7.2%; a realised 9% gives $\alpha = 1.8\%$
- A rejection may mean that the market proxy is not efficient (A2)
  - zero-beta CAPM (Black (1972)): $\mu_i - \gamma = \beta_i(\mu_m - \gamma)$, γ: the mean of a zero-beta portfolio

Classical, White and Newey–West Standard Errors

SD: standard deviation; SE: standard error.



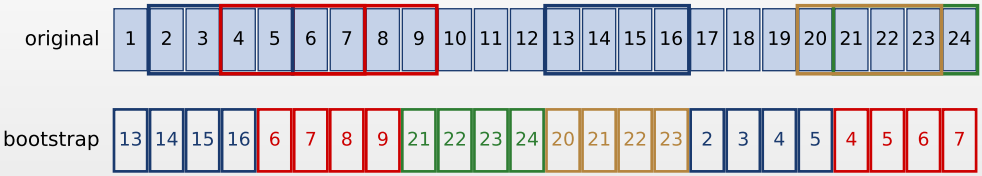
- Chart: 2 000 simulated samples with autocorrelated residuals that are larger in volatile months
  - the classical standard error (0.077) understates the true standard deviation of $\hat{\beta}$ (0.132)
  - the White and Newey–West errors are close to the true value

- OLS of $\hat{\boldsymbol{\theta}} = (\hat{\alpha}, \hat{\beta})^\top$ on $\mathrm{x}_t = (1, f_t)^\top$; X: rows $\mathrm{x}_t^\top$; $\hat{\varepsilon}_t$: residuals

$$\widehat{\mathrm{Var}}(\hat{\boldsymbol{\theta}}) = (\mathrm{X}^\top \mathrm{X})^{-1} \hat{\mathrm{S}} (\mathrm{X}^\top \mathrm{X})^{-1}$$

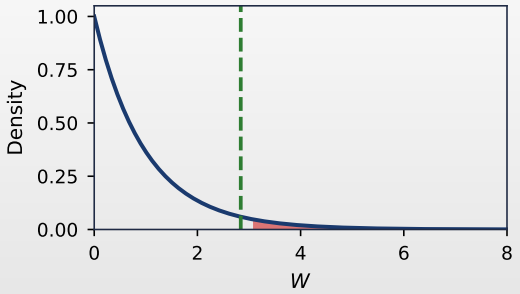
- Choices of $\hat{\mathrm{S}}$
  - classical: $\hat{\sigma}_\varepsilon^2 \mathrm{X}^\top \mathrm{X}$ (constant variance, no autocorrelation)
  - White (White (1980)): $\sum_t \hat{\varepsilon}_t^2 \mathrm{x}_t \mathrm{x}_t^\top$, robust to heteroskedasticity
  - Newey–West (NW; Newey & West (1987)), a heteroskedasticity and autocorrelation consistent (HAC) estimator: adds $\sum_{l=1}^{L} \left(1 - \frac{l}{L+1}\right) \sum_t \hat{\varepsilon}_t \hat{\varepsilon}_{t-l} (\mathrm{x}_t \mathrm{x}_{t-l}^\top + \mathrm{x}_{t-l} \mathrm{x}_t^\top)$
  - lag rule $L = \lfloor 4(T/100)^{2/9} \rfloor$; $T = 331$: $L = 5$

Bootstrap Intervals



- **Moving-block bootstrap** (Künsch (1989)): draw blocks of b consecutive months (here $b = 4$) and concatenate them to length T
  - resample the pairs (R^e_t, f_t) together; re-estimate; repeat B times; the 2.5% and 97.5% quantiles form a 95% percentile interval
  - the dependence inside each block (autocorrelation, volatility clustering) is kept; block starts may overlap
- Variants
  - **stationary bootstrap** (Politis & Romano (1994)): random block lengths, geometric with mean b; **i.i.d.** bootstrap (independent and identically distributed draws): single months ($b = 1$)
  - bootstrap *under* H_0: rebuild data with the null imposed; $p = (1 + \#\{W^* \geq W\})/(B + 1)$

The GRS Test



- **GRS** (Gibbons–Ross–Shanken test, Gibbons, Ross & Shanken (1989)): H_0: $\alpha_1 = \cdots = \alpha_N = 0$ jointly

$$W = \frac{T-N-1}{N} \cdot \frac{\hat{\boldsymbol{\alpha}}^\top \hat{\boldsymbol{\Sigma}}^{-1} \hat{\boldsymbol{\alpha}}}{1 + \hat{\mu}_m^2 / \hat{\sigma}_m^2} \sim F(N, T-N-1)$$

- Symbols
  - N: number of test assets; T: number of months; $\hat{\boldsymbol{\alpha}}$: the N alphas; $\hat{\boldsymbol{\Sigma}}$: ML residual covariance
  - $SR_m = \hat{\mu}_m / \hat{\sigma}_m$: the factor's Sharpe ratio (ML); F: the Fisher distribution, exact for i.i.d. Normal errors
- Reading
  - large W: the alphas are far from 0 relative to their noise; reject when W exceeds the 95% quantile of F
  - example: $T = 121$, $N = 2$, $\hat{\boldsymbol{\alpha}}^\top \hat{\boldsymbol{\Sigma}}^{-1} \hat{\boldsymbol{\alpha}} = 0.05$, $SR_m = 0.2$: $W = 59 \times 0.05 / 1.04 = 2.84 < 3.07$, $p = 0.062$

Normality of the Residuals

- The F distribution of GRS assumes residuals with the Normal distribution (B7)
  - heavy tails change the distribution of W; the bootstrap under H_0 (slide "Bootstrap Intervals") does not require the Normal distribution, but assumes i.i.d. residuals independent of the factor
- **Pearson kurtosis** of a residual series

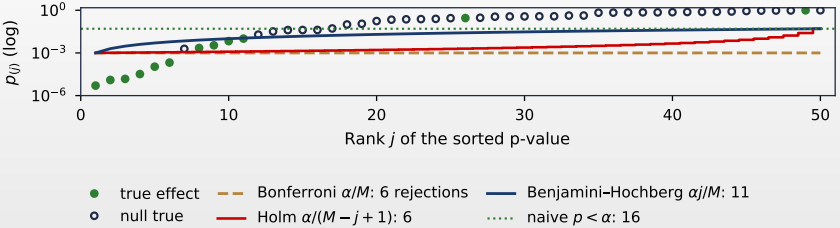
$$\text{Kurt} = \frac{\frac{1}{T}\sum_t \hat{\varepsilon}_t^4}{\hat{\sigma}^4}, \qquad \hat{\sigma}^2 = \frac{1}{T}\sum_t \hat{\varepsilon}_t^2$$

  - Kurt $= 3$ for the Normal distribution; Kurt > 3: heavy tails, extreme months more frequent
- **Q–Q plot** (quantile–quantile plot): sorted standardised residuals against Normal quantiles
  - points on the 45-degree line: Normal; an S shape: heavy tails (Seminar 1)
- **Wald test**, general form: a vector of estimates $\hat{\boldsymbol{\theta}}$ with covariance $\hat{\mathrm{V}}$

$$W = \hat{\boldsymbol{\theta}}^\top \hat{\mathrm{V}}^{-1} \hat{\boldsymbol{\theta}} \to \chi^2_q \text{ under } H_0 : \boldsymbol{\theta} = 0$$

  - q: the number of restrictions; χ^2_q: the chi-squared distribution with q degrees of freedom
- **What do you think**: two alphas have $t_1 = 1.5$ and $t_2 = -1.5$, and the residuals of the two assets have correlation $\rho = 0.9$; does the joint Wald test reject H_0: $\alpha_1 = \alpha_2 = 0$ at 5%?
  - **Answer**: yes: $W = (t_1^2 + t_2^2 - 2\rho\, t_1 t_2)/(1-\rho^2) = 8.55/0.19 = 45.0$, far above 5.99, the 95% quantile of χ^2_2
  - two highly correlated assets with alphas of opposite signs: the difference of the alphas has a small standard error

Multiple Testing, FWER and FDR



- With M tests some p-values are small by chance (hundreds of published factors, Harvey, Liu & Zhu (2016))
  - **FWER** (family-wise error rate): the probability of at least one false rejection
  - Bonferroni controls the FWER by rejecting when $p_i \leq \alpha/M$
  - Holm (Holm (1979)) compares, in order, $p_{(j)}$ with $\alpha/(M-j+1)$ and stops at the first non-rejection
  - **FDR** (false discovery rate): the expected share of false discoveries among the rejections
  - **Benjamini–Hochberg** (BH; Benjamini & Hochberg (1995)): find the largest j with $p_{(j)} \leq \alpha j/M$ and reject $H_{(1)}, \ldots, H_{(j)}$; valid for independent or positively dependent tests
  - Benjamini–Yekutieli (BY; Benjamini & Yekutieli (2001)): the BH thresholds divided by $\sum_{i=1}^{M} 1/i$; valid under any dependence
  - chart: 12 true effects among 50 tests; BH reports 11 discoveries (one false), Holm 6

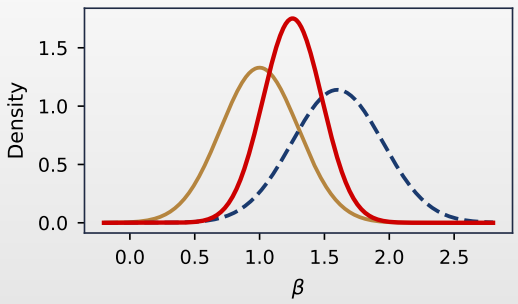
Momentum and Look-Ahead Bias

- **Momentum** (Jegadeesh & Titman (1993)): rank stocks on the past return, buy the winners, sell the losers
  - 12–1 signal: the compound return of months $t-12, \ldots, t-2$, known at the end of $t-1$

$$s_{i,t} = \prod_{k=2}^{12} (1 + R_{i,t-k}) - 1$$

  - month $t-1$ is skipped because of **short-term reversal**: at a one-month horizon, returns reverse (Jegadeesh (1990))
  - the MOM factor (Kenneth French Data Library): a 2×3 sort re-formed monthly, on size (NYSE median) and on $s_{i,t}$ (NYSE 30th and 70th percentiles); MOM $= \frac{1}{2}(S_W + B_W) - \frac{1}{2}(S_L + B_L)$, with S/B small/big firms and W/L winners/losers
  - the portfolio is held during month t; its return averages the simple returns of month t (Returns and Portfolios)
- **Look-ahead bias**: the signal uses information not yet available at formation
  - test: the weights must not change if any price after the end of $t-1$ changes
- Other pitfalls of a backtest (C1, C2)
  - **survivorship bias**: today's blue chips are chosen with hindsight; losers that left the index are missing
  - short-sale restrictions, costs and liquidity reduce what can be earned; with 3 stocks per side, one stock can drive the result of a month

Bayesian Shrinkage of Beta

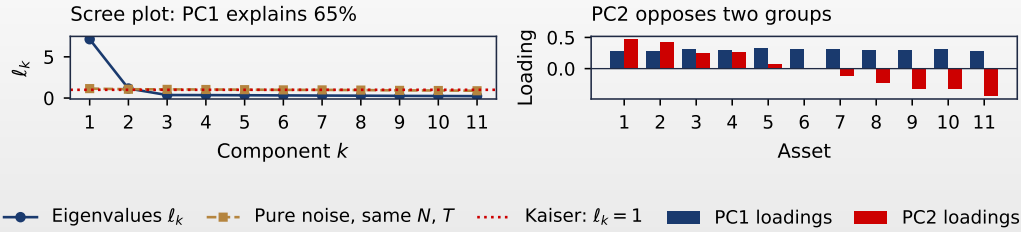


- Prior $\beta_i \sim N(\bar{\beta}, \sigma_{\beta}^2)$; estimate $\hat{\beta}_i \mid \beta_i \sim N(\beta_i, \text{se}_i^2)$
  - $\bar{\beta}$, σ_{β}^2: mean and true dispersion of the betas across assets; se_i: the standard error of $\hat{\beta}_i$
- **Vasicek** (Vasicek (1973)): the posterior mean shrinks $\hat{\beta}_i$ toward $\bar{\beta}$

$$\beta_i^V = w_i \hat{\beta}_i + (1 - w_i)\bar{\beta}, \qquad w_i = \frac{\sigma_{\beta}^2}{\sigma_{\beta}^2 + \text{se}_i^2}$$

- Reading and estimation
  - noisy $\hat{\beta}_i$ (large se_i): small w_i, strong shrinkage; chart: $w = 0.42$
  - law of total variance: $\text{Var}(\hat{\beta}) = \text{Var}(\mathbb{E}[\hat{\beta} \mid \beta]) + \mathbb{E}[\text{Var}(\hat{\beta} \mid \beta)]$, used to estimate σ_{β}^2 (A3)
  - **Blume** (Blume (1971)): betas regress toward 1; the practitioners' rule $0.33 + 0.67\hat{\beta}_i$ is a fixed pull toward 1

Principal Components



- PCA (principal component analysis); the first and second principal components (PC1, PC2)
- Standardise each return series (Z: $T \times N$), then decompose the correlation matrix

$$\mathrm{C} = \tfrac{1}{T}\mathrm{Z}^\top \mathrm{Z} = \sum_{k=\mathbf{1}}^{N} \ell_k \boldsymbol{\gamma}_k \boldsymbol{\gamma}_k^\top, \qquad \ell_{\mathbf{1}} \geq \cdots \geq \ell_N, \quad \textstyle\sum_k \ell_k = N$$

- ℓ_k: the eigenvalue, equal to the variance of component k (written ℓ so that it is not confused with the risk premium λ); the component explains ℓ_k / N of the total variance
- $\boldsymbol{\gamma}_k$: the eigenvector of norm 1 (the loadings); $\mathrm{Z}\boldsymbol{\gamma}_k$: the component scores, i.e. its time series
  - chart: 11 simulated assets, a market factor and a weaker one; PC1: all loadings positive; PC2 opposes two groups

How Many Factors?

- **Kaiser** (Kaiser (1960)): keep the components with $\ell_k > 1$ (more variance than one standardised asset)
- **Bai–Ng** (Bai & Ng (2002)): penalised information criterion

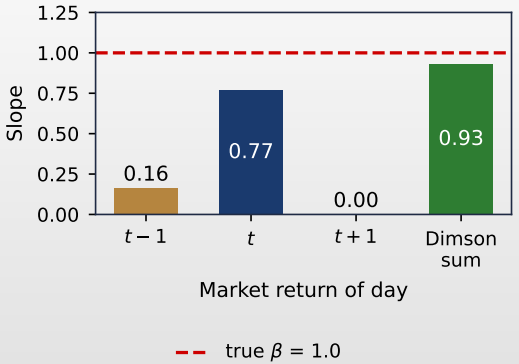
$$IC_{p\mathbf{2}}(k) = \ln V(k) + k\,\frac{N+T}{NT}\ln\min(N,T), \qquad V(k) = \frac{1}{N}\sum_{j>k}\ell_j$$

  - $V(k)$: the variance left after k factors; choose $\hat{k} = \arg\min_{\mathbf{0}\le k\le k_{\max}} IC_{p\mathbf{2}}(k)$
  - designed for large N and T; with $N = 11$ the penalty is weak and $\hat{k}$ tends to $k_{\max}$
- **Ahn–Horenstein** (Ahn & Horenstein (2013)): eigenvalue ratio

$$\hat{k} = \arg\max_{\mathbf{1}\le k\le k_{\max}} \frac{\ell_k}{\ell_{k+\mathbf{1}}}$$

  - the largest drop in the scree plot; in the chart on slide "Principal Components", at $k = 1$
- Pure noise (amber in the chart on slide "Principal Components") has all eigenvalues near 1: compare with it before naming a factor

Stale Prices and the Dimson Beta



- **Stale prices**: an illiquid stock does not trade every day, so its price reacts to market news with a delay
  - the same-day OLS beta is biased toward 0
- **Dimson beta** (Dimson (1979)): add the lagged and the leading market return and sum the slopes

$$r_{i,t} = a + b_{-\mathbf{1}} r_{m,t-\mathbf{1}} + b_{\mathbf{0}} r_{m,t} + b_{+\mathbf{1}} r_{m,t+\mathbf{1}} + e_t$$

$$\beta^D = b_{-\mathbf{1}} + b_{\mathbf{0}} + b_{+\mathbf{1}}$$

- Chart (simulated, true beta 1, the stock trades on 75% of days)
  - same-day slope 0.77; the lag adds 0.16; Dimson 0.93, close to 1
  - the lead and lag must be formed on the full calendar, before any day is removed

Exercise Map

Exercise	Question	Type	Model
A1	how many of ten candidate factors are discoveries?	Solved	–
A2	why does the tangency portfolio price every asset by beta?	Solved	–
B1	how precisely are XLK's beta and alpha estimated?	Solved	–
B2	do nine sector ETFs reject the CAPM?	Solved	A2, B1
C1	what is wrong in an AI-written momentum backtest?	Proposed	B1, primer
After the lecture	A3 [Solved], B3 [Solved], B4 [Proposed]: Vasicek beta, PCA and the number of factors, BVB betas	Solved; B4: Proposed	B1, A3
Extension packs	A2 Extended, A4–A8, B5–B11, C2, C3 (optional, after the seminar)	Solved or Proposed (marked on each)	named 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: estimation and inference on real data, each ending with an interpretation question; Part C: open questions and the critique of an artificial intelligence (AI) answer
- The Newey–West lag count stated in each exercise is the replication convention, not the optimal choice for every sample

A1 [Solved]: Ten Candidate Factors

- **Question**: how many of ten new factors are discoveries once we account for testing ten at a time?
- **Inputs**: a new example (not the vector of Chapter 3); level $\alpha = 5\%$, $M = 10$ two-sided tests of zero mean
- **Tasks**: complete the worksheet
  1. Count the naive discoveries $p \le 0.05$ and the Bonferroni discoveries $p \le 0.005$.
  2. Apply Holm: compare $p_{(j)}$ with $0.05/(M - j + 1)$ and stop at the first failure.
  3. Apply Benjamini–Hochberg (BH): find the largest j with $p_{(j)} \le 0.005j$ and reject $H_{(1)}, \ldots, H_{(j)}$.
  4. State which error rate each procedure controls and under which dependence; as an extension, also apply Benjamini–Yekutieli (BY).

j	1	2	3	4	5	6	7	8	9	10
$p_{(j)}$	0.001	0.004	0.009	0.012	0.021	0.035	0.048	0.060	0.20	0.55
Holm $0.05/(M - j + 1)$	0.0050	...								
BH $0.005j$	0.005	...								

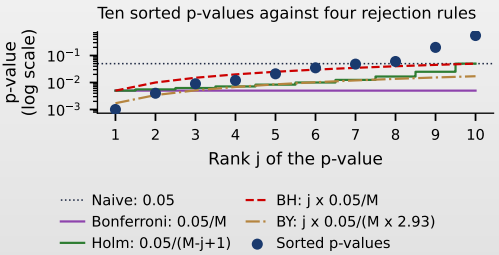
- **Report**: the completed worksheet and the four counts of discoveries
- Notebook: section A1

Back Exercise Map

A1 [Solved]: Solution

j	1	2	3	4	5	6	7	8	9	10
$p_{(j)}$	0.001	0.004	0.009	0.012	0.021	0.035	0.048	0.060	0.20	0.55
Holm	**0.0050**	**0.0056**	0.0063	0.0071	0.0083	0.0100	0.0125	0.0167	0.025	0.05
BH	**0.005**	**0.010**	**0.015**	**0.020**	**0.025**	0.030	0.035	0.040	0.045	0.050

- Naive **7**; Bonferroni **2**; Holm: $0.009 > 0.0063$ is the first failure, so **2**
- BH: $j = 5$ is the largest rank with $p_{(j)} \leq 0.005j$ ($0.021 \leq 0.025$), so **5**
  - for $j = 6, \ldots, 10$: $0.035 > 0.030$; $0.048 > 0.035$; $0.060 > 0.040$; $0.20 > 0.045$; $0.55 > 0.050$
- Error rate controlled
  - Bonferroni and Holm: FWER, under any dependence
  - BH: FDR, under independence or positive dependence
  - BY (the BH thresholds divided by $\sum_{i=1}^{10} 1/i = 2.93$): FDR, under any dependence; **1** discovery
- Interpretation: a factor search tolerates a known share of false factors, so FDR is the relevant target



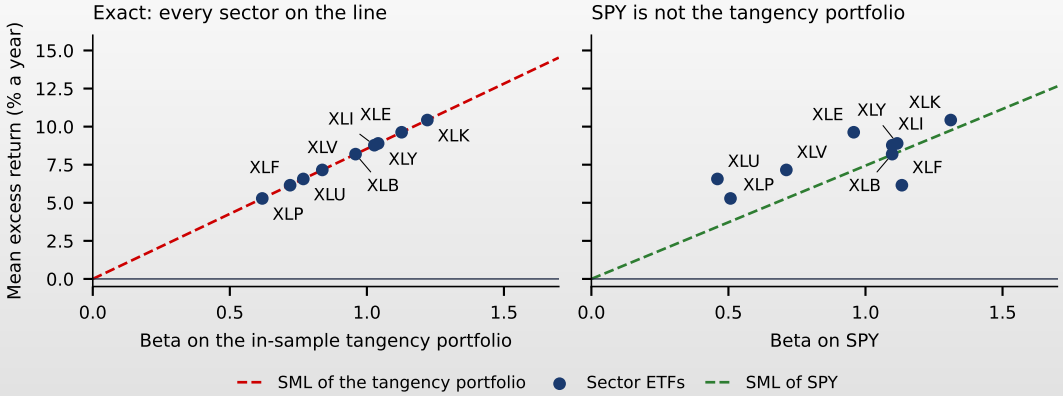
MFM_ch3_seminar

A2 [Solved]: The CAPM from the Tangency Portfolio

- **Question**: if the market is the tangency portfolio, why must every asset satisfy $\mu_i = \beta_i \mu_m$?
- **Inputs**: N risky assets with expected excess returns $\boldsymbol{\mu}$ ($N \times 1$) and positive definite covariance $\boldsymbol{\Sigma}$; a risk-free asset
- **Assumption**: the market is the tangency portfolio $\mathrm{w}_m = \boldsymbol{\Sigma}^{-1}\boldsymbol{\mu}/c$, with the scalar $c = 1^\top \boldsymbol{\Sigma}^{-1}\boldsymbol{\mu} \neq 0$; $R^e_m = \mathrm{w}_m^\top \mathrm{R}^e$
- **Tasks**
  1. Express the covariance vector $\mathrm{Cov}(\mathrm{R}^e, R^e_m) = \boldsymbol{\Sigma}\mathrm{w}_m$ as a function of $\boldsymbol{\mu}$ and c.
  2. Compute $\mu_m = \mathrm{w}_m^\top \boldsymbol{\mu}$ and $\sigma^2_m = \mathrm{w}_m^\top \boldsymbol{\Sigma} \mathrm{w}_m$ and find the relation between them.
  3. Combine steps 1 and 2 to prove that $\mu_i = \beta_i \mu_m$ for every asset, with $\beta_i = \mathrm{Cov}(R^e_i, R^e_m)/\sigma^2_m$.
  4. Check the result on the 9 sector ETFs (monthly, 1999–2026): replace $\boldsymbol{\mu}$, $\boldsymbol{\Sigma}$ by sample moments and compare $\hat{\mu}_i$ with $\hat{\beta}_i \hat{\mu}_m$, first for the sample tangency portfolio, then for SPY (SPDR S&P 500 ETF Trust).
- **Report**: the derivation in three steps, a chart of the two checks and one sentence
- Notebook: section A2

Back Exercise Map

A2 [Solved]: Chart



MFM_ch3_seminar

A2 [Solved]: Steps of the Proof

1. $\boldsymbol{\Sigma} \mathrm{w}_m = \boldsymbol{\mu}/c$, so $\mathrm{Cov}(\mathrm{R}^e, R_m^e) = \boldsymbol{\mu}/c$
2. $\sigma_m^2 = \mathrm{w}_m^\top \boldsymbol{\Sigma} \mathrm{w}_m = \mu_m/c$, hence $c = \mu_m/\sigma_m^2$
3. $\boldsymbol{\mu} = c\ \mathrm{Cov}(\mathrm{R}^e, R_m^e) = \boldsymbol{\beta}\, \mu_m$: the SML
4. Sample tangency portfolio (mean 8.55% a year): largest gap $|\hat{\mu}_i - \hat{\beta}_i \hat{\mu}_m| \approx 3 \times 10^{-17}\%$ a year (numerically zero); SPY: gaps up to 3.1% a year (XLU +3.1%, XLE +2.5%, XLF −2.3%)

⊡ Any efficient portfolio prices all assets exactly; a CAPM test asks whether the market proxy is efficient

B1 [Solved]: Beta and Alpha of XLK

- **Questions**
  - what are the CAPM beta and alpha of XLK, the technology exchange-traded fund (ETF)?
  - how precisely are they known?
- **Data**: XLK and Mkt-RF, monthly excess returns, January 1999 – July 2026 ($T = 331$), as in the Setup slide
- **Tasks**
  1. Estimate $R^e_{\text{XLK},t} = \alpha + \beta\,\text{Mkt-RF}_t + \varepsilon_t$ by OLS; annualise α and its standard errors by multiplying by 12.
  2. Compute classical, White (heteroskedasticity-robust) and Newey–West standard errors (s.e.) of α and β, with $L = 5$ lags.
  3. Test $H_0 : \alpha = 0$ at 5% with the Newey–West t-statistic.
  4. Apply a moving-block bootstrap (Künsch (1989)): resample blocks of 12 consecutive months of the pairs (R^e_t, f_t) and re-estimate α in 2000 draws, with seed 42; report the 95% percentile interval.
  5. Plot the data with the fitted line, the residuals over time and the bootstrap distribution.
- **Report**: a table (estimate, three standard errors), $t(\alpha)$, R^2, the interval and the chart
- **Interpretation**: the HAC error of beta is twice the classical one: which feature of the residuals causes it?
- Notebook: section B1

Back Exercise Map

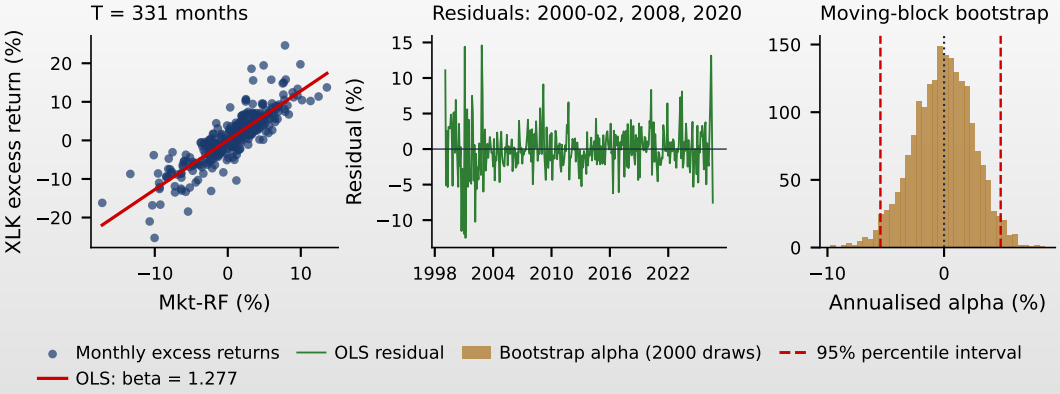
B1 [Solved]: Estimates and Interpretation

	Estimate	Classical s.e.	White s.e.	Newey–West s.e.
α (annualised)	0.33%	2.27%	2.32%	2.45%
β	1.277	0.041	0.054	0.082

1. $R^2 = 0.742$; monthly $\hat{\alpha} = 0.027\%$; annualise the unrounded value: $12 \times 0.0273509\% = 0.328211\% \approx 0.33\%$ a year
2. Lags $L = \lfloor 4 \times 3.31^{2/9} \rfloor = \lfloor 5.2 \rfloor = 5$; robust 95% interval of β: $1.277 \pm 1.96 \times 0.082 = [1.12, 1.44]$ instead of $[1.20, 1.36]$
3. $t(\alpha) = 0.33/2.45 = 0.13$, $p = 0.89$: do not reject H_0; no evidence of a non-zero CAPM alpha
4. Bootstrap: $\alpha \in [-5.5\%, 4.9\%]$, standard deviation of the draws 2.61%; the interval contains 0

- Classical errors assume $\mathrm{Var}(\varepsilon_t \mid f_t)$ constant and no serial correlation
  - the sandwich estimator weights each squared residual by $(f_t - \bar{f})^2$
- White raises the s.e. of β from 0.041 to 0.054: heteroskedasticity
  - residuals are large when the market moves a lot (2000–2002, 2008, 2020)
- Newey–West raises it further to 0.082: serial correlation
  - the products $f_t \varepsilon_t$ are serially correlated, so the cross terms of the HAC estimator are positive
- The point estimates are the same; only the uncertainty changes

B1 [Solved]: Chart



MFM_ch3_seminar

B2 [Solved]: GRS Test on Sector ETFs

- **Question**: do the nine sector ETFs jointly reject the CAPM?
- **Data**: 9 sector ETFs, month-end adjusted prices aligned on common months; monthly excess returns, January 1999 – July 2026 ($T = 331$, $N = 9$); factor Mkt-RF
- **Tasks**
  1. Estimate the 9 time-series regressions on Mkt-RF; report the annual alphas with Newey–West intervals and the betas.
  2. Build the residual matrix $\hat{\mathrm{E}}$ ($T \times 9$), $\hat{\boldsymbol{\Sigma}} = \hat{\mathrm{E}}^{\top}\hat{\mathrm{E}}/T$ and the quadratic form $\hat{\boldsymbol{\alpha}}^{\top}\hat{\boldsymbol{\Sigma}}^{-1}\hat{\boldsymbol{\alpha}}$ with monthly alphas.
  3. Compute $\hat{\mu}_m$, $\hat{\sigma}_m$ (ML), then the GRS statistic W (slide "The GRS Test") and its p-value from $F(9, 321)$.
  4. Repeat the test on the 25 size × B/M portfolios over the same 331 months.
- **Report**: the coefficient table, the three inputs of W, W, p, the chart and the same-period comparison
- **Interpretation**: the 25 portfolios reject the CAPM (GRS = 4.20 over 1963–2026); why do the sectors not reject it?
- Model: A2 (efficiency), B1 (regression); Notebook: section B2

Back Exercise Map

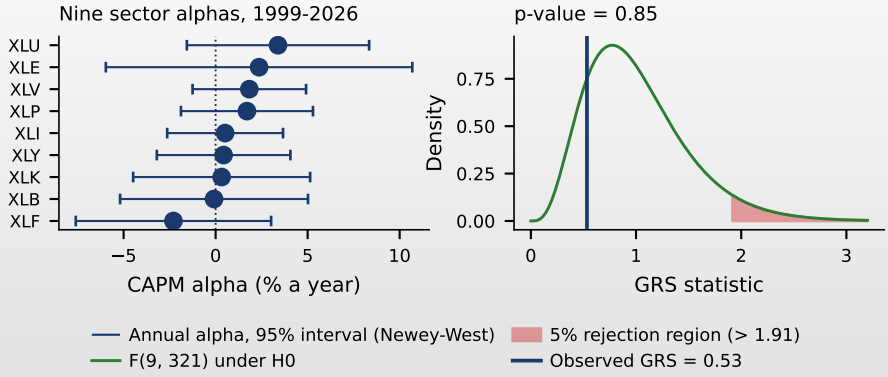
B2 [Solved]: From Residuals to W

	XLB	XLE	XLF	XLI	XLK	XLP	XLU	XLV	XLY
$\hat{\alpha}$ (% a year)	-0.1	2.4	-2.3	0.5	0.3	1.7	3.4	1.8	0.4
NW s.e. (% a year)	2.6	4.2	2.7	1.6	2.5	1.8	2.5	1.6	1.9
$\hat{\beta}$	1.05	0.92	1.07	1.05	1.28	0.45	0.40	0.67	1.07

1. No alpha is significant alone; the largest $|t|$ is XLU's, $3.39/2.53 = 1.34$
2. $\hat{\mathrm{E}}$: 331×9; $\hat{\boldsymbol{\Sigma}}$: 9×9; monthly alphas (%): $-0.007, 0.197, -0.191, 0.043, 0.027, 0.142, 0.282, 0.153, 0.036$; $\hat{\boldsymbol{\alpha}}^\top \hat{\boldsymbol{\Sigma}}^{-1} \hat{\boldsymbol{\alpha}} = 0.0153$
3. $\hat{\mu}_m = 0.659\%$, $\hat{\sigma}_m = 4.51\%$, $SR_m = 0.146$; $W = \frac{321}{9} \times \frac{0.0153}{1+0.146^2} = 35.67 \times 0.0150 = 0.53$, $p = 0.85$
4. Same 331 months, 25 size $\times$ B/M portfolios: GRS $= 2.91$, $p = 9 \times 10^{-6}$

- Interpretation: the non-rejection comes from the test assets, not from the sample; sectors mix small and big, value and growth
- Not rejecting is not accepting the CAPM: choose test assets that spread the characteristics suspected to be priced

B2 [Solved]: Chart and Interpretation



- $W = 0.53$ lies in the body of $F(9, 321)$, far below the 5% critical value 1.91: do not reject H_0: all $\alpha_i = 0$
- Sharpe-ratio form (A4): the sectors raise the monthly Sharpe ratio only from 0.146 to 0.191 (annual 0.51 → 0.66)

MFM_ch3_seminar

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

- **Prompt** sent to an AI assistant: "P holds cleaned month-end adjusted prices of 10 BVB blue chips, one column per stock. Build a 12-month momentum long–short portfolio, top 3 minus bottom 3 with equal weights, and report its annual mean and volatility"
- **Answer**: the code below and the conclusion

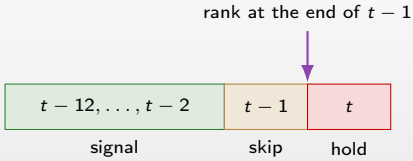
```
R   = np.log(P).diff()                          # monthly returns
mom = P.pct_change(12)                          # past 12-month return
ls  = {}
for t in R.index[13:]:
    s = mom.loc[t].dropna().sort_values()       # rank the stocks
    ls[t] = R.loc[t, s.index[-3:]].mean() - R.loc[t, s.index[:3]].mean()
ls = pd.Series(ls)
print(12 * ls.mean(), np.sqrt(12) * ls.std())  # annual mean, volatility
```

- **Conclusion of the AI**: "Momentum earns about 49% a year with a volatility of 16%. This agrees with Fama & French (2012), who document momentum on the Bucharest Stock Exchange."

Back Exercise Map

C1 [Proposed]: Critique an AI Answer

- **Strategy intended by the student** (the contract)
  - 12–1 momentum signal (slide "Momentum and Look-Ahead Bias"); a stock is eligible if it has a signal and a price at the end of month $t-1$
  - long the top 3, short the bottom 3, equal weights, using the simple returns of month t
- **Data**: 10 blue chips (TLV, SNP, BRD, TGN, FP, SNG, SNN, EL, DIGI, TEL; H2O, listed in July 2023, is excluded), month-end prices, December 2015 – August 2026
  - FP prices are adjusted for the special dividend of 7 September 2023 (RON 1.7225 per share)
- **Tasks**
  1. Audit the timing of the signal, the aggregation of returns across stocks and the citation; list every error you find.
  2. For each error, say how you would detect it: a test, a published number or the paper itself.
  3. Report corrected code, annual mean and volatility, sample size and the mean t-statistic (Newey–West, 3 lags).
  4. Log each error in the format of `AI_ERRORS.md`: request, answer, error, how found, correction.
- **Report**: corrected code, four statistics and `AI_ERRORS.md`



- **Interpretation**
  - does "about 49% a year" survive?
  - at which hurdle is the corrected premium significant: 1.96 (a single test) or 3 (the Harvey, Liu & Zhu (2016) hurdle for new factors)?
- AI code usually runs without an error message: the mistakes are in the finance, not in the syntax
- Model: B1; Notebook: section C1

A3 [Solved]: Vasicek as a Posterior Mean

- **Question**: why does the Vasicek rule shrink a noisy beta towards the cross-sectional mean?
- **Data**: 9 sector ETFs, monthly excess returns on SPY, January 1999 – December 2012; OLS betas with Newey–West errors ($L = 4$ lags, the rule on the slide "Classical, White and Newey–West Standard Errors")
  - XLK: $\hat{\beta} = 1.415$, se $= 0.139$; beta estimated later, on 2013–2026: 1.172
  - across the 9 sectors: $\bar{\beta} = 0.949$, $\mathrm{Var}_{cs}(\hat{\beta}) = 0.1159$, mean squared standard error $\overline{\mathrm{se}^2} = 0.0077$
- **Assumptions**: prior $\beta_i \sim N(\bar{\beta}, \sigma_\beta^2)$; likelihood $\hat{\beta}_i \mid \beta_i \sim N(\beta_i, \mathrm{se}_i^2)$ with se_i known
- **Tasks**
  1. Multiply prior and likelihood, complete the square in β and state the posterior distribution: its mean β_i^V and its variance.
  2. Use the law of total variance to show that $\mathrm{Var}_{cs}(\hat{\beta}) = \sigma_\beta^2 + \mathbb{E}[\mathrm{se}^2]$, and write the empirical-Bayes estimator of σ_β^2.
  3. Compute σ_β^2, the weight w, β^V and the posterior standard deviation for XLK, and the beta adjusted by the practitioners' rule $0.33 + 0.67\hat{\beta}$ (the Blume beta, slide "Bayesian Shrinkage of Beta").
  4. Compare the absolute errors of the raw, Vasicek and Blume betas against the later estimate 1.172.
- **Report**: the posterior, the numbers of step 3 and the three errors
- Notebook: section A3

Back Exercise Map

A3 [Solved]: The Posterior

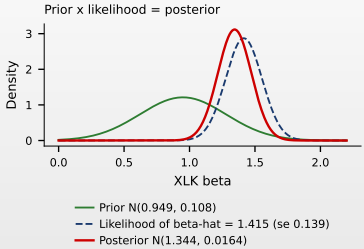
1. Posterior $\propto \exp\{-\frac{(\beta-\bar{\beta})^2}{2\sigma_{\beta}^2}\}\exp\{-\frac{(\hat{\beta}-\beta)^2}{2\,\mathrm{se}^2}\}$; the exponent is

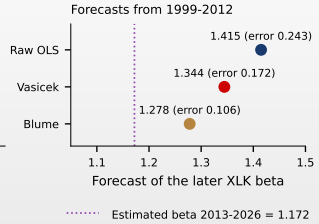
$$-\tfrac{1}{2}\Big[\beta^2\Big(\tfrac{1}{\sigma_{\beta}^2}+\tfrac{1}{\mathrm{se}^2}\Big)-2\beta\Big(\tfrac{\bar{\beta}}{\sigma_{\beta}^2}+\tfrac{\hat{\beta}}{\mathrm{se}^2}\Big)\Big]+\text{const}$$

   - Completing the square: a Normal posterior with variance $v=(1/\sigma_{\beta}^2+1/\,\mathrm{se}^2)^{-1}$
   - and mean $v(\bar{\beta}/\sigma_{\beta}^2+\hat{\beta}/\,\mathrm{se}^2)=w\hat{\beta}+(1-w)\bar{\beta}$, with $w=\sigma_{\beta}^2/(\sigma_{\beta}^2+\mathrm{se}^2)$

2. Law of total variance: $\mathrm{Var}(\hat{\beta})=\mathrm{Var}(\mathbb{E}[\hat{\beta}\mid\beta])+\mathbb{E}[\mathrm{Var}(\hat{\beta}\mid\beta)]=\sigma_{\beta}^2+\mathbb{E}[\mathrm{se}^2]$
   - Empirical Bayes: $\hat{\sigma}_{\beta}^2=\max\{\mathrm{Var}_{cs}(\hat{\beta})-\overline{\mathrm{se}^2},\,0\}$
   - Using $\mathrm{Var}_{cs}(\hat{\beta})$ itself counts sampling noise as true dispersion and shrinks too little

A3 [Solved]: Solution





3. $\hat{\sigma}^2_\beta = 0.1159 - 0.0077 = 0.1082$; $se^2 = 0.0193$; $w = 0.849$; $\beta^V = 0.849 \times 1.415 + 0.151 \times 0.949 = 1.344$; posterior variance 0.0164 (s.d. 0.128); Blume 1.278
4. Errors against 1.172: raw 0.243, Vasicek 0.172, Blume 0.106

- XLK's beta is precise, so Vasicek shrinks it little; Blume moves every beta one third of the way to 1
- The later beta is itself an estimate; A5 (extension) asks why regression to the mean exceeds sampling noise

MFM_ch3_seminar

B3 [Solved]: PCA and the Number of Factors

- **Questions**
  - how many common factors drive the daily returns of the US (United States) sectors?
  - what are these factors?
- **Data**: 11 sector ETFs, daily log returns on common days, 20 June 2018 – 18 September 2026 ($T = 2073$); SPY held out for comparison
- **Tasks**
  1. Standardise each column: Z is $T \times 11$; compute $\mathrm{C} = \mathrm{Z}^\top \mathrm{Z}/T$, its eigenvalues $\ell_1 \geq \cdots \geq \ell_{11}$ and unit-length eigenvectors, signed so that their loadings sum to a positive number.
  2. Report the shares $\ell_k/11$ of PC1–PC3 and the loadings of PC1 and PC2; build the scores $\mathrm{Z}\gamma_1$ (γ_1: the PC1 eigenvector) and correlate them with SPY.
  3. Estimate the number of factors, with $k_{\max} = 5$, by Kaiser, Bai–Ng (IC_{p2}) and Ahn–Horenstein (eigenvalue ratio), as defined on slide "How Many Factors?"
- **Report**: shares, loadings, the correlation with SPY, the three estimates of k and a chart
- **Interpretation**: PC2 opposes utilities and staples to technology
  - would PC2 earn a premium?
  - what would you need to test a PC2 premium?
- Notebook: section B3

Back Exercise Map

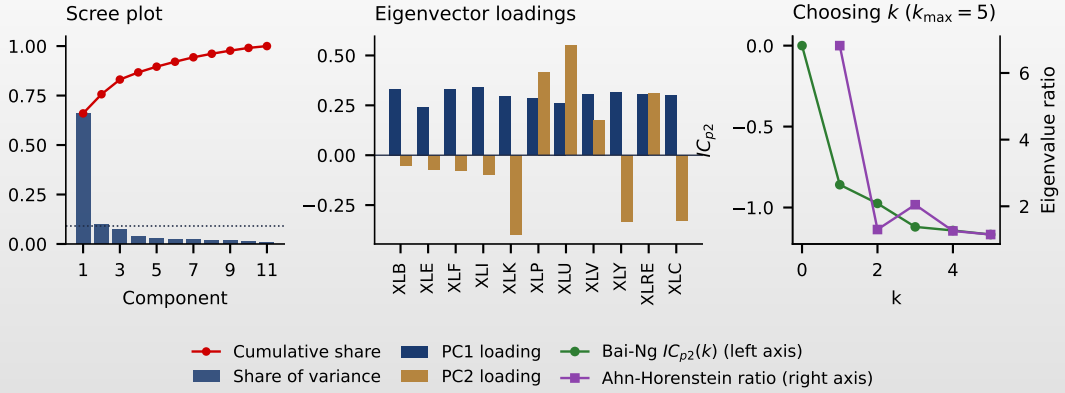
B3 [Solved]: Components

1. Eigenvalues $7.26, 1.06, 0.82, 0.40, 0.32, \ldots$ (they sum to 11)
2. Shares: PC1 66.0%, PC2 9.7%, the third component (PC3) 7.4%; the first three explain 83.1%
   - PC1 loadings all positive (0.24–0.34): an almost equal-weighted portfolio; corr(PC1 scores, SPY) $= 0.954$: PC1 is strongly market-related, but not identical to SPY
   - PC2: XLU $+0.55$, XLP $+0.42$, XLRE $+0.31$ against XLK -0.40, XLY -0.33: defensive versus growth sectors
   - The eigenvectors are portfolio weights on standardised returns, not betas; a return beta on PC1 needs a regression of each sector on the scores
3. Number of factors: Kaiser 2, Bai–Ng 5, Ahn–Horenstein 1 (computation: slide "B3: How Many Factors?")
   - the estimators disagree: with $N = 11$ the large-N asymptotics do not apply; report the range, not a single k

⊡ Interpretation: PCA explains covariance, not average returns, so PC2 need not be a priced factor

- test: Fama–MacBeth with PC2 betas on many assets, or a long–short portfolio sorted on PC2 loadings, its mean tested with Newey–West errors

B3 [Solved]: Chart



MFM_ch3_seminar

B3 [Solved]: How Many Factors?

- Kaiser: $\hat{k} = 2$
  - two eigenvalues exceed 1: 7.26 and 1.06
- Bai–Ng: $\hat{k} = 5 = k_{\max}$
  - the penalty is $\frac{N+T}{NT} \ln 11 = 0.219$ per factor
  - $IC_{p2}(k)$ for $k = 0, \ldots, 5$: $0, -0.86, -0.97, -1.12, -1.14, -1.17$; it falls all the way to $k_{\max}$
- Ahn–Horenstein: $\hat{k} = 1$
  - the ratios ℓ_k / ℓ_{k+1} for $k = 1, \ldots, 5$: $6.82, 1.30, 2.05, 1.26, 1.15$; the maximum is at $k = 1$
- To report: one dominant, market-related factor; Kaiser keeps a second component, but does not establish whether it is significant; Bai–Ng selects the imposed limit $k_{\max}$: inconclusive without sensitivity to $k_{\max}$

B4 [Proposed]: CAPM Betas of BVB Blue Chips

- **Questions**
  - how sensitive are the BVB blue chips to the BET index?
  - how do stale prices bias daily betas?
- **Data**: 11 blue chips (adjusted closing prices) and the BET index, daily log returns from 6 January 2015 (or the listing date) to 18 September 2026
  - each stock is aligned separately with BET, on common days
  - the daily risk-free rate is treated as a constant, so it does not change beta

Checking example: FP around 7 September 2023

Date	Close	r from close	adjusted r	BET_t
5 Sep	1.918	−0.52%	−0.52%	−1.05%
6 Sep	1.984	3.38%	3.38%	1.96%
7 Sep	0.381	−165.0%	5.85%	1.36%
8 Sep	0.386	1.30%	1.30%	1.02%
11 Sep	0.391	1.16%	1.16%	1.34%

- 7 September 2023 is the ex-date of a special dividend of RON 1.7225 per share
- The close falls by 81%, but the shareholder receives the dividend: the total return of the day is +6.0%
- An unadjusted capital event creates a fake return: use adjusted prices and check every day with $|r_t| > 0.5$

Back Exercise Map

B4 [Proposed]: BVB Betas — Tasks and Interpretation

- **Tasks**
  1. Form the lead and the lag of BET on the full calendar of each pair, then drop days with $|r_t| > 0.5$ for the stock or BET; report the dropped days.
  2. Estimate each beta by OLS on BET, with Newey–West errors ($L = \lfloor 4(T/100)^{2/9} \rfloor$ lags) and R^2.
  3. Compute empirical-Bayes Vasicek betas (A3), the 11 stocks as the cross-section.
  4. Estimate the Dimson betas, with BET at $t-1$, t and $t+1$ (slide "Stale Prices and the Dimson Beta").
- **Report**: a table of the three beta estimates (OLS, Vasicek, Dimson), with 95% Newey–West intervals for OLS in the chart
- **Interpretation**
  - which stocks does the Dimson correction move most?
  - why those stocks?
- Model: B1, A3; Notebook: section B4

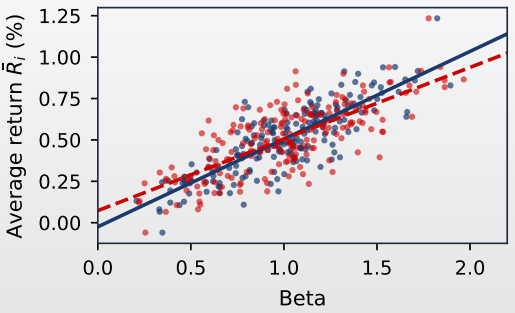
Two-Pass (Fama–MacBeth) Regressions

- **Pass 1** (time series): for each asset, estimate the betas $\hat{\boldsymbol{\beta}}_i$ ($K \times 1$) by OLS; $\hat{\mathrm{B}}$: $N \times K$
- **Pass 2** (Fama & MacBeth (1973); cross-section, every month t): regress the N returns on the betas

$$R^e_{i,t} = \lambda_{\mathbf{0},t} + \hat{\boldsymbol{\beta}}_i^\top \boldsymbol{\lambda}_t + e_{i,t}, \qquad \hat{\boldsymbol{\lambda}} = \frac{1}{T}\sum_{t=\mathbf{1}}^{T} \hat{\boldsymbol{\lambda}}_t$$

  - $\boldsymbol{\lambda}$: the risk premia (prices of risk) of the factors
  - $\lambda_{\mathbf{0}}$: the intercept; zero if investors can borrow at the risk-free rate, and in the Black model the excess of the zero-beta rate over the risk-free rate
  - for a traded factor the model also implies $\lambda = \mathbb{E}[f]$
- Standard errors
  - Fama–MacBeth: $\mathrm{se}(\hat{\lambda}) = \mathrm{sd}(\hat{\lambda}_t)/\sqrt{T}$, the standard deviation of the monthly slopes divided by $\sqrt{T}$
  - if the series $\hat{\lambda}_t$ is autocorrelated, use Newey–West on it
  - the Shanken correction (Shanken (1992)) accounts for the betas being estimated: $\mathrm{Var}_{\mathbf{sh}} = (1+c)(\mathrm{Var}_{FM} - \boldsymbol{\Sigma}_f/T) + \boldsymbol{\Sigma}_f/T$, with $c = \hat{\boldsymbol{\lambda}}^\top \boldsymbol{\Sigma}_f^{-\mathbf{1}} \hat{\boldsymbol{\lambda}}$ and $\boldsymbol{\Sigma}_f$ the factor covariance
  - Kan–Robotti–Shanken (Kan, Robotti & Shanken (2013)): standard errors computed by GMM (generalized method of moments), valid also when the model is misspecified
  - misspecification-robust is not HAC: with autocorrelated moments, add Newey–West lags to the moment covariance (`two_pass_gmm(..., lags=6)`)

Errors in Variables and Useless Factors



- true β_i: slope 0.53
- estimated $\hat{\beta}_i$: slope 0.43

- **Errors in variables**: pass 2 uses $\hat{\beta}_i = \beta_i + \eta_i$, with estimation noise η_i
  - the slope on a noisy regressor is attenuated toward 0 by the factor $\text{Var}_{cs}(\beta)/[\text{Var}_{cs}(\beta) + \text{Var}(\eta)]$
  - chart: 200 simulated assets, $T = 60$; true slope 0.53, on estimated betas 0.43
  - remedies: portfolios (smaller $\text{Var}(\eta)$), long samples, the Shanken correction for the standard error
- **Useless factor** (Kan & Zhang (1999)): a factor unrelated to returns has true betas equal to 0
  - its estimated betas are pure noise, yet pass 2 can report a large, "significant" premium (A7)
  - check first that the betas differ from 0 jointly (pass 1)

Cross-Sectional Fit, OLS and GLS

- Fit of mean returns $\boldsymbol{\mu}$ ($N \times 1$) by $\mathsf{X} = [1, \hat{\mathsf{B}}]$

$$R^{2}_{OLS} = 1 - \frac{\mathsf{e}^{\top}\mathsf{e}}{\sum_i (\mu_i - \bar{\mu})^{2}}, \qquad \mathsf{e} = \boldsymbol{\mu} - \mathsf{X}\hat{\boldsymbol{\lambda}}^{+}$$

  - $\hat{\boldsymbol{\lambda}}^{+} = (\mathsf{X}^{\top}\mathsf{X})^{-1}\mathsf{X}^{\top}\boldsymbol{\mu}$: the intercept $\hat{\lambda}_{0}$ and the premia $\hat{\boldsymbol{\lambda}}$; $\bar{\mu}$: the average of the μ_i
  - with few, strongly factor-driven test assets (25 size $\times$ B/M) almost any model with the right factors reaches a high R^{2}_{OLS} (Lewellen, Nagel & Shanken (2010))

- **GLS**: weight the pricing errors by the inverse covariance $\hat{\mathsf{V}}^{-1}$ of the returns

$$R^{2}_{GLS} = 1 - \frac{\mathsf{e}_{G}^{\top}\hat{\mathsf{V}}^{-1}\mathsf{e}_{G}}{(\boldsymbol{\mu} - \bar{\mu}_{G}1)^{\top}\hat{\mathsf{V}}^{-1}(\boldsymbol{\mu} - \bar{\mu}_{G}1)}$$

  - it measures the fit of the portfolios that matter for the maximum Sharpe ratio, not only of the 25 averages
  - $\mathsf{e}_{G} = \boldsymbol{\mu} - \mathsf{X}\hat{\boldsymbol{\lambda}}^{+}_{GLS}$, $\hat{\boldsymbol{\lambda}}^{+}_{GLS} = (\mathsf{X}^{\top}\hat{\mathsf{V}}^{-1}\mathsf{X})^{-1}\mathsf{X}^{\top}\hat{\mathsf{V}}^{-1}\boldsymbol{\mu}$; $\bar{\mu}_{G} = 1^{\top}\hat{\mathsf{V}}^{-1}\boldsymbol{\mu}/1^{\top}\hat{\mathsf{V}}^{-1}1$: the GLS mean
  - remedy of Lewellen, Nagel & Shanken (2010): add test assets with other characteristics (industries) and report both R^{2}

Comparing Factor Models by Sharpe Ratios

- Maximum squared Sharpe ratio attainable with the K factors of a model (Barillas & Shanken (2018))

$$SR^2 = \boldsymbol{\mu}_f^\top \boldsymbol{\Sigma}_f^{-1} \boldsymbol{\mu}_f$$

  - $\boldsymbol{\mu}_f$, $\boldsymbol{\Sigma}_f$: means and covariance of the factors; the better model prices more assets and has the higher SR^2
- **Nested** models (one contains the other): a **spanning regression**

$$f_{\mathbf{new},t} = a + \mathrm{b}^\top \mathrm{f}_t + u_t$$

  - $a \neq 0$: the new factor raises the attainable Sharpe ratio; $a = 0$: it is spanned (redundant); test a with Newey–West
- **Non-nested** models (Barillas, Kan, Robotti & Shanken (2020)): bootstrap the difference $SR_A^2 - SR_B^2$
  - use the same resampled months for both models, so that the difference keeps their correlation

Rolling and Historical Betas

- **Rolling beta**: estimate $\beta_{i,t}$ on the last H months up to $t-1$ only (C2: $H = 36$, at least 18)
  - a beta-sorted portfolio formed at $t-1$ is then implementable, without look-ahead
- **Low-beta anomaly** (Frazzini & Pedersen (2014)): low-beta stocks earn more than the SML predicts, high-beta stocks less
  - BAB (betting against beta): leverage the low-beta side and scale the high-beta side so that the portfolio beta is 0
  - an unscaled low-minus-high spread has a negative beta: part of its return is market exposure
- **Regression to the mean**: betas estimated in one period move toward 1 in the next
  - regress later betas on earlier ones across assets: a slope below 1 measures the pull toward the mean (A5)

Tools of the Replication Project (Pack 8)

- **Factor themes**: hierarchical clustering of the factors (Ward's method, Ward (1963))
  - distance $d_{jk} = 1 - \text{corr}(\hat{\varepsilon}_j, \hat{\varepsilon}_k)$ between CAPM residuals; Ward merges the two clusters whose union raises the within-cluster variance least; Ward needs Euclidean distances, for standardised series $\sqrt{2(1-\rho_{jk})}$ (`scipy` with $1-\rho$ directly is not equivalent)
  - the dendrogram shows the merges; cutting it at a height gives the themes
- **Hierarchical Bayes** for many alphas (Jensen, Kelly & Pedersen (2023))

$$\hat{\boldsymbol{\alpha}} \sim N\big(0,\ \mathrm{M}\mathrm{M}^{\top}\tau_c^{\mathbf{2}} + \mathrm{I}\tau_w^{\mathbf{2}} + \hat{\boldsymbol{\Sigma}}/T\big)$$

  - M: the 0/1 matrix showing which theme each factor belongs to
  - τ_c: the standard deviation of alphas across themes; τ_w: the standard deviation of factor alphas within a theme
  - $\hat{\boldsymbol{\Sigma}}/T$: the sampling noise of the estimated alphas
  - the posterior mean of each alpha borrows strength from its theme, as Vasicek borrows from the cross-section (slide "Bayesian Shrinkage of Beta")
- **Post-publication decay**: regress the out-of-sample alphas on the in-sample ones
  - a slope below 1: part of the published alpha was noise or was arbitraged away

Extension Packs: What Each One Needs

Pack	Exercises	Prerequisites
1. Pricing restrictions	A2 Extended [Solved], A4, A5	A2, A3, B2
2. Two-pass inference	A6 [Solved], A7, B5, B6, B7	B1, B2, lecture Section 6
3. Model comparison	B8	B1, lecture Section 7
4. Multiple testing on real alphas	B9	A1, B1
5. Latent factors	A8, B10	B3
6. Daily factor ETFs	B11	B1
7. BVB factor premia	C2	B4, C1
8. Replication project	C3	A3, A1, lecture case study

- Each pack has its own notebook section and starter code; do them in order within a pack
- Useful for the team project: C2 and C3 are possible project starts

A2 Extended [Solved]: Other Efficient Portfolios

- **Question**: does the beta pricing equation need the market portfolio, or only an efficient benchmark?
- **Inputs**: as in A2; two benchmarks
  - a CML portfolio, $a\mathrm{w}_m$ in risky assets ($a > 0$) and $1 - a$ in the risk-free asset
  - a frontier portfolio of risky assets only (gross returns R, means m): the first-order condition of $\min \mathrm{w}^\top \boldsymbol{\Sigma} \mathrm{w}$ subject to $\mathrm{w}^\top \mathrm{m} = m^*$, $\mathrm{w}^\top 1 = 1$ gives $\boldsymbol{\Sigma}\mathrm{w} = \kappa_1 \mathrm{m} + \kappa_2 1$, with κ_1, κ_2 the Lagrange multipliers of the two constraints
- **Tasks**
  1. Show that the CML portfolio has $\beta_{i,p} = \beta_i / a$ and $\mu_p = a\mu_m$, so $\mu_i = \beta_{i,p}\mu_p$ again.
  2. Assume $\kappa_1 \neq 0$ (the portfolio is not the global minimum-variance portfolio). From the first-order condition, show that $m_i - \gamma = \beta_{i,p}(m_p - \gamma)$ with $\gamma = -\kappa_2/\kappa_1$, the zero-beta rate.
  3. State what this implies for a test of the CAPM with a market proxy.
- **Report**: the two proofs and one sentence on what a CAPM test checks
- Model: A2

Back Exercise Map

A2 Extended [Solved]: Solution

1. The excess return of the CML portfolio is aR_m^e: $\mathrm{Cov}(R_i^e, aR_m^e) = a\,\mathrm{Cov}(R_i^e, R_m^e)$ and $\mathrm{Var}(aR_m^e) = a^2\sigma_m^2$, so $\beta_{i,p} = \beta_i/a$; with $\mu_p = a\mu_m$, $\beta_{i,p}\mu_p = \beta_i\mu_m = \mu_i$
2. $\mathrm{Cov}(\mathrm{R}, R_p) = \boldsymbol{\Sigma}\mathrm{w} = \kappa_1\mathrm{m} + \kappa_2 1$ and $\sigma_p^2 = \mathrm{w}^\top\boldsymbol{\Sigma}\mathrm{w} = \kappa_1 m_p + \kappa_2$
   - Hence $\kappa_1(\mathrm{m} - \gamma 1) = \mathrm{Cov}(\mathrm{R}, R_p)$ and $\kappa_1(m_p - \gamma) = \sigma_p^2$; divide: $\mathrm{m} - \gamma 1 = \boldsymbol{\beta}_p(m_p - \gamma)$
   - γ is the mean of every portfolio with zero beta on p: the zero-beta CAPM of Black (1972)
3. Any mean–variance efficient benchmark prices all assets exactly by beta; a rejection means that the proxy is not efficient for the test assets, which is a statement about the proxy as much as about the model (Roll's critique, Roll (1977))

A4 [Proposed]: GRS Equals a Sharpe-Ratio Gain

- **Question**: why does GRS measure how much the test assets improve the market's Sharpe ratio?
- **Notation**: $Y_t = (f_t, \mathrm{R}_t^{e\top})^\top$, $(N+1)\times 1$; ML moments $\boldsymbol{\mu}_Y = (\mu_f, \boldsymbol{\mu}_R^\top)^\top$ and $\mathrm{V} = \begin{pmatrix} \sigma_f^2 & \boldsymbol{\sigma}_{Rf}^\top \\ \boldsymbol{\sigma}_{Rf} & \boldsymbol{\Sigma}_R \end{pmatrix}$, all divided by T; $\boldsymbol{\beta} = \boldsymbol{\sigma}_{Rf}/\sigma_f^2$, $\boldsymbol{\alpha} = \boldsymbol{\mu}_R - \boldsymbol{\beta}\mu_f$, $\boldsymbol{\Sigma} = \boldsymbol{\Sigma}_R - \boldsymbol{\beta}\boldsymbol{\beta}^\top \sigma_f^2$
- **Supplied identity** (block inverse with the Schur complement $\boldsymbol{\Sigma}$ of σ_f^2)

$$\mathrm{V}^{-1} = \begin{pmatrix} 1/\sigma_f^2 + \boldsymbol{\beta}^\top \boldsymbol{\Sigma}^{-1}\boldsymbol{\beta} & -\boldsymbol{\beta}^\top \boldsymbol{\Sigma}^{-1} \\ -\boldsymbol{\Sigma}^{-1}\boldsymbol{\beta} & \boldsymbol{\Sigma}^{-1} \end{pmatrix}$$

- **Tasks**
  1. Expand $SR^2(f, \mathrm{R}) = \boldsymbol{\mu}_Y^\top \mathrm{V}^{-1} \boldsymbol{\mu}_Y$ into four terms and show that it equals $\mu_f^2/\sigma_f^2 + \boldsymbol{\alpha}^\top \boldsymbol{\Sigma}^{-1}\boldsymbol{\alpha}$.
  2. Write the GRS statistic W as a function of $SR^2(f)$ and $SR^2(f, \mathrm{R})$.
  3. The 25 size $\times$ B/M portfolios: $T = 757$, $N = 25$, monthly $SR(f) = 0.134$, $SR(f, \mathrm{R}) = 0.405$; compute W and decide at 5% (critical value of $F(25, 731)$: 1.52).
  4. Compare the observed gain with the gain expected under H_0, $N(1 + SR^2(f))/(T - N - 3)$.
- **Report**: the algebra, W, the decision and one sentence
- Model: A2, B2; Notebook: section A4

Back Exercise Map

A5 [Proposed]: Blume's Slope as a Shrinkage Weight

- **Question**: is the regression of later betas on earlier betas a Vasicek estimator?
- **Simplifying assumptions**: true betas constant, $\hat{\beta}_{i,k} = \beta_i + \eta_{i,k}$ in periods $k = 1, 2$; the noises are independent of β_i and of each other, with a **common** variance s^2; β_i has mean $\bar{\beta}$ and variance σ^2_β
- **Tasks**
  1. Derive, as $N \to \infty$, the limits of the OLS slope b and intercept a of $\hat{\beta}_2$ on $\hat{\beta}_1$.
  2. Show that $a + b\hat{\beta}_1$ is then the Vasicek beta with $w = b$.
  3. With A3's $\sigma^2_\beta = 0.10822$, $s^2 = 0.00773$, $\bar{\beta} = 0.9495$, compute the implied b and a; compare them with the fitted 0.7242 and 0.2304 and plot both lines.
  4. The fitted slope has classical OLS standard error 0.1137 ($N = 9$): compute $(0.7242 - b)/0.1137$ as an illustration and list the explanations that remain possible.
- **Report**: the two limits, the implied b and a, the chart of the two lines and the list of explanations
- Model: A3; Notebook: section A5

	$\hat{\beta}_1$ 1999–2012	se$_1$	$\hat{\beta}_2$ 2013–2026
XLB	1.174	0.083	1.026
XLE	0.904	0.089	1.051
XLF	1.217	0.120	1.029
XLI	1.122	0.057	1.078
XLK	1.415	0.139	1.172
XLP	0.451	0.064	0.578
XLU	0.463	0.086	0.455
XLV	0.716	0.047	0.705
XLY	1.083	0.063	1.169

Back Exercise Map

A6 [Solved]: Errors in Variables in Pass 2

- **Question**: how much does pass 2 attenuate a premium because the betas are estimated?
- **Setting**: one factor, $R_{it} = a_i + \beta_i f_t + \varepsilon_{it}$; ε_{it} i.i.d. over t and i, variance σ_ε^2, independent of f and of the β_i; the β_i i.i.d. across assets with variance $\mathrm{Var}_{cs}(\beta)$
- **Three premia**: λ, the true price of risk; $\mu_f = \mathbb{E}[f_t]$; $\lambda_T = \lambda + \bar{f} - \mu_f$, the realised premium in the sample. Pricing restriction: $a_i = \lambda_0 + \beta_i(\lambda - \mu_f)$
- Work conditionally on the factor path $f_1, \ldots, f_T$; pass 2: the OLS regression of $\bar{R}_i$ on $(1, \hat{\beta}_i)$, $\hat{\beta}_i = \beta_i + \eta_i$, $\eta_i = \sum_t (f_t - \bar{f})\varepsilon_{it} / \sum_t (f_t - \bar{f})^2$
- **Tasks**
  1. Show that $\mathrm{Var}(\eta_i \mid f) = \sigma_\varepsilon^2 / (T \hat{\sigma}_f^2)$ and that $\mathrm{Cov}(\eta_i, \bar{\varepsilon}_i \mid f) = 0$.
  2. Show that, for $N \to \infty$ and T fixed, the pass-2 slope converges to $\lambda_T \ \mathrm{Var}_{cs}(\beta) / [\mathrm{Var}_{cs}(\beta) + \sigma_\varepsilon^2 / (T \hat{\sigma}_f^2)]$.
  3. Compute the attenuation factor for the two panels below, with $\mathrm{Var}_{cs}(\beta) = \mathrm{Var}_{cs}(\hat{\beta}) - \bar{\sigma}_\varepsilon^2 / (T \hat{\sigma}_f^2)$, and explain what the Shanken correction adjusts.
- **Report**: the two proofs, the two attenuation factors and one sentence on the Shanken correction; Notebook: section A6

Panel (CAPM)	T	$\mathrm{Var}_{cs}(\hat{\beta})$	$\bar{\sigma}_\varepsilon^2$ (%² per month)	$\hat{\sigma}_f^2$ (%² per month)
25 size × B/M, 1963–2026	757	0.01905	9.33	19.87
9 sector ETFs, 1999–2026	331	0.09247	13.45	20.36

Back Exercise Map

A6 [Solved]: Derivation

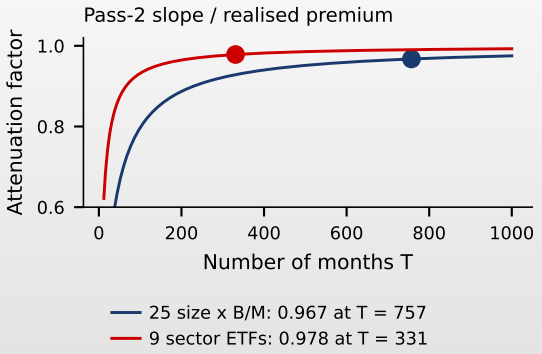
1. Given f, η_i is a linear combination of the ε_{it} with weights $(f_t - \bar{f})/\sum_s (f_s - \bar{f})^2$
   - $\text{Var}(\eta_i \mid f) = \sigma_\varepsilon^2 \sum_t (f_t - \bar{f})^2 / [\sum_s (f_s - \bar{f})^2]^2 = \sigma_\varepsilon^2 / (T \hat{\sigma}_f^2)$
   - $\text{Cov}(\bar{\varepsilon}_i, \eta_i \mid f) = \sigma_\varepsilon^2 \frac{1}{T} \sum_t (f_t - \bar{f}) / \sum_s (f_s - \bar{f})^2 = 0$, because $\sum_t (f_t - \bar{f}) = 0$
2. With the restriction, $\bar{R}_i = \lambda_0 + \lambda_T \beta_i + \bar{\varepsilon}_i$: the realised factor mean $\bar{f}$ is common to all assets and does not average out as $N \to \infty$
   - Given f, $(\beta_i, \eta_i, \bar{\varepsilon}_i)$ are i.i.d. across i: $\widehat{\text{Cov}}_{cs}(\bar{R}, \hat{\beta}) \to \lambda_T \, \text{Var}_{cs}(\beta)$ and $\widehat{\text{Var}}_{cs}(\hat{\beta}) \to \text{Var}_{cs}(\beta) + \sigma_\varepsilon^2 / (T \hat{\sigma}_f^2)$
   - The ratio is the stated limit: attenuation of λ_T towards zero; $\lambda_T - \lambda = \bar{f} - \mu_f$ vanishes only as $T \to \infty$

A6 [Solved]: Numbers

3. 25 portfolios: noise $9.33/(757 \times 19.87) = 0.00062$; $\mathrm{Var}_{cs}(\beta) = 0.01905 - 0.00062 = 0.01843$; factor $0.01843/0.01905 = 0.967$

   9 sectors: noise $13.45/(331 \times 20.36) = 0.00200$; $\mathrm{Var}_{cs}(\beta) = 0.09048$; factor 0.978

- The bias is only 2–3% here: diversified portfolios have small residual variance and the sample is long; single stocks have a much larger σ_ε^2, so the attenuation grows
- Shanken (1992) corrects the standard error of $\hat{\lambda}$ for the noise in $\hat{\beta}$ (a variance correction); it does not remove the finite-T bias of the point estimate



MFM_ch3_seminar

A7 [Proposed]: A Useless Factor in Pass 2

- **Question**: why can a factor unrelated to returns look priced in pass 2?
- **Assumptions**: N fixed, $T \to \infty$; R_t i.i.d. over time with means μ_i, nonsingular covariance $\mathbf{\Sigma}$, finite fourth moments; a factor g_t, i.i.d., mean 0, variance σ_g^2, independent of all returns, so $\beta_{i,g} = 0$; in the proof g is the only factor in pass 2
- **Tasks**
  1. Show that $\sqrt{T}\,\hat{\beta}_{i,g} \to N(0, \sigma_i^2/\sigma_g^2)$: $\hat{\beta}_{i,g}$ is of order $T^{-1/2}$ (numerator and denominator separately).
  2. Write $\hat{\lambda}_g = \widehat{\mathrm{Cov}}_{cs}(\bar{R}_i, \hat{\beta}_{i,g}) / \widehat{\mathrm{Var}}_{cs}(\hat{\beta}_{i,g})$ with $z = \sqrt{T}\,\hat{\boldsymbol{\beta}}_g$; show $\hat{\lambda}_g$ is of order $\sqrt{T}$ when the μ_i differ.
  3. Case of equal means: using the joint limit of $(\sqrt{T}(\bar{R} - \boldsymbol{\mu}), z)$, show that $\hat{\lambda}_g$ converges to a non-degenerate random variable, not 0.
  4. The simulation of Chapter 3 adds g to the CAPM on the 25 portfolios (2000 replications): $\lambda_g = 0$ is rejected in 59.6% of the replications with Fama–MacBeth errors and in 26.6% with Shanken errors, while the first-pass Wald test of $\beta_{i,g} = 0$ rejects in 6.7%; explain these rates and state which first-pass test protects against them.
- **Report**: the three asymptotic results and the explanation of the rejection rates
- Model: A6; Notebook: section A7

B5 [Proposed]: Fama–MacBeth on ETFs

- **Question**: are there market, size and value premia in 16 US ETFs over 2013–2026?
- **Data**: monthly excess returns, August 2013 – July 2026 ($T = 156$), a 156×16 panel
  - 9 sector ETFs
  - 7 factor ETFs: MTUM (momentum), VLUE (value), QUAL (quality), USMV (minimum volatility), IWM (small firms, Russell 2000), IWD (large value firms, Russell 1000 Value), IWF (large growth firms, Russell 1000 Growth)
  - the FF3 factors: Mkt-RF, SMB and HML from the three-factor data set of the Kenneth French Data Library (156×3)
- **Tasks**
  1. Pass 1: estimate the full-sample regression of each ETF on the three factors; report $\hat{B}$ (16×3).
  2. Pass 2: for each month, estimate the OLS regression of the 16 returns on $X = [1, \hat{B}]$; show the computation of one month in full (fitted against actual), then average the 156 slopes; annualise by multiplying by 12.
  3. Compute the Fama–MacBeth, Newey–West (on $\hat{\lambda}_t$, 6 lags) and Shanken standard errors (slide "Two-Pass (Fama–MacBeth) Regressions"; for the intercept, border the factor covariance with a zero row and column: $\mathrm{Var}_{Sh}(\hat{\lambda}_0) = (1 + c)\,\mathrm{Var}_{FM}(\hat{\lambda}_0)$); as an extension, compute the Kan–Robotti–Shanken errors with the notebook helper `two_pass_gmm`.
  4. Compare each $\hat{\lambda}$ with the time-series mean of its factor in a chart and test $\lambda = 0$ at 5%.
- **Report**: $\hat{B}$, the month computed in full, the table of premia with their standard errors, the chart and the test decisions
- **Interpretation**: is a non-significant premium here evidence that size and value are not priced?
- Model: B1, B2 and Chapter 3, Section 6; Notebook: section B5

Back Exercise Map

B6 [Proposed]: A Lewellen–Nagel–Shanken Check

- **Question**: does FF3's high cross-sectional R^2 survive broader test assets and a GLS metric (Lewellen, Nagel & Shanken (2010))?
- **Data**: 25 size $\times$ B/M and 30 industry portfolios (Kenneth French Data Library), monthly excess returns, July 1963 – July 2026 ($T = 757$); FF3 factors
- **Definitions**: μ, X, $\hat{V}$, R^2_{OLS} and R^2_{GLS} as on slide "Cross-Sectional Fit, OLS and GLS"
- **Tasks**
  1. Baseline: FF3 on the 25 portfolios; plot actual against fitted means; report both R^2.
  2. Add the 30 industries and repeat; also impose $\lambda = \mathbb{E}[f]$ with no intercept and report R^2_{OLS} with the same denominator.
  3. Uncertainty: resample months i.i.d. (500 draws, seed 42), re-estimate both passes and $\hat{V}$ in each draw; report 95% intervals.
  4. Extension: repeat the analysis for the CAPM and FF5 on both sets of assets.
- **Report**: the table of the two R^2 with their intervals, on both sets of assets, and the chart
- **Interpretation**: does the FF3 R^2 survive once industries are added?
- Model: B2, Chapter 3, Section 6; Notebook: section B6

Back Exercise Map

B7 [Proposed]: GRS without Normality

- **Question**: is the rejection of the CAPM on the 25 portfolios an artefact of assuming residuals with the Normal distribution?
- **Data**: 25 size $\times$ B/M portfolios, CAPM with Mkt-RF, monthly, July 1963 – July 2026 ($T = 757$, $N = 25$)
- **Tasks**
  1. Diagnostics: estimate the 25 regressions; compute the Pearson kurtosis of each residual series (slide "Normality of the Residuals") and its average; draw a Q–Q plot.
  2. Exact test: compute GRS with the ML $\hat{\boldsymbol{\Sigma}}$ and its p-value from $F(25, 731)$.
  3. Null bootstrap: resample whole rows of residuals (keeping their cross-correlation), rebuild $\mathrm{R}_t^* = \hat{\boldsymbol{\beta}} f_t + \boldsymbol{\varepsilon}_t^*$ (so $\alpha = 0$), refit and recompute GRS 2000 times, seed 42; report $p = (1 + \#\{W^* \geq W\})/2001$ and the 95% bootstrap critical value.
  4. Extension: compute the robust Wald test with the moments $(1, f_t)^\top \otimes \boldsymbol{\varepsilon}_t$, parameters ordered $(\boldsymbol{\alpha}, \boldsymbol{\beta})$; $W = \hat{\boldsymbol{\alpha}}^\top \hat{\mathrm{V}}_\alpha^{-1} \hat{\boldsymbol{\alpha}} \to \chi^2_{25}$ with the sandwich $\mathrm{D}^{-1}\hat{\mathrm{S}}\mathrm{D}^{-1}/T$, $\mathrm{D} = (\mathrm{Z}^\top \mathrm{Z}/T) \otimes \mathrm{I}_{25}$, $\mathrm{Z} = [1, \mathrm{f}]$, and $\hat{\mathrm{S}}$ White or Newey–West (6 lags).
  5. Explain why GRS cannot be computed for the 100 size $\times$ B/M portfolios over August 2018 – July 2026 ($T = 96$).
- **Report**: the average kurtosis, the Q–Q plot, W with the p-values of the three methods and the bootstrap critical value
- **Interpretation**: does the departure of the residuals from the Normal distribution change the decision of the GRS test?
- Model: B2; Notebook: section B7

Back Exercise Map

B8 [Proposed]: Comparing Models by Sharpe Ratios

- **Questions**
  - which factor model spans the highest Sharpe ratio?
  - is the difference significant (Barillas & Shanken (2018); Barillas, Kan, Robotti & Shanken (2020))?
- **Data**: monthly Mkt-RF, SMB, HML, RMW, CMA, MOM (Kenneth French Data Library), July 1963 – July 2026 ($T = 757$) and January 2000 – July 2026 ($T = 319$)
- **Membership**: CAPM (Mkt-RF); FF3 and Carhart use the three-factor SMB, FF5 and FF5 + MOM the five-factor one; covariances divide by $T - 1$
- **Tasks**
  1. Compute the maximum squared Sharpe ratio SR^2 of each model (slide "Comparing Factor Models by Sharpe Ratios"), annualise ($\times 12$) and report $\sqrt{SR^2}$.
  2. Nested models, FF5 + MOM against FF5: regress MOM on the five FF5 factors and test the intercept (spanning regression, Newey–West).
  3. Non-nested models, FF5 against Carhart: apply a stationary bootstrap (Politis & Romano (1994)) with a mean block of 6 months and 2000 draws, seed 42, with the same indices for both models; report the 95% interval of the SR^2 difference.
  4. Extension: compute the spanning t of each of the six factors on the other five, in both periods.
- **Report**: the Sharpe ratios in both periods, the spanning t and the bootstrap interval
- **Interpretation**
  - why can the percentile interval of a nested difference not test equality?
  - which factors add no significant alpha after 2000?
- Model: B1; Notebook: section B8

Back Exercise Map

B9 [Proposed]: Multiple Testing of 25 Alphas

- **Question**: how many of the 25 size $\times$ B/M portfolios have a genuine FF3 alpha?
- **Data**: 25 value-weighted size $\times$ B/M portfolios (the "25 Portfolios Formed on Size and Book-to-Market" data set of the Kenneth French Data Library) and the FF3 factors (Fama & French (1993)), monthly excess returns, July 1963 – July 2026 ($T = 757$)
- **Tasks**
  1. For one portfolio (small-growth), show the full computation: FF3 regression, Newey–West t of α (6 lags), two-sided $p = 2(1 - \Phi(|t|))$, Φ: the distribution function of the standard Normal distribution.
  2. Repeat for all 25; draw a 5×5 heatmap of the annual alphas with their t-statistics.
  3. Sort the p-values and apply the A1 worksheet: naive, Bonferroni, Holm, Benjamini–Hochberg; extension: Benjamini–Yekutieli.
  4. Count the portfolios with $|t| > 3$ (Harvey, Liu & Zhu (2016)), a heuristic hurdle, not a procedure with a guaranteed error rate.
- **Report**: the small-growth computation, the heatmap, the number of discoveries of each procedure and the number of portfolios with $|t| > 3$
- **Interpretation**: the 25 tests are strongly correlated
  - which methods remain valid?
  - which test uses the correlation directly?
- Model: A1; Notebook: section B9

Back Exercise Map

A8 [Proposed]: Equicorrelated Assets

- **Question**: how does the share of the first component depend on the average correlation and on N?
- **Stage 1**, $N = 3$, $\rho = 0.5$: $\mathrm{C} = \begin{pmatrix} 1 & 0.5 & 0.5 \\ 0.5 & 1 & 0.5 \\ 0.5 & 0.5 & 1 \end{pmatrix}$
  1. Verify that $\mathbf{v_1} = (1,1,1)/\sqrt{3}$, $\mathbf{v_2} = (1,-1,0)/\sqrt{2}$, $\mathbf{v_3} = (1,1,-2)/\sqrt{6}$ are orthonormal eigenvectors and find the eigenvalues.
  2. Compute the PC1 share $\ell_{\mathbf{1}}/3$.
- **Stage 2**, N assets, $\mathrm{C} = (1-\rho)\mathrm{I} + \rho 1 1^{\top}$
  3. Show that $\ell_{\mathbf{1}} = 1 + (N-1)\rho$ and that all other eigenvalues equal $1-\rho$; find the limit of the PC1 share as $N \to \infty$.
  4. Invert the share formula to get the equicorrelation-equivalent ρ of a PC1 share s; apply it to the BVB blue chips ($N = 10$, $s = 0.425$, B10) and the US sectors ($N = 11$, $s = 0.660$, B3).
  5. Extension: for a general C with average off-diagonal correlation $\bar{\rho}$, show that $1 + (N-1)\bar{\rho} \le \ell_{\mathbf{1}}$.
- **Report**: the eigenvalues, the formula of the PC1 share with its limit and the two values of ρ
- Model: B3; Notebook: section A8

Back Exercise Map

B10 [Proposed]: PCA of BVB Blue Chips

- **Question**: is the first factor of the Romanian blue chips as strong as in the US sectors?
- **Data**: 10 blue chips (TLV, SNP, BRD, TGN, FP, SNG, SNN, EL, DIGI, TEL; H2O excluded, listed in 2023) and BET, aligned on common days from 7 June 2017 to 18 September 2026; daily log returns; drop days with any $|r_t| > 0.5$ and report how many; BET is held out of the PCA
- **Tasks**
  1. Apply the B3 recipe: standardise ($T \times 10$), eigenvalues, variance shares, PC1 loadings.
  2. Correlate the PC1 scores with the BET returns.
  3. Estimate the number of factors by Kaiser, Bai–Ng and Ahn–Horenstein, with $k_{\max} = 5$.
  4. Compare with the US sectors on the same period (from 20 June 2018): PC1 shares and scree plots side by side.
- **Report**: the number of days dropped, the shares, the PC1 loadings, the correlation with BET, the three estimates of k and the scree plots
- **Interpretation**
  - what does a weaker first factor imply for diversification within the BVB?
  - how much of the gap comes from comparing stocks with sector portfolios?
- Model: B3; Notebook: section B10

Back Exercise Map

B11 [Proposed]: Do Factor ETFs Have Alpha?

- **Question**: do MTUM and QUAL earn an alpha beyond their six-factor exposures?
- **Data**: adjusted prices of MTUM (18 Apr 2013 – 31 Jul 2026) and QUAL (18 Jul 2013 – 31 Jul 2026); daily FF5 (Fama & French (2015)), MOM (Carhart (1997)) and RF, same end date
- **Calendar rule**: keep an ETF return only if the previous ETF price dates from the immediately preceding trading day of the factor calendar
  - Example: MTUM has prices on 18 and 23 April 2013 but none on 19 and 22 April; the return 18 $\rightarrow$ 23 April covers three factor days and is dropped
- **Tasks**
  1. Build daily excess returns with the calendar rule; report n before and after alignment.
  2. Estimate $R_t^e = \alpha + b\,\text{Mkt-RF}_t + s\,\text{SMB}_t + h\,\text{HML}_t + r\,\text{RMW}_t + c\,\text{CMA}_t + m\,\text{MOM}_t + \varepsilon_t$; annualise α by multiplying by 252.
  3. Compute the Newey–West standard error (10 lags) and t of α; apply a moving-block bootstrap of the pairs, with blocks of 20 days and 1000 draws, seed 42, and report the 95% interval.
  4. Plot the two intervals for each ETF.
- **Report**: n before and after alignment, the six loadings, α with its s.e. and t, the two intervals and the chart
- **Interpretation**: why must a 3-day ETF return never be regressed on a 1-day factor return?
- Model: B1; Notebook: section B11

Back Exercise Map

C2 [Proposed]: Do Factor Premia Exist on the BVB? — Setting

- **Question**: do momentum (Jegadeesh & Titman (1993)) and low beta earn premia among BVB blue chips?
- **Data**: month-end adjusted prices of the 10 blue chips of C1 and BET, December 2015 – August 2026 (September 2026 incomplete); monthly simple returns; FP prices are adjusted for the special dividend of 7 September 2023 (RON 1.7225 per share)
- **Rules fixed in advance**
  - signal and beta use information up to the end of $t-1$ (timeline of C1)
  - at least 6 eligible stocks; ties broken by ticker order
  - weights stay those fixed at formation; a missing holding return (suspension, missing quote, delisting) is not dropped from the leg mean: on delisting include the actual terminal value; otherwise mark the month incomplete and report how often this happens

Ranking for August 2026

Stock	Signal	Leg
EL	129.9%	long
TGN	127.3%	long
DIGI	113.2%	long
SNG	88.9%	
BRD	87.8%	
SNN	75.5%	
TEL	66.1%	
FP	51.0%	short
TLV	48.5%	short
SNP	46.6%	short

- Signal: the compounded return of August 2025 – June 2026, known at the end of July 2026

Back Exercise Map

C2 [Proposed]: Do Factor Premia Exist on the BVB? — Tasks

- **Tasks**
  1. Build 12–1 momentum (C1 contract): long the top 3, short the bottom 3, equal weights.
  2. Build the low-beta portfolio: estimate 36-month rolling betas on BET (at least 18 months, ending at $t-1$); long the 3 lowest, short the 3 highest; an unscaled spread, not the beta-neutral BAB of Frazzini & Pedersen (2014).
  3. Report periods, T, annual mean, volatility and Newey–West t (3 lags); compare t with 1.96 and 3; plot the cumulative spreads.
  4. Discuss survivorship (today's blue chips chosen with hindsight), short-selling constraints, costs and liquidity.
  5. Extension: identify public sources of book equity and shares outstanding, needed for a BVB value factor.
- **Deliverables**: a Quantlet folder in your repository (script, data source and short description) and one slide with the conclusion and its main limitation
- **Assessment criteria**
  - signals formed only with the information available at the end of month $t-1$
  - the Newey–West t statistic compared with 1.96 and with 3, with the decision stated
  - the limitations discussed: survivorship bias, short-selling restrictions, costs
  - the C1 AI critique documented in `AI_ERRORS.md`
- Model: B4, C1; Notebook: Part C

C3 [Proposed]: Do US Factors Replicate? A Staged Project (1/2)

- **Context**: the case study of Chapter 3, Jensen, Kelly & Pedersen (2023): 82.4% of US factors keep a significant CAPM alpha under a hierarchical empirical-Bayes model, against 75.6% with Benjamini–Yekutieli and 35% in Hou, Xue & Zhang (2020) on raw returns (Figure 1)
- **Data**: value-weighted portfolios sorted on a single characteristic, from the Kenneth French Data Library, July 1963 – July 2026, and Mkt-RF from the three-factor data set
- Each factor buys one end of the sort and sells the other, for 15 characteristics
  - 30% at each end: size, book-to-market, operating profitability, investment, earnings/price, cash flow/price, dividend/price
  - 20% at each end: accruals, net share issues, market beta, variance, residual variance
  - 10% at each end: prior returns of months 12–2 (the 12–1 momentum, slide "Momentum and Look-Ahead Bias"), 1–0 and 60–13
- **Stage 1, construction and frequentist rates**
  1. Build the 15 factors, long the side that earned more in the original paper; cite that paper for each sign and record its sample end year and publication date in the factor table.
  2. Regress each factor on Mkt-RF, rescale it to a monthly residual volatility of $10\%/\sqrt{12}$ (Section II.A of the paper), report $\hat{\alpha}$, t and the OLS replication rate ($t \geq 1.96$).
  3. Apply Benjamini–Yekutieli at 5% and report the critical $|t|$ and the rate.

Back Exercise Map

C3 [Proposed]: Do US Factors Replicate? A Staged Project (2/2)

- **Stage 2, themes and the hierarchical model**
  4. Build the Ward dendrogram on the distance $1-$ correlation of the CAPM residuals (authors: R `hclust(method="ward.D")` on $1-\rho$; in SciPy, `linkage(method="ward")` on the condensed distances $\sqrt{2(1-\rho)}$, the R equivalent being `ward.D2`; state the distance–algorithm pair); compare it with the paper's themes (`Cluster Labels.csv`, authors' code).
  5. Build M ($15 \times J$), the theme-membership matrix; estimate τ_c, τ_w by maximum likelihood of $\hat{\boldsymbol{\alpha}} \sim N(0, \mathrm{M}\mathrm{M}^\top \tau_c^2 + \mathrm{I}\tau_w^2 + \hat{\boldsymbol{\Sigma}}/T)$, report a boundary estimate $\tau_c = 0$ if it occurs; compute the posterior means and variances (Proposition 4) and the empirical-Bayes rate (posterior $z \geq 1.96$).
- **Stage 3, sensitivity**
  6. Compute the empirical-Bayes rate with a single theme.
  7. Plot the rate as a function of τ_c and compute the Bayesian FDR (Proposition 5).
- **Stage 4, post-sample decay**
  8. Estimate the alphas within the original sample years (from July 1963) and after them, with at least 60 months in each.
  9. Regress the post-sample on the in-sample alphas by feasible GLS with the inverse factor covariance and compare the slope with 0.26 (Figure 8).
- **Report**: a factor table (sample end, publication date, alpha, t, theme, posterior mean, z), the rates next to Figure 1, $\hat{\tau}$ next to Table I, the dendrogram and the slope; keep the paper's numbers and the course numbers apart
- Model: A3, A1 and the case study of Chapter 3

Acronyms used in this material

- **AI**: Artificial Intelligence
- **BAB**: Betting Against Beta
- **BET**: Bucharest Exchange Trading (index)
- **BH**: Benjamini–Hochberg (procedure)
- **BVB**: Bursa de Valori București — *the Bucharest Stock Exchange*
- **CAPM**: Capital Asset Pricing Model
- **CMA**: Conservative Minus Aggressive (investment factor)
- **CML**: Capital Market Line
- **ETF**: Exchange-Traded Fund
- **FDR**: False Discovery Rate
- **FF3**: Fama–French three-factor model
- **FF5**: Fama–French five-factor model
- **FM**: Fama–MacBeth (regressions)
- **FWER**: Family-Wise Error Rate
- **GLS**: Generalised Least Squares
- **GMM**: Generalized Method of Moments
- **GRS**: Gibbons–Ross–Shanken (test)
- **HAC**: Heteroskedasticity and Autocorrelation Consistent (standard errors)
- **HML**: High Minus Low (value factor)
- **KRS**: Kan–Robotti–Shanken (standard errors)
- **ML**: Maximum Likelihood
- **MOM**: Momentum (factor)
- **NW**: Newey–West (standard errors)
- **OLS**: Ordinary Least Squares
- **PC1**: first Principal Component
- **PC2**: second Principal Component
- **PC3**: third Principal Component
- **PCA**: Principal Component Analysis
- **RF**: Risk-Free rate
- **RMW**: Robust Minus Weak (profitability factor)
- **SMB**: Small Minus Big (size factor)
- **SML**: Security Market Line
- **US**: United States

References (1/4)

- Ahn, S.C., Horenstein, A.R. (2013). Eigenvalue ratio test for the number of factors. *Econometrica*, 81(3), 1203–1227.
- Bai, J., Ng, S. (2002). Determining the number of factors in approximate factor models. *Econometrica*, 70(1), 191–221.
- Barillas, F., Kan, R., Robotti, C., Shanken, J. (2020). Model comparison with Sharpe ratios. *Journal of Financial and Quantitative Analysis*, 55(6), 1840–1874.
- Barillas, F., Shanken, J. (2018). Comparing asset pricing models. *Journal of Finance*, 73(2), 715–754.
- Benjamini, Y., Hochberg, Y. (1995). Controlling the false discovery rate: a practical and powerful approach to multiple testing. *Journal of the Royal Statistical Society B*, 57(1), 289–300.
- Benjamini, Y., Yekutieli, D. (2001). The control of the false discovery rate in multiple testing under dependency. *Annals of Statistics*, 29(4), 1165–1188.
- Black, F. (1972). Capital market equilibrium with restricted borrowing. *Journal of Business*, 45(3), 444–455.
- Blume, M.E. (1971). On the assessment of risk. *Journal of Finance*, 26(1), 1–10.
- Campbell, J.Y., Lo, A.W., MacKinlay, A.C. (1997). The Econometrics of Financial Markets. Princeton University Press.
- Carhart, M.M. (1997). On persistence in mutual fund performance. *Journal of Finance*, 52(1), 57–82.

References (2/4)

- Dimson, E. (1979). Risk measurement when shares are subject to infrequent trading. *Journal of Financial Economics*, 7(2), 197–226.
- Fama, E.F., French, K.R. (1993). Common risk factors in the returns on stocks and bonds. *Journal of Financial Economics*, 33(1), 3–56.
- Fama, E.F., French, K.R. (2012). Size, value, and momentum in international stock returns. *Journal of Financial Economics*, 105(3), 457–472.
- Fama, E.F., French, K.R. (2015). A five-factor asset pricing model. *Journal of Financial Economics*, 116(1), 1–22.
- Fama, E.F., MacBeth, J.D. (1973). Risk, return, and equilibrium: empirical tests. *Journal of Political Economy*, 81(3), 607–636.
- Frazzini, A., Pedersen, L.H. (2014). Betting against beta. *Journal of Financial Economics*, 111(1), 1–25.
- Gibbons, M.R., Ross, S.A., Shanken, J. (1989). A test of the efficiency of a given portfolio. *Econometrica*, 57(5), 1121–1152.
- Härdle, W.K., Simar, L. (2019). Applied Multivariate Statistical Analysis, 5th ed. Springer.
- Harvey, C.R., Liu, Y., Zhu, H. (2016). . . . and the cross-section of expected returns. *Review of Financial Studies*, 29(1), 5–68.
- Holm, S. (1979). A simple sequentially rejective multiple test procedure. *Scandinavian Journal of Statistics*, 6(2), 65–70.

References (3/4)

- Hou, K., Xue, C., Zhang, L. (2020). Replicating anomalies. *Review of Financial Studies*, 33(5), 2019–2133.
- Jegadeesh, N. (1990). Evidence of predictable behavior of security returns. *Journal of Finance*, 45(3), 881–898.
- Jegadeesh, N., Titman, S. (1993). Returns to buying winners and selling losers: implications for stock market efficiency. *Journal of Finance*, 48(1), 65–91.
- Jensen, T.I., Kelly, B., Pedersen, L.H. (2023). Is there a replication crisis in finance? *Journal of Finance*, 78(5), 2465–2518.
- Kaiser, H.F. (1960). The application of electronic computers to factor analysis. *Educational and Psychological Measurement*, 20(1), 141–151.
- Kan, R., Robotti, C., Shanken, J. (2013). Pricing model performance and the two-pass cross-sectional regression methodology. *Journal of Finance*, 68(6), 2617–2649.
- Kan, R., Zhang, C. (1999). Two-pass tests of asset pricing models with useless factors. *Journal of Finance*, 54(1), 203–235.
- Künsch, H.R. (1989). The jackknife and the bootstrap for general stationary observations. *Annals of Statistics*, 17(3), 1217–1241.
- Lewellen, J., Nagel, S., Shanken, J. (2010). A skeptical appraisal of asset pricing tests. *Journal of Financial Economics*, 96(2), 175–194.
- Markowitz, H. (1952). Portfolio selection. *Journal of Finance*, 7(1), 77–91.
- Newey, W.K., West, K.D. (1987). A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix. *Econometrica*, 55(3), 703–708.

References (4/4)

- Politis, D.N., Romano, J.P. (1994). The stationary bootstrap. *Journal of the American Statistical Association*, 89(428), 1303–1313.
- Roll, R. (1977). A critique of the asset pricing theory's tests. Part I: On past and potential testability of the theory. *Journal of Financial Economics*, 4(2), 129–176.
- Shanken, J. (1992). On the estimation of beta-pricing models. *Review of Financial Studies*, 5(1), 1–33.
- Sharpe, W.F. (1964). Capital asset prices: a theory of market equilibrium under conditions of risk. *Journal of Finance*, 19(3), 425–442.
- Vasicek, O.A. (1973). A note on using cross-sectional information in Bayesian estimation of security betas. *Journal of Finance*, 28(5), 1233–1239.
- Ward, J.H. (1963). Hierarchical grouping to optimize an objective function. *Journal of the American Statistical Association*, 58(301), 236–244.
- White, H. (1980). A heteroskedasticity-consistent covariance matrix estimator and a direct test for heteroskedasticity. *Econometrica*, 48(4), 817–838.
- Data: Kenneth R. French Data Library, mba.tuck.dartmouth.edu/pages/faculty/ken.french/data_library.html.