








Statistics of Financial Markets

Seminar 1: Data, returns and indicators



Daniel Traian PELE

Bucharest University of Economic Studies
IDA Institute Digital Assets
Blockchain Research Center
AI4EFin Artificial Intelligence for Energy Finance
Romanian Academy, Institute for Economic Forecasting
MSCA Digital Finance

Acronyms used in this material

- **AI**: Artificial Intelligence
- **BET**: Bucharest Exchange Trading (index)
- **BNR**: Banca Națională a României (the National Bank of Romania)
- **BVB**: Bursa de Valori București (the Bucharest Stock Exchange)
- **CAGR**: Compound Annual Growth Rate
- **CDF**: Cumulative Distribution Function
- **CI**: Confidence Interval
- **EODHD**: EOD Historical Data (market-data provider)
- **MDD**: Maximum DrawDown
- **OMV**: Österreichische Mineralölverwaltung — Austrian Mineral Oil Administration (OMV group)
- **SE**: Standard Error
- **SR**: Sharpe Ratio

Today's Question and Seminar Structure

- Question of the seminar
  - **how do we measure, from real data, how much an investment grew and how risky it was?**
- Exercises: in the order of the Exercise Map (next slide)
- Outputs to produce
  - the answers to the A exercises; the table or figure each B exercise asks for
  - a short argued answer to C1; the errors found in the AI answer of C2 and the corrected answer
  - the interpretation sentences each exercise asks for
- Nothing is handed in: the seminar is for practice; the solutions of the [Proposed] exercises are discussed in class
- Notebook for today: open the seminar notebook in Google Colab; each task names its section
- Contact: Prof. Daniel Traian Pele, danpele@ase.ro

Exercise Map

Task	Question	Type	Model
A1, A2	do simple and log returns add up over time?	Solved, Proposed	A1
A3, A4	how do a dividend, a split and a consolidation change past prices?	Solved, Proposed	A3
A5	from daily moments to annual mean, volatility, CAGR and drag	Solved	–
A6, A7	Sharpe, Sortino, MDD and Calmar from a short price path	Solved, Proposed	A6
B1	can we trust a EUR/RON series?	Solved	–
B2, B3	is the mean return of the BET (and of other markets) significantly positive?	Solved, Proposed	B2
B4, B5	is the Sharpe ratio gap between TLV and the BET significant?	Solved, Proposed	B4
B6, B7	what do drawdowns and the volatility drag show on real data?	Solved, Proposed	B6
C1	do Sharpe ratios persist?	Proposed	B4, B6
C2	what is wrong in an AI answer?	Proposed	–

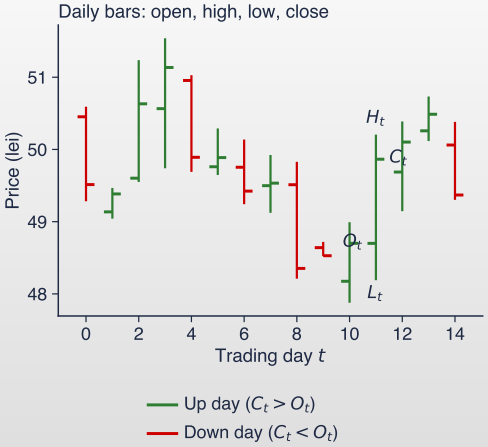
- **[Solved]**: full solution in the slides and in the notebook, a model to follow; **[Proposed]**: you solve it, following the model
- **A**: calculations on paper; **B**: real data, in the notebook; **C**: open questions and the critique of an AI answer

Data Used

Series	Source	Price used	Period
BET, S&P 500, DAX	EODHD	close	2015–2026
BET (B6)	EODHD, official BVB values	close	1997–2026
Bitcoin	EODHD	close, 7 days a week	2015–2026
TLV, SNP, BRD, TGN, SNG, SNN	EODHD	adjusted close	2015–2026
EUR/RON reference rate	BNR	daily fixing	2005–2026
EUR/RON series (B1 only)	EODHD	close	2005–2026

- Weekends and repeated holiday closes are dropped (except Bitcoin); each series on its own calendar, joint analyses on common days
- In the notebook: `load_close('bet')`, `log_returns('sp500')`; no account or key is needed

Prerequisites for Today (1/14): the Daily Bar



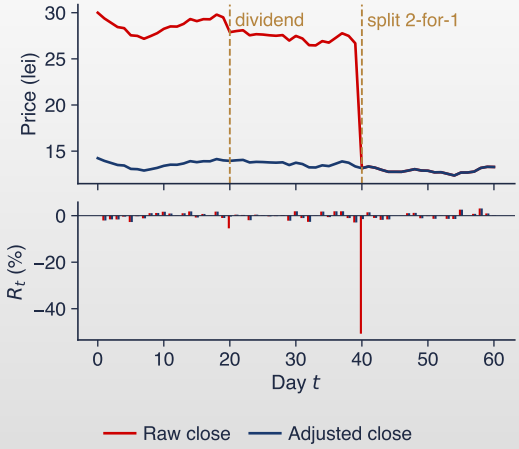
- **Daily bar** of day t: open O_t, high H_t, low L_t, close C_t and volume
  - O_t: first trade of the day; C_t: last trade (or the closing auction)
  - volume: the number of shares (or the value) traded that day
- A valid bar always satisfies

$$L_t \leq \min(O_t, C_t) \leq \max(O_t, C_t) \leq H_t$$

- The price used in this seminar
  - $P_t = C_t$ for indices, exchange rates and crypto assets
  - P_t = adjusted close for stocks (next slide)
- In the chart: green bars close above the open, red bars below it

SFM_ch1_seminar

Prerequisites for Today (2/14): Corporate Actions and Adjusted Prices



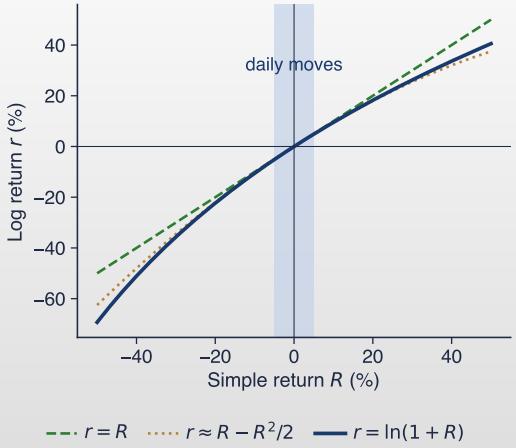
- A **corporate action** changes the price, not the value of a holding
  - cash dividend D on day t: $F = 1 - D/P_{t-1}$, with P_{t-1} the close of the day before
  - split, m new shares for n old ones: $F = n/m$ (2-for-1: $F = 1/2$)
  - consolidation of n old shares into one: $F = n$
- **Adjusted price** of an earlier day s

$$P_s^{\text{adj}} = P_s \prod_{e:\ t_e > s} F_e$$

- t_e, F_e: the day and the factor of event e; the product runs over the events after day s
  - prices after the last event stay unchanged
- In the chart: the raw close falls by half at the split, a fall that never happened; the adjusted close has no jump

SFM_ch1_seminar

Prerequisites for Today (3/14): Simple and Log Returns



- **Simple** and **log return** of day t

$$R_t = \frac{P_t}{P_{t-1}} - 1, \qquad r_t = \ln \frac{P_t}{P_{t-1}} = \ln(1 + R_t)$$

- ln: the natural logarithm; back from r_t to R_t: $R_t = e^{r_t} - 1$
  - $R_t = 0.01$: the price rose by 1% over the day
- Second-order approximation (Taylor expansion of $\ln(1 + R)$)

$$r_t \approx R_t - \frac{R_t^2}{2}$$

- so always $r_t < R_t$ (for $R_t \neq 0$), with a gap of about $R_t^2/2$
  - for daily moves (blue band, $|R_t| < 5\%$) the two are almost equal

SFM_ch1_seminar

Prerequisites for Today (4/14): Returns over Time and across Assets

- Over k periods: $R_t(k)$ and $r_t(k)$ are the returns from P_{t-k} to P_t
  - $j = 0, \ldots, k-1$ counts the days back from t
- Simple returns **compound**, log returns **add up**

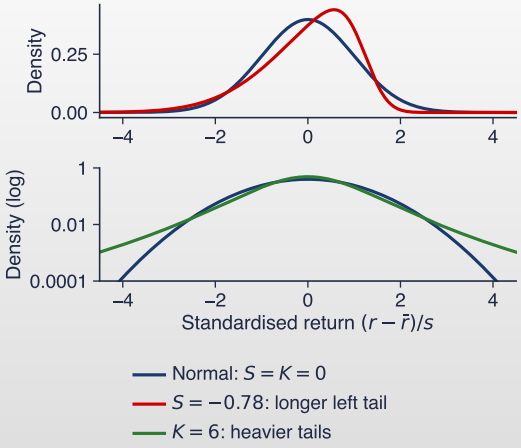
$$1 + R_t(k) = \prod_{j=0}^{k-1}(1 + R_{t-j}) = \frac{P_t}{P_{t-k}}, \qquad r_t(k) = \sum_{j=0}^{k-1} r_{t-j} = \ln \frac{P_t}{P_{t-k}}$$

- **Portfolio**: weights w_i, the shares of wealth in asset i, with $\sum_i w_i = 1$, set at the start of the period
  - R_i, r_i: the simple and the log return of asset i over the period
- Simple returns **add across assets**, log returns do not

$$R_p = \sum_i w_i R_i, \qquad r_p = \ln(1 + R_p) = \ln\left(\sum_i w_i e^{r_i}\right) \neq \sum_i w_i r_i$$

- Rule: log returns for sums over time, simple returns for portfolios

Prerequisites for Today (5/14): Mean, Standard Deviation and Shape



- From n returns $r_1, \ldots, r_n$: sample mean and standard deviation

$$\bar{r} = \frac{1}{n}\sum_{t=1}^{n} r_t, \qquad s = \sqrt{\frac{1}{n-1}\sum_{t=1}^{n}(r_t - \bar{r})^2}$$

- Central moments $m_k = \frac{1}{n}\sum_t (r_t - \bar{r})^k$; shape measures

$$S = \frac{m_3}{m_2^{3/2}}, \qquad K = \frac{m_4}{m_2^2} - 3$$

- S: **skewness**; $S < 0$: a longer left tail (large falls)
  - K: **excess kurtosis**; $K > 0$: heavier tails than the Normal distribution
  - for the Normal distribution $S = K = 0$; the bottom panel uses a log scale to show the tails

SFM_ch1_seminar

Prerequisites for Today (6/14): Annualisation and CAGR

- A: the number of observations per year
  - about 252 for an exchange, 365 for Bitcoin, 12 for monthly data; from the data: $A = n/Y$
  - Y: the length of the sample in years
- **Annual mean** and **annual volatility**

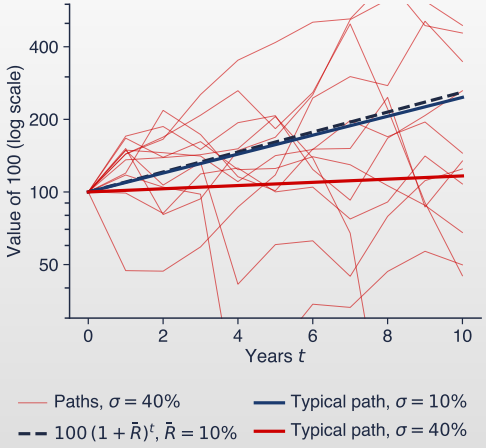
$$\mu_{\textbf{year}} = A\,\bar{r}, \qquad \sigma_{\textbf{year}} = \sqrt{A}\,s$$

- for uncorrelated returns the variances of the A periods add up: $\mathrm{Var}(r_1 + \cdots + r_A) = As^2$
  - so the standard deviation grows only with $\sqrt{A}$
- **CAGR** (compound annual growth rate)

$$\mathrm{CAGR} = \left(\frac{P_T}{P_0}\right)^{1/Y} - 1 = e^{\mu_{\textbf{year}}} - 1$$

- P_0, P_T: the first and the last price
  - the constant yearly rate that turns P_0 into P_T; $\mu_{\textbf{year}}$ is the annual mean of the log returns

Prerequisites for Today (7/14): Volatility Drag



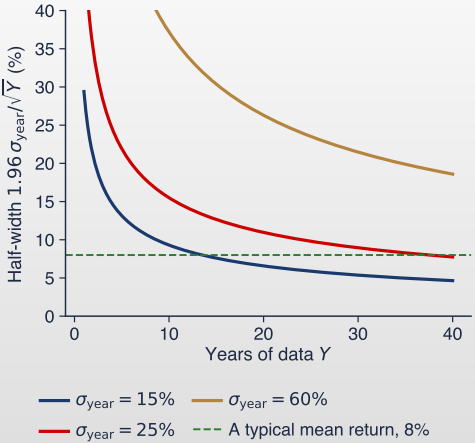
- Annual mean of simple returns $\bar{R}$ and annual mean of log returns μ_{year}

$$\bar{R} \approx \mu_{\text{year}} + \frac{\sigma^2}{2}$$

- σ: the annual volatility; $\sigma^2/2$: the **volatility drag**
  - it follows from $r \approx R - R^2/2$, averaged over the year
  - the growth of wealth (CAGR) follows μ_{year}, not $\bar{R}$
- In the chart: two assets with the same $\bar{R} = 10\%$ a year
  - $\sigma = 10\%$: the typical path (median) grows almost like $(1+\bar{R})^t$
  - $\sigma = 40\%$: the drag is $0.4^2/2 = 8\%$ a year, and the typical path barely grows

SFM_ch1_seminar

Prerequisites for Today (8/14): Standard Error and Confidence Interval



- **Standard error** (SE) of the annual mean, from Y years of data

$$\text{SE} = \frac{\sigma_{\text{year}}}{\sqrt{Y}}$$

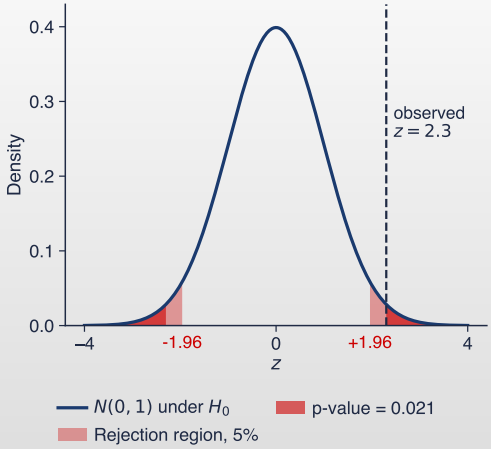
- SE: the typical error of the estimated mean around the true mean
  - it shrinks only with more years, not with more frequent data
- **95% confidence interval** (CI)

$$\hat{\mu} \pm 1.96\,\text{SE}$$

- $\hat{\mu}$: the estimate; $1.96 = \Phi^{-1}(0.975)$, the 97.5% quantile of the standard Normal distribution
  - Φ: the CDF (cumulative distribution function) of $N(0,1)$
- In the chart: with $\sigma_{\text{year}} = 25\%$, even 40 years give a half-width of 7.7%, about a typical mean return

SFM_ch1_seminar

Prerequisites for Today (9/14): The z Test and the p-value



- Test of H_0: mean $= 0$, against the alternative mean $\neq 0$

$$z = \frac{\hat{\mu}}{\text{SE}}$$

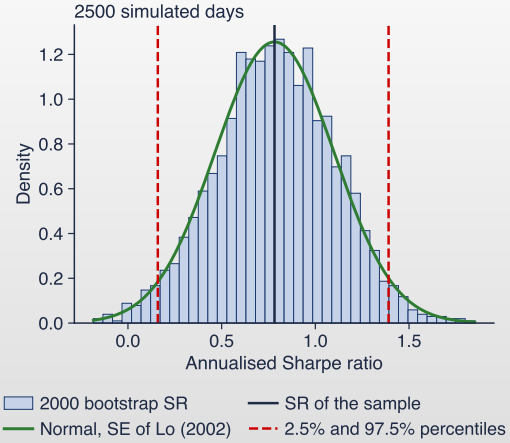
- z: the estimate measured in standard errors; under H_0, $z \approx N(0, 1)$
  - reject H_0 at 5% if $|z| > 1.96$ (red regions)
- **p-value**: the probability, under H_0, of a $|z|$ at least as large as the observed one

$$p = 2\left(1 - \Phi(|z|)\right)$$

- $p < 0.05$ exactly when $|z| > 1.96$
  - in the chart: $z = 2.3$ gives $p = 0.021$, so H_0 is rejected at 5%

SFM_ch1_seminar

Prerequisites for Today (10/14): the Sharpe Ratio and Its Precision



- **Sharpe ratio** Sharpe (1994): excess return per unit of volatility

$$\mathrm{SR} = \frac{\mu - r_f}{\sigma}, \qquad \mathrm{SR}_{\mathbf{year}} = \sqrt{A}\,\mathrm{SR}$$

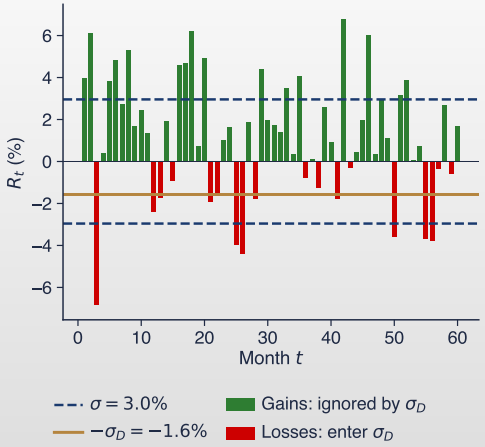
- μ, σ: mean and volatility of the returns; r_f: the risk-free rate (here $r_f = 0$)
  - the annual SR is $\sqrt{A}$ times the daily one (A for the mean divided by $\sqrt{A}$ for the volatility)
- SE of the per-period SR for i.i.d. returns Lo (2002)

$$\mathrm{SE} = \sqrt{\frac{1 + \mathrm{SR}^2/2}{n}}$$

- i.i.d.: independent and identically distributed; n: number of returns; times $\sqrt{A}$ for the annual SR
  - **bootstrap**: draw n days with replacement, recompute SR; repeat B times; the 2.5% and 97.5% percentiles give the 95% interval

SFM_ch1_seminar

Prerequisites for Today (11/14): the Sortino Ratio



- **Downside deviation**: only the losses count

$$\sigma_D = \sqrt{\frac{1}{n}\sum_{t=1}^{n} \min(R_t, 0)^2}$$

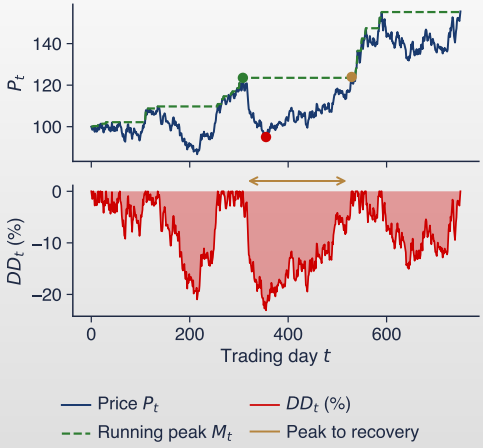
- $\min(R_t, 0)$: equals R_t for a loss and 0 for a gain
  - σ_D measures the risk of falls, not of gains
- **Sortino ratio** Sortino & van der Meer (1991), annualised like the Sharpe ratio

$$\text{Sortino} = \sqrt{A}\,\frac{\bar{R}}{\sigma_D}$$

- $\bar{R}$: the mean return per period ($r_f = 0$)
  - few or small losses: $\sigma_D < \sigma$, so Sortino $>$ Sharpe
  - in the chart: the green gains count in σ, but not in σ_D

SFM_ch1_seminar

Prerequisites for Today (12/14): Drawdown



- **Running peak** and **drawdown** on day t

$$M_t = \max_{s \le t} P_s, \qquad \mathrm{DD}_t = \frac{P_t}{M_t} - 1 \le 0$$

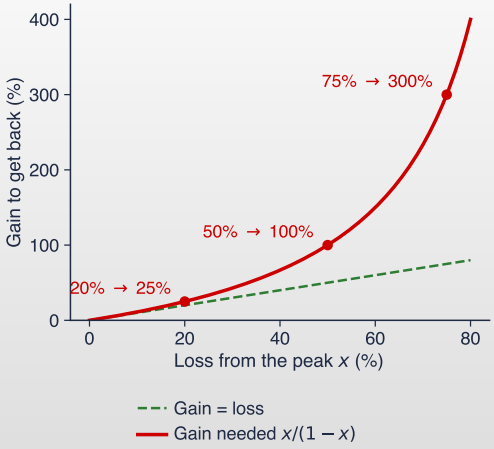
- M_t: the highest price up to day t; $\mathrm{DD}_t = 0$ on every new maximum
  - the red area: the investor is **underwater**, below the earlier peak
- **MDD** (maximum drawdown): the worst peak-to-trough fall

$$\mathrm{MDD} = \min_t \mathrm{DD}_t$$

- Three dates describe it
  - the peak (green), the trough (red) and the recovery (amber), the first day with $P_t \ge$ the old peak
  - the amber arrow: the time from the peak to the recovery

SFM_ch1_seminar

Prerequisites for Today (13/14): Recovery and the Calmar Ratio



- After a loss x from the peak, the gain needed to get back

$$g = \frac{x}{1-x}$$

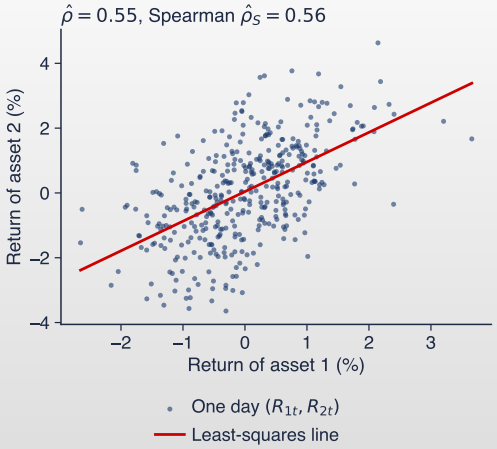
- from $(1-x)(1+g) = 1$: the gain applies to a smaller base
  - $x = 20\%$ needs 25%; $x = 50\%$ needs 100%; $x = 75\%$ needs 300%
- **Calmar ratio**: growth per unit of the worst fall

$$\text{Calmar} = \frac{\text{CAGR}}{|\text{MDD}|}$$

- both over the same period; a higher value is better
  - the volatility ignores the order of returns; the MDD and the Calmar ratio depend on it

SFM_ch1_seminar

Prerequisites for Today (14/14): Correlation and Rank Correlation



- **Correlation** of two return series on the same days

$$\hat{\rho} = \frac{\sum_t (x_t - \bar{x})(y_t - \bar{y})}{\sqrt{\sum_t (x_t - \bar{x})^2 \sum_t (y_t - \bar{y})^2}}$$

- $\hat{\rho} \in [-1, 1]$: the strength of the linear link; 0 means no linear link
  - correlated series: resample the days **in pairs** in the bootstrap, to keep $\hat{\rho}$
- **Spearman rank correlation**: the same formula applied to the ranks

$$\hat{\rho}_S = 1 - \frac{6 \sum_i d_i^2}{n(n^2 - 1)}$$

- d_i: the difference between the two ranks of item i; n: the number of items
  - rank 1 = the smallest value; $\hat{\rho}_S = 1$: identical rankings, -1: reversed rankings

SFM_ch1_seminar

A1: Simple and Log Returns over Two Days [Solved]

Task

- A share closes at 50, 55 and 48 lei on three consecutive days.
- 1. Compute the two daily simple returns and the two daily log returns.
- 2. Compute the two-day simple return from the prices and by compounding.
- 3. Check whether the sum of the simple returns and the sum of the log returns equal the two-day returns.
- Report: four daily returns, two two-day returns, one sentence on which returns add up.

Solution

- 1. $R_1 = 55/50 - 1 = 10.00\%$, $R_2 = 48/55 - 1 = -12.73\%$; $r_1 = \ln(55/50) = 9.53\%$, $r_2 = \ln(48/55) = -13.61\%$
- 2. $48/50 - 1 = -4.00\%$; $(1 + R_1)(1 + R_2) - 1 = -4.00\%$, the same
- 3. $R_1 + R_2 = -2.73\% \neq -4.00\%$; $r_1 + r_2 = -4.08\% = \ln(48/50)$
- Log returns add up over time; simple returns must be compounded.

Back Exercise list

A2: A Three-Day Path [Proposed]

Solution discussed in class

Task

- A share closes at 80, 100, 75 and 90 lei on four consecutive days. Model: A1.
- 1. Compute the three daily simple returns and the three daily log returns.
- 2. Compute the three-day simple return in two ways: from the first and last price, and by compounding.
- 3. Explain why the sum of the simple returns overstates the three-day return.
- Report: six daily returns, the three-day simple and log returns, one sentence of explanation.

A3: A Dividend and a Split [Solved]

Task

- Closes on days 0–6: 30, 31, 30.5, 29.2, 29.8, 10.1, 10.3 lei. On day 3 the share pays a dividend of 1.5 lei; on day 5 it splits 3-for-1.
- 1. Compute the adjustment factor of each event.
- 2. Compute the adjusted prices of days 0–6.
- 3. Compare the raw and the adjusted return of day 5.
- Report: two factors, seven adjusted prices, two returns of day 5.

Solution

- 1. Dividend: $F = 1 - 1.5/30.5 = 0.9508$; split: $F = 1/3 = 0.3333$
- 2. Days 0–2 are multiplied by both factors, 0.3169: 9.51, 9.83, 9.67; days 3–4 by $1/3$: 9.73, 9.93; days 5–6 unchanged: 10.10, 10.30
- 3. Raw: $10.1/29.8 - 1 = -66.1\%$, a fall that never happened; adjusted: $10.10/9.93 - 1 = 1.7\%$

Back Exercise list

A4: A Dividend and a Consolidation [Proposed]

Solution discussed in class

Task

- Closes on days 0–5: 2.05, 2.10, 1.95, 1.98, 19.90, 20.40 lei. On day 2 the share pays a dividend of 0.18 lei; on day 4, ten old shares are consolidated into one. Model: A3.
- 1. Compute the adjustment factor of each event.
- 2. Compute the adjusted prices of days 0–5.
- 3. Compute the raw log return of day 4 and say what a volatility estimate would look like if it were kept.
- Report: two factors, six adjusted prices, one return and one sentence.

Back Exercise list

A5: From Daily Moments to Annual Indicators [Solved]

Task

- A stock has daily log returns with mean 0.06% and standard deviation 1.5%; it trades 250 days a year; you have 10 years of data.
- 1. Compute the annual mean log return and the annual volatility.
- 2. Compute the CAGR and the annual mean of simple returns implied by the volatility drag.
- 3. Compute the standard error and the 95% CI of the annual mean log return.
- Report: five numbers and one sentence on the width of the interval.

Solution

- 1. $250 \times 0.06\% = 15.0\%$; $\sqrt{250} \times 1.5\% = 23.7\%$
- 2. CAGR $= e^{\mathbf{0.15}} - 1 = 16.2\%$; $\sigma^{\mathbf{2}}/2 = 0.0563/2 = 2.81\%$, so the arithmetic mean is $15.0\% + 2.81\% \approx 17.8\%$
- 3. SE $= 23.7\%/\sqrt{10} = 7.5\%$; CI $[0.3\%, 29.7\%]$
- Ten years of data barely exclude zero: mean returns are imprecise.

Back Exercise list

A6: Indicators of a Short Price Path [Solved]

Task

- A fund's month-end values: 100, 104, 101, 107, 99, 103, 110 (six months, $A = 12$, $r_f = 0$).
- 1. Compute the six monthly simple returns, their mean and standard deviation.
- 2. Compute the annualised Sharpe ratio and Sortino ratio.
- 3. Compute the MDD, the CAGR and the Calmar ratio.
- Report: a table of returns, five indicators, one sentence on why Sortino exceeds Sharpe.

Solution

- 1. Returns (%): $+4.00$; -2.88; $+5.94$; -7.48; $+4.04$; $+6.80$; mean 1.74%, s.d. 5.66%
- 2. SR $= \sqrt{12} \times 1.74/5.66 = 1.06$; $\sigma_D = 3.27\%$, Sortino $= \sqrt{12} \times 1.74/3.27 = 1.84$
- 3. MDD $= 99/107 - 1 = -7.48\%$; CAGR $= 1.10^2 - 1 = 21.0\%$; Calmar $= 2.81$
- Only two of six months are losses, so $\sigma_D < \sigma$; with half a year of data all these numbers are very imprecise.

Back Exercise list

A7: Indicators of Another Fund [Proposed]

Solution discussed in class

Task

- Month-end values: 200, 190, 205, 215, 180, 196, 210, 222 (seven months, $A = 12$, $r_f = 0$). Model: A6.
- 1. Compute the monthly simple returns, their mean and standard deviation.
- 2. Compute the annualised Sharpe and Sortino ratios.
- 3. Compute the MDD, the CAGR and the Calmar ratio.
- 4. Compare with the fund of A6: which indicators rank the two funds differently?
- Report: five indicators and one sentence comparing the two funds.

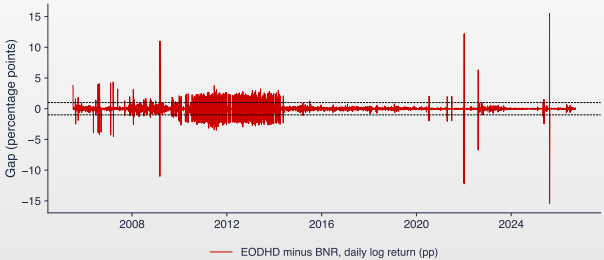
Back Exercise list

B1: Can We Trust a EUR/RON Series? [Solved]

- **Question**: do the EUR/RON series from EODHD and the BNR reference rate describe the same exchange rate?
- **Inputs**: the two daily EUR/RON series on their common days, 2005–2026
- **Tasks**
  1. Compute the daily log returns of both series on the common days.
  2. Count the days on which the two returns differ by more than one percentage point.
  3. Compute the annualised volatility of each series, and of the EODHD series without those days.
  4. Interpretation: which series is suitable for measuring the currency risk of a EUR loan? Justify the choice.
- **Report**: the number of days, three volatilities and one sentence of interpretation
- Notebook: section B1

Back Exercise list

B1: Solution [Solved]



- 5 154 common days; on 412 of them (8.0%) the gap exceeds one percentage point; most in 2011 (107 days)
- Largest gaps: 13 August 2025 (EODHD 15.4%, BNR −0.10%), 14 August 2025 (EODHD −15.4%, BNR −0.01%)
- Volatility: BNR 4.5%, EODHD 13.4%, EODHD without the bad days 4.7%
- Interpretation: the BNR rate, the primary source of the number; the EODHD series would multiply the measured currency risk by 3.0

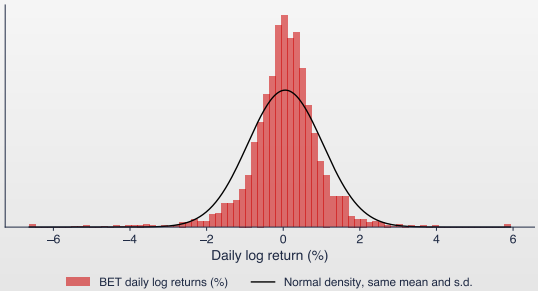
SFM_ch1_seminar

B2: Is the Mean Return of the BET Positive? [Solved]

- **Question**: is the mean daily return of the BET significantly different from zero over 2015–2026?
- **Inputs**: BET closes, 2015–2026, daily log returns in %
- **Tasks**
  1. Compute the mean, standard deviation, skewness, excess kurtosis, minimum (with its date) and maximum of the daily log returns.
  2. Annualise the mean and the standard deviation with the actual number of observations per year.
  3. Compute the standard error of the annual mean, its 95% CI and the z statistic of H_0: mean $= 0$.
  4. Interpretation: does the histogram look like the Normal distribution?
  5. Interpretation: what does the confidence interval tell an investor?
- **Report**: a table of statistics, the CI, z and the p-value, two sentences of interpretation
- Notebook: section B2

Back Exercise list

B2: Solution [Solved]



- $n = 2\,925$ days, 250 a year; mean 0.053%, s.d. 0.98%
- Skewness -1.41, excess kurtosis 17.8; worst day -11.9% on 19 December 2018, best 6.8%
- Annual: mean 13.2%, volatility 15.5%; SE $= 15.5/\sqrt{11.7} = 4.5\%$
- 95% CI [4.3%, 22.1%]; $z = 2.91$, p-value 0.004: we reject a zero mean at 5%
- Interpretation: a taller centre and heavier tails than the Normal distribution; the mean is positive, but anywhere between 4.3% and 22.1% a year

SFM_ch1_seminar

B3: Four More Markets [Proposed]

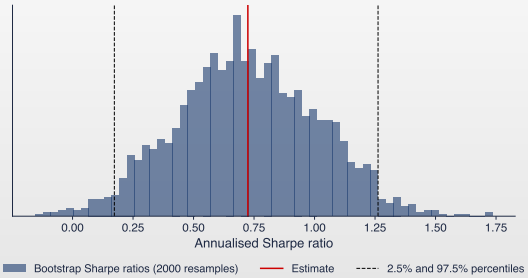
- **Question**: which of the S&P 500, DAX, Bitcoin and Banca Transilvania have a mean return significantly above zero?
- **Inputs**: daily log returns, 2015–2026, each on its own calendar; model: B2
- **Tasks**
  1. Repeat steps 1–3 of B2 for each of the four series.
  2. Put the annual mean, the volatility, the CI and the p-value of all five series (with the BET) in one table.
  3. Interpretation: why is the interval of Bitcoin so much wider than that of the S&P 500, although both cover the same years?
- **Report**: one table and two sentences
- Notebook: section B3

B4: How Precise Is a Sharpe Ratio? [Solved]

- **Question**: how precisely is the Sharpe ratio of the S&P 500 estimated from 2015–2026 data?
- **Inputs**: S&P 500 closes, daily simple returns, $r_f = 0$
- **Tasks**
  1. Compute the daily mean, standard deviation and Sharpe ratio, and annualise the Sharpe ratio.
  2. Compute the standard error of Lo (2002) and the 95% CI.
  3. Bootstrap the Sharpe ratio with 2000 resamples of the days and take the 2.5% and 97.5% percentiles.
  4. Interpretation: could the true Sharpe ratio of the S&P 500 be below 0.25?
- **Report**: the Sharpe ratio, two intervals and one sentence
- Notebook: section B4

Back Exercise list

B4: Solution [Solved]



- $n = 2\,943$, $A = 251.5$; daily mean 0.0508%, s.d. 1.113%, daily SR 0.0457
- Annual SR $= 0.0457 \times 15.86 = 0.72$
- Lo: SE $= \sqrt{(1 + 0.0457^2/2)/2\,943} \times 15.86 = 0.29$; CI [0.15, 1.30]
- Bootstrap: [0.17, 1.26], very close
- Interpretation: yes; every value from 0.17 to 1.26 is compatible with the data

SFM_ch1_seminar

B5: Is Banca Transilvania Better than the BET? [Proposed]

- **Question**: is the Sharpe ratio of TLV significantly higher than that of the BET index?
- **Inputs**: BET close and TLV adjusted close, 2015–2026, joined on common days first; model: B4
- **Tasks**
  1. Join the two price series on their common days, then compute daily simple returns.
  2. Compute both annualised Sharpe ratios with their Lo standard errors.
  3. Bootstrap the difference of the two Sharpe ratios, resampling the same days for both series.
  4. Interpretation: why must the days be resampled in pairs?
- **Report**: two Sharpe ratios, the difference with its interval, one sentence
- Notebook: section B5

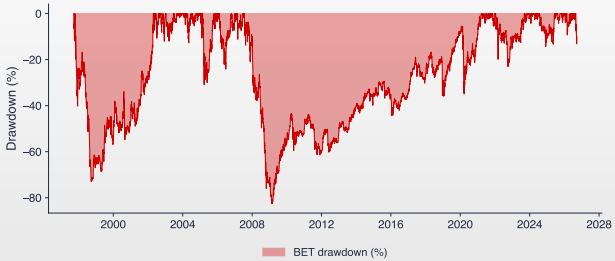
Back Exercise list

B6: The Deepest Fall of the BET [Solved]

- **Question**: what were the depth and the duration of the worst fall of the BET since its launch?
- **Inputs**: official BET closes since 19 September 1997; for the Calmar ratio, 2015–2026
- **Tasks**
  1. Compute the drawdown series and the maximum drawdown with the dates of the peak, the trough and the recovery.
  2. Compute the gain needed to recover from the trough and the share of days spent more than 20% below the peak.
  3. Compute the CAGR, the MDD and the Calmar ratio over 2015–2026.
  4. Interpretation: what does the drawdown add to what the volatility already tells us?
- **Report**: the MDD with three dates, two percentages, three indicators, one sentence
- Notebook: section B6

Back Exercise list

B6: Solution [Solved]



- MDD −82.5%: from 10 814 points on 24 July 2007 to 1 887 on 25 February 2009; the old peak was regained on 16 March 2021, after 13.6 years
- From the trough, a gain of 473% was needed; 56% of all days were more than 20% below the peak
- 2015–2026: CAGR 14.1%, MDD −31.1% (23 January 2020 to 23 March 2020), Calmar $= 14.1/31.1 = 0.45$
- Interpretation: volatility ignores the order of returns; the drawdown shows how long an investor stays under water

SFM_ch1_seminar

B7: Bitcoin, OMV Petrom and the Volatility Drag [Proposed]

- **Question**: what do the drawdowns and the volatility drag show for Bitcoin and SNP?
- **Inputs**: Bitcoin close and SNP adjusted close, 2015–2026; model: B6
- **Tasks**
  1. Compute the MDD of each series with the dates of the peak, the trough and the recovery.
  2. Compute the CAGR and the Calmar ratio of each series.
  3. Compute the arithmetic annual mean of simple returns, the annual mean of log returns and their difference.
  4. Compare the difference with $\sigma^2/2$, using the annual volatility of simple returns.
  5. Interpretation: for which asset does the drag matter for a long-term investor?
- **Report**: a table with MDD, dates, CAGR, Calmar, the two means, the drag and $\sigma^2/2$, one sentence
- Notebook: section B7

C1: Do Sharpe Ratios Persist? [Proposed]

- **Question**: does a BVB stock with a high Sharpe ratio in 2015–2020 also have a high one in 2021–2026?
- **Inputs**: adjusted closes of TLV, SNP, BRD, TGN, SNG, SNN, 2015–2026; models: B4, B6
- **Tasks**
  1. Compute the annualised Sharpe ratio and its standard error for each stock in each half.
  2. Compute the Spearman rank correlation between the Sharpe ratios of the two halves.
  3. Propose one way to make the answer more reliable (more stocks, other indicators, other windows).
  4. Interpretation: what can six stocks tell us about persistence?
- **Report**: a table, the rank correlation with its p-value, a short plan for a project
- Notebook: section C1

Back Exercise list

C2: Audit an AI Answer [Proposed]

- A student asked an AI assistant: "Compare the risk and the performance of Bitcoin and the BET in 2015–2026." The answer:
- (a) Bitcoin's daily volatility is 3.49%, so its annual volatility is 3.49% x sqrt(252) = 55.4%.
- (b) Bitcoin's Sharpe ratio (1.05) is higher than the BET's (0.93), which proves that Bitcoin is the better investment.
- (c) Bitcoin's maximum drawdown is its worst daily return, −46.5%.
- (d) For a 50/50 portfolio, the log return is exactly the average of the two log returns.
- (e) For BVB stocks such as TLV I used the close column, which is the official price.
- **Tasks**
  1. For each statement, say whether it is correct; if not, give the correct statement and, where possible, the correct number from the notebook (section C2).
- **Report**: a list of five verdicts with one line of justification each.

Points to Remember

- Check the data before computing: source, adjusted prices, a second series
- Log returns add over time; simple returns add across assets
- Mean returns and Sharpe ratios come with wide intervals: report them
- The drawdown and the volatility drag show what the volatility alone hides
- An AI answer is a draft: check every factor, definition and column

After the Seminar

- Lecture 1 develops each topic of today: data quality, returns, annualisation, descriptive statistics and performance indicators
- C1 can grow into a team project: more stocks, rolling windows, other indicators
- Reading: Franke, Härdle & Hafner (2019), Ch. 11; exercises with solutions: Borak, Härdle & López-Cabrera (2013)

References

- Borak, S., Härdle, W.K., López-Cabrera, B. (2013). *Statistics of Financial Markets: Exercises and Solutions*, 2nd ed. Springer.
- Franke, J., Härdle, W.K., Hafner, C.M. (2019). *Statistics of Financial Markets: An Introduction*, 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.
- Lo, A.W. (2002). The statistics of Sharpe ratios. *Financial Analysts Journal*, 58(4), 36–52.
- Sharpe, W.F. (1994). The Sharpe ratio. *Journal of Portfolio Management*, 21(1), 49–58.
- Sortino, F.A., van der Meer, R. (1991). Downside risk. *Journal of Portfolio Management*, 17(4), 27–31.