








Modelling Financial Markets

Seminar 16: Digital Assets and DeFi — Hands-on



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

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

What measures digital-asset risk?

- Questions of the seminar
  - **can the main numbers on digital assets be reproduced from the data?**
  - **how precise are they?**
- Deliverables
  - the answers to the A exercises
  - one table and one figure per B exercise
  - the corrected C1 code
  - one interpretation sentence per exercise
- Assessment criteria
  - correct derivations
  - inference: SE (Standard Error), CI (Confidence Interval), decision
  - interpretation
  - critique of the AI (Artificial Intelligence) answer
- Notebook for today: open the seminar notebook in Colab
  - each task names its notebook section
- Seminar: Prof. Daniel Traian Pele, danpele@ase.ro; Antoaneta Amza, PhD candidate, antoaneta.amza@csie.ase.ro

Assets and Units

- Decentralised finance (DeFi): financial services through contracts executed on a blockchain
- A liquidity pool holds the reserves used for swaps; stablecoins are tokens pegged to the dollar (USD), e.g. USD Coin, Tether, Dai
- Tickers used
  - IBIT (iShares Bitcoin Trust, BlackRock) and ETHA (iShares Ethereum Trust): the spot ETFs (exchange-traded funds) on Bitcoin and on Ether
  - MSTR: the shares of Strategy Inc. (formerly MicroStrategy), listed on Nasdaq
- The formulas in What You Need for Today use decimal units; the exercise data use log returns in percent, $r_t = 100 \ln(P_t / P_{t-1})$
- Example: from 100 to 110, the simple return is 0.10 (10%) and the log return 0.0953 (9.53%)
- The risk-free rate uses the same units and the same period as the return
- Risk-free rate: in C1, DTB3 (the 3-month Treasury bill secondary-market rate) from FRED (Federal Reserve Economic Data); C3 uses the monthly rate of the Kenneth French Data Library

Data Used

Series	Price used	Calendar	Sample	Used in
Bitcoin, Ether, Solana	daily close	every day, 7 a week	2018 (Solana: Apr 2020) – 18 Sep 2026	B1, B2, B7
USD Coin, Tether, Dai	daily close and daily low	every day	2021 – 18 Sep 2026	B5, B6, A3
Spot ETFs (IBIT, ETHA)	adjusted close	US trading days	from launch (11 Jan 2024; 23 Jul 2024)	B2, B8
Strategy, Coinbase	adjusted close	Friday closes	16 Apr 2021 – 18 Sep 2026	B3, B4
S&P 500; Nasdaq 100 ETF (QQQ), gold ETF (GLD), long US Treasury ETF (TLT)	index: daily close; ETFs: adjusted close	US trading days	as the paired asset	B1, B3, C2
3-month Treasury bill	FRED DTB3	US business days	2021–2025	C1

- Daily prices of the course data, one series per asset (field used: close / adjusted close)
- Two series with different calendars (crypto 7 days, shares 5 days): align the *prices* on common trading days first, then compute returns

Setup: Run and Verify

```
import numpy as np, pandas as pd
URL = ('https://raw.githubusercontent.com/'
       'danpele/MFM/main/data/market/')
def load(sym, field='close'):
    return pd.read_csv(URL + sym + '.csv',
        index_col='date', parse_dates=True)[field]
btc = load('BTC-USD.CC')
btc = btc.loc['2018-01-01':'2026-09-18']
r = 100 * np.log(btc).diff().dropna()
print(len(r), r.index[0].date(), r.index[-1].date())
```

- Expected output: 3 182 2018-01-02 2026-09-18
- Crypto trades 7 days a week: 3 182 daily returns in 8.7 years, about 365 a year; annualise with $\sqrt{365}$, not $\sqrt{252}$
- Further checks: USD Coin 2 087 daily closes from 1 Jan 2021; IBIT 673 common trading days and 140 weeks; Strategy 284 weeks
- If a number of observations differs, stop and check the asset, the price used and the dates

Returns, Volatility and the Sharpe Ratio

- Details: Chapter 16
- **Log** and **simple** return of period t (P_t: price at the end of period t)

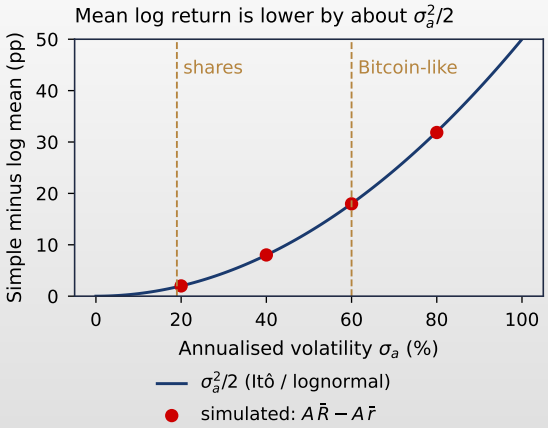
$$r_t = \ln \frac{P_t}{P_{t-\mathbf{1}}}, \qquad R_t = \frac{P_t}{P_{t-\mathbf{1}}} - 1 = e^{r_t} - 1$$

  - Log returns add over time; simple returns add across assets in a portfolio
- **Annualising** with A return periods per year: 365 for crypto (it trades every day), 252 for shares, 52 for weekly data

$$\mu_a = A\,\bar{R}, \qquad \sigma_a = \sqrt{A}\,s, \qquad SR = \frac{\mu_a - r_f}{\sigma_a}$$

  - $\bar{R}$, s: sample mean and standard deviation of the period returns; r_f: annual risk-free rate, e.g. a 3-month Treasury bill
  - SR: Sharpe ratio, excess return per unit of volatility

Log and Simple Mean Returns



- For a lognormal price the two means differ by half the variance

$$\mu_a^{\log} \approx \mu_a^{\text{simple}} - \tfrac{1}{2}\sigma_a^2$$

- Chart: the gap in percentage points against the annual volatility; points from simulated daily returns, $A = 365$
- Negligible for shares, large for crypto
  - $\sigma_a = 20\%$: 2 points; $\sigma_a = 60\%$: 18 points
  - A Sharpe ratio uses the mean simple return

A Market-Value Index and Its Divisor

- **Market-value index** (e.g. CRIX, the CRypto IndeX, Trimborn & Härdle (2018)): total value of the basket, scaled by a divisor

$$I_t = \frac{\sum_i P_{i,t}\, Q_{i,\tau}}{D_\tau}, \qquad w_{i,t} = \frac{P_{i,t} Q_{i,\tau}}{\sum_j P_{j,t} Q_{j,\tau}}$$

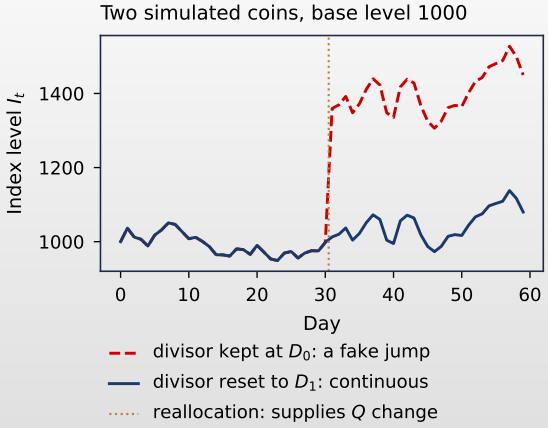
  - $P_{i,t}$: price of coin i on day t; $Q_{i,\tau}$: its circulating supply, fixed during the reallocation period τ (e.g. a month)
  - $w_{i,t}$: market-value weight; within a period the index return is $\sum_i w_{i,t-\mathbf{1}} R_{i,t}$

- **Divisor** D_τ: sets the level; reset at each reallocation

$$D_{\tau+\mathbf{1}} = \frac{\sum_i P_{i,t}\, Q_{i,\tau+\mathbf{1}}}{I_t} \quad \Longrightarrow \quad I_t \text{ unchanged by the new basket}$$

  - At a reallocation the basket changes, prices do not: no return may appear
  - Example: one coin, $P = 100$, $Q = 10$, $D = 1$: $I = 1\,000$; Q becomes 12 at the same price: $D = 1\,200/1\,000 = 1.2$ keeps $I = 1\,000$

Why the Divisor Must Change



- Two simulated coins; on day 30 the supply of the second rises from 50 to 80
- Divisor kept at D_0 (red)
  - The index jumps by about 34% on a day without any price move: a fake return
- Divisor reset to D_1 (blue)
  - The level is continuous; afterwards the new weights apply

A Constant-Product Pool

- **Automated market maker** (Uniswap v2, Adams, Zinsmeister & Robinson (2020)): a smart contract holds x Ether and y USD Coin and trades with anyone

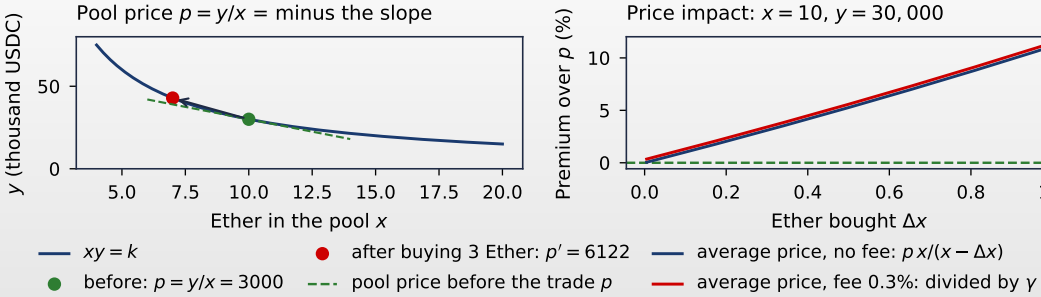
$$x\,y = k, \qquad p = \frac{y}{x}$$

  - k: the invariant, constant without fees; p: pool price of Ether in USD Coin, minus the slope of the curve $y = k/x$
- **A trade with a fee** f on the input, $\gamma = 1 - f$: paying Δy buys Δx

$$(x - \Delta x)(y + \gamma\,\Delta y) = k \quad\Longrightarrow\quad \Delta y = \frac{1}{\gamma}\Big(\frac{k}{x - \Delta x} - y\Big)$$

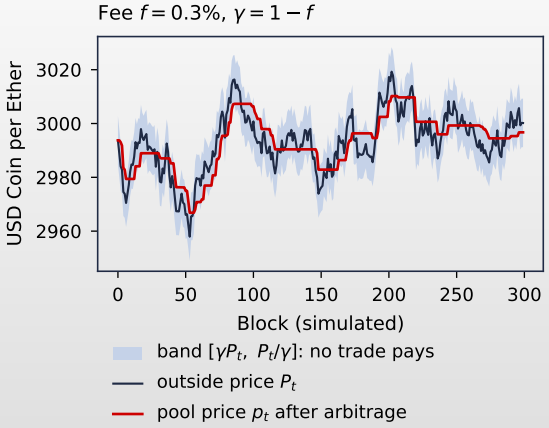
  - Δx: Ether leaving the pool; Δy: USD Coin paid in; $\gamma\Delta y$: the part that enters the invariant; the fee stays in the pool
  - Example: $x = 10$, $y = 30\,000$ ($p = 3\,000$), $f = 0.3\%$; buying 1 Ether: $\gamma\Delta y = 300\,000/9 - 30\,000 = 3\,333.33$, $\Delta y = 3\,343.36$

Price Impact in a Pool



- Left: buying 3 of the 10 Ether moves the pool along $xy = k$; the pool price doubles, from 3 000 to 6 122
- Right: without a fee the average price is $\Delta y/\Delta x = p\,x/(x - \Delta x)$, above p for any trade; the fee multiplies it by $1/\gamma$

Fees and the No-Arbitrage Band



- Outside price P: the price on the other markets; arbitrageurs trade against the pool when it pays
- Marginal prices of a tiny trade in the pool
  - Marginal buying and selling prices: you derive them in A2
  - No trade pays while $\gamma P \leq p \leq P/\gamma$
- Chart: simulated P_t; the pool price moves only when P_t pushes it out of the band
- Question for the class: does arbitrage bring the pool price exactly to the market price?
- Answer: no; the effective price reaches only the edge of the band: γP when the market price rises, P/γ when it falls
- The fee retained in the reserve leaves the reserve ratio just inside the band (A2); the chart tracks the effective price only

Stablecoins and a Threshold AR

- **Stablecoin**: a token meant to trade at 1 USD (USD Coin, Tether, Dai), redeemable or backed by reserves
  - The peg can break: USD Coin fell below 0.90 USD intraday in March 2023, when part of its reserves sat in a failed United States (US) bank; check the daily low, not only the close
  - Deviation in basis points (1 bp = 0.01%): $d_t = 10^4(P_t - 1)$
- **Threshold autoregressive (AR) model** (TAR, Balke & Fomby (1997)): arbitrage costs money, so small deviations may persist and large ones revert fast

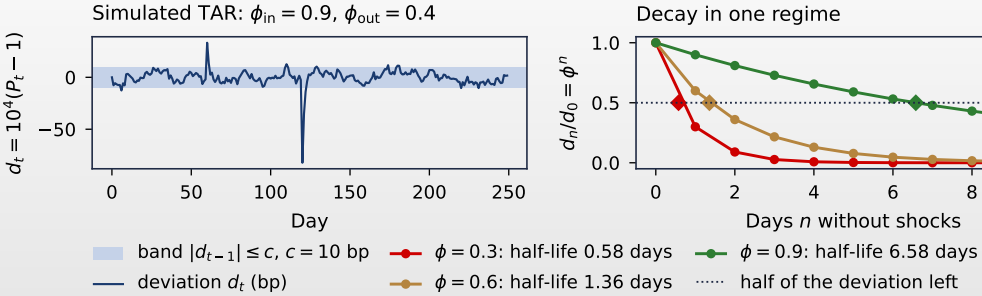
$$d_t = \begin{cases} \phi_{\text{in}}\, d_{t-1} + \varepsilon_t, & |d_{t-1}| \le c \text{ (inside the band)} \\ \phi_{\text{out}}\, d_{t-1} + \varepsilon_t, & |d_{t-1}| > c \end{cases}$$

  - c: threshold, the half-width of the band (bp); ϕ_{in}, ϕ_{out}: persistence inside and outside, $0 < \phi < 1$; ε_t: shock with mean 0
- **Half-life**: without shocks, $d_n = \phi^n d_0$ after n days in one regime

$$\phi^{h_{1/2}} = \tfrac{1}{2} \quad \Longrightarrow \quad h_{1/2} = \frac{\ln(1/2)}{\ln \phi}$$

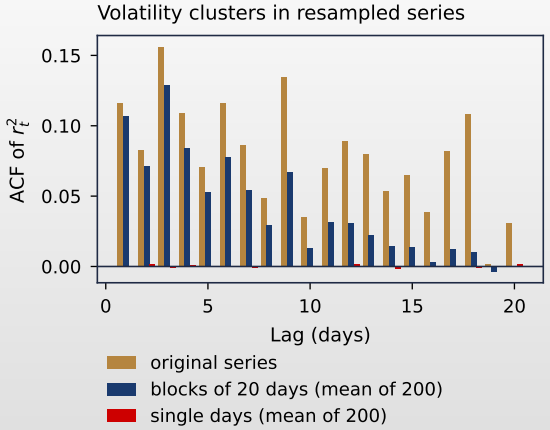
  - Example: $\phi_{\text{out}} = 0.5$, $d_0 = -100$ bp, $c = 20$ bp: $100 \cdot 0.5^n \le 20$ needs $n \ge \ln 0.2 / \ln 0.5 = 2.32$: inside the band on day 3; half-life 1 day

A Threshold AR in Pictures



- Left: inside the band ($\phi_{\text{in}} = 0.9$) the deviation wanders; after a large shock it returns fast ($\phi_{\text{out}} = 0.4$) until it re-enters the band
- Right: ϕ^n for three values of ϕ; the diamonds mark the half-lives $\ln 0.5 / \ln \phi$; a smaller ϕ means a faster return to par

Block Bootstrap and the Weekend Ratio



- **Moving-block bootstrap** Künsch (1989)
  - Glue random blocks of b consecutive days into a series of the original length; recompute the statistic; repeat B times
  - 95% CI: the 2.5% and 97.5% percentiles; blocks keep volatility clustering, single days destroy it (chart)
- **Weekend ratio** (B1)

$$\mathrm{VR} = \frac{s^2_{\text{Sat, Sun}}}{s^2_{\text{Mon–Fri}}}$$

  - Sample variances of the returns on weekend days and on weekdays; $\mathrm{VR} < 1$: calmer weekends
  - CI of a change: percentiles of $\mathrm{VR}^*_{\text{after}} - \mathrm{VR}^*_{\text{before}}$ from two independent bootstraps

Spot ETFs, Tracking and Two Stationarity Tests

- Spot Bitcoin ETF (IBIT, since 11 January 2024): a US-listed fund that holds Bitcoin; it trades on US exchange days, Bitcoin every day
- **Tracking**: log ratio of the ETF price to the Bitcoin price

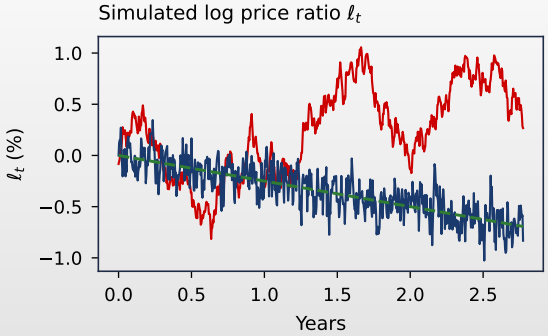
$$\ell_t = \ln \frac{P_t^{\mathrm{ETF}}}{P_t^{\mathrm{BTC}}} = a - \delta\, t + e_t$$

  - δ: fee drift, about 0.25% a year paid in Bitcoin BlackRock (2026); e_t: stationary timing noise; t in years
- **ADF (Augmented Dickey–Fuller)** Dickey & Fuller (1979): H_0 unit root

$$\Delta\ell_t = a + g\, t + b\, \ell_{t-1} + \textstyle\sum_{j=1}^{p} c_j \Delta\ell_{t-j} + u_t, \qquad t_b = \hat{b}/\mathrm{SE}(\hat{b})$$

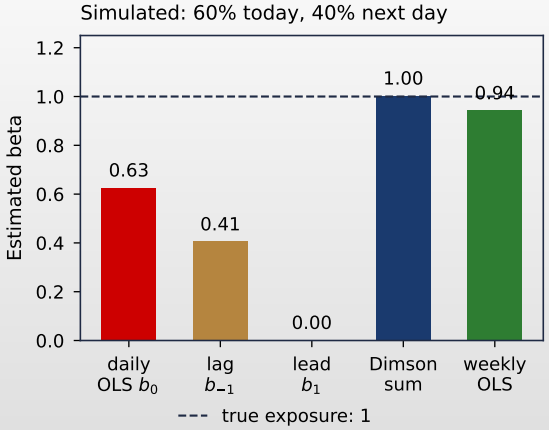
  - H_0: $b = 0$; reject if t_b is below the Dickey–Fuller critical value (about -3.4 at 5% with a trend)
- **KPSS (Kwiatkowski–Phillips–Schmidt–Shin test)** Kwiatkowski et al. (1992): H_0 trend stationarity
  - $\mathrm{KPSS} = \sum_t S_t^2/(T^2\hat{\sigma}^2_{\mathrm{LR}})$, S_t: partial sums of the residuals from a trend fit, $\hat{\sigma}^2_{\mathrm{LR}}$: their long-run variance; large values reject
  - Opposite null hypotheses: use both

Trend-Stationary or Unit Root?



- Simulated ℓ_t, both with the same fee trend of -0.25% a year
- Blue: trend + AR(1) noise
  - ADF rejects the unit root, KPSS does not reject stationarity: the two tests agree
  - The trend slope estimates the fee
- Red: trend + random walk
  - ADF does not reject, KPSS rejects: the drift is buried in noise

Asynchronous Closes and the Dimson Beta



- The ETF closes at 16:00 New York time, the daily Bitcoin price at another hour: part of a Bitcoin move reaches the ETF the next day
- **Dimson beta** Dimson (1979): add a lag and a lead

$$r_t^{\text{ETF}} = a + \sum_{j=-\mathbf{1}}^{\mathbf{1}} b_j \, r_{t+j}^{\text{BTC}} + u_t, \qquad \beta_D = b_{-\mathbf{1}} + b_{\mathbf{0}} + b_{\mathbf{1}}$$

  - $b_{\mathbf{0}}$: same-day slope, below 1 here; $b_{-\mathbf{1}}$: response to yesterday's Bitcoin return
- Chart: 60% of each move today, 40% the next day; the Dimson sum and weekly returns recover the exposure 1

HAC Standard Errors and the Wald Test

- **Regression** $y_t = \mathrm{x}_t^\top \boldsymbol{\beta} + u_t$ with OLS (Ordinary Least Squares) estimate $\hat{\boldsymbol{\beta}}$
  - y_t: a return; x_t: the regressors (a 1 and factor returns); u_t: error
- **HAC covariance** (heteroskedasticity and autocorrelation consistent; Newey–West Newey & West (1987)): robust to heteroskedasticity and autocorrelation up to lag L

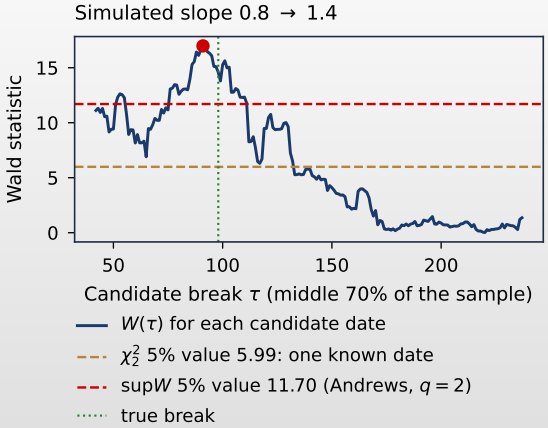
$$\hat{V} = Q^{-1}\hat{S}\,Q^{-1}, \quad Q = \sum_t \mathrm{x}_t\mathrm{x}_t^\top, \quad \hat{S} = \hat{\Gamma}_0 + \sum_{l=1}^{L}\Big(1 - \frac{l}{L+1}\Big)(\hat{\Gamma}_l + \hat{\Gamma}_l^\top)$$

  - $\hat{\Gamma}_l = \sum_t \hat{u}_t\hat{u}_{t-l}\mathrm{x}_t\mathrm{x}_{t-l}^\top$; $\hat{u}_t$: OLS residuals; SE: square roots of the diagonal of $\hat{V}$
- t-**test** of H_0: $\beta = b_0$: $t = (\hat{\beta} - b_0)/\mathrm{SE}$; $|t| > 1.96$ rejects at 5%
- **Wald test** of q restrictions H_0: $R\boldsymbol{\beta} = \mathrm{r}$

$$W = (R\hat{\boldsymbol{\beta}} - \mathrm{r})^\top (R\hat{V}R^\top)^{-1}(R\hat{\boldsymbol{\beta}} - \mathrm{r}) \;\sim\; \chi^2_q \text{ under } H_0$$

  - R: $q \times k$ matrix that picks the restricted combinations; large W: the estimates are far from H_0, measured in standard errors

A Break at an Unknown Date



- Break at τ: $\boldsymbol{\beta_1}$ up to τ, $\boldsymbol{\beta_2}$ after; $W(\tau)$: Wald test of $\boldsymbol{\beta_1} = \boldsymbol{\beta_2}$
  - Known date: compare $W(\tau)$ with χ^2_q, q = number of coefficients that may change
- Unknown date Andrews (1993)

$$\sup W = \max_{\mathbf{0.15}T \le \tau \le \mathbf{0.85}T} W(\tau)$$

  - The maximum of many W exceeds χ^2_q by chance: own critical values, e.g. 13.95 for $q = 3$ (7.81 for χ^2_3)
  - Date $\hat{\tau}$: smallest total sum of squared residuals; its 95% CI Bai (1997)
- Chart: one simulated slope change; $W(\tau)$ crosses 5.99 at many dates, sup W is judged against 11.70 (the Andrews critical value for $q = 2$; a simulated grid gives 11.55)

After the Lecture: Testing a Threshold

- Estimating c (B5): grid of candidate thresholds; for each, OLS in both regimes; keep the c with the smallest total sum of squared residuals (SSR)
  - Each regime must keep at least 15% of the days
  - White SE: the HAC formula with $L = 0$, robust to heteroskedasticity only
- Test of $\phi_{\text{in}} = \phi_{\text{out}}$

$$W(c) = \frac{(\hat{\phi}_{\text{in}} - \hat{\phi}_{\text{out}})^2}{\hat{v}_{\text{in}} + \hat{v}_{\text{out}}}, \qquad \sup_c W(c)$$

  - $\hat{v}$: squared White SE of each slope; under H_0, c is not identified, so χ^2_1 is wrong (the Davies problem)
  - Fixed-regressor wild bootstrap Hansen (1996): $d_t^* = \hat{e}_t \eta_t$, $\eta_t \sim N(0, 1)$, same d_{t-1}; recompute sup W^*

For the Extension Packs: Pools and Hedges

- **Weighted pool** (Balancer type, A5–A7): $x^{w} y^{\mathbf{1}-w} = k$, pool price $p = \frac{w}{\mathbf{1}-w}\frac{y}{x}$, w: value weight of Ether
  - Impermanent loss after a price ratio $r = P_{\mathbf{1}}/P_{\mathbf{0}}$: pool value over the value of holding the initial tokens, minus 1
  - Price $dP_t = \mu P_t\, dt + \sigma P_t\, dW_t$ (geometric Brownian motion); Itô's lemma: $dV = V' dP + \frac{\mathbf{1}}{\mathbf{2}} V'' \sigma^{\mathbf{2}} P^{\mathbf{2}} dt$
- **Hedge and safe haven** (C2) Baur & Lucey (2010)
  - Correlation $\rho = \mathrm{Cov}(r_{\mathbf{1}}, r_{\mathbf{2}})/(\sigma_{\mathbf{1}}\sigma_{\mathbf{2}})$; Forbes–Rigobon Forbes & Rigobon (2002): $\rho^{*} = \rho/\sqrt{1+\delta(1-\rho^{\mathbf{2}})}$, δ: relative rise of the market variance
  - Quantile regression Koenker & Bassett (1978): $\beta(\tau)$ minimises $\sum_t \rho_\tau(y_t - x_t^{\top}\boldsymbol{\beta})$, $\rho_\tau(u) = u(\tau - 1\{u<0\})$
- More in Chapter 16: stylised facts of crypto returns, stablecoin runs, loss-versus-rebalancing (LVR) and the crypto factor case study

Exercise Map

Exercise	Question	Type	Model
A1	how does a divisor keep an index continuous?	Solved	–
A2	when does arbitrage against a pool pay, and how much?	Solved	–
B1	is Bitcoin's weekend lull new since the ETFs?	Solved	–
A3	how many days does USD Coin need to return inside its band?	Solved	–
B2	does IBIT track Bitcoin one for one, and is the fee visible?	Solved	–
B3	when did Strategy's Bitcoin beta change?	Solved	–
C1	what is wrong in an AI answer?	Proposed	B1
B4	when did Coinbase's Bitcoin beta change? (homework)	Proposed	B3
B5	is the USD Coin peg a threshold process? (after the lecture)	Solved	A3
Extension packs	A4–A8, B6–B8, C2, C3 (optional, after the lecture)	Solved and Proposed (marked on each task)	named on each task

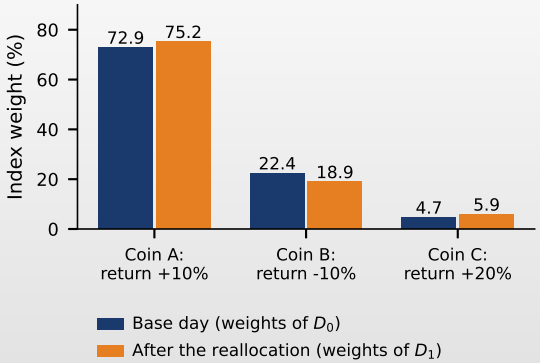
- **[Solved]**: full solution in the slides and in the notebook, a model to copy; **[Proposed]**: you solve it, the solution is discussed in class
- Part A: derivations on paper; Part B: estimation and inference on real data; Part C: open questions and the critique of an AI answer

A1 [Solved]: A Three-Coin Index and Its Divisor — Task

- **Question**: how does a market-value index stay continuous when its quantities change?
- **Data and inputs**
  - Coins A, B, C on the base day: prices $P_0 = (60\,000;\ 3\,000;\ 150)$ USD, circulating supplies $Q = (19.5;\ 120;\ 500)$ million
  - Market-value (Laspeyres) index I_t with divisor D_τ (index divisor, What You Need for Today); quantities fixed within the month
- **Tasks**
  1. Compute the market values, the weights and the divisor D_0 that sets $I_0 = 1000$.
  2. At month end the prices are $P_1 = (66\,000;\ 2\,700;\ 180)$: compute I_1 and check that $I_1/1000 - 1$ is the weighted average of the three returns.
  3. At the reallocation the supplies become $(19.52;\ 120.1;\ 560)$ million: compute the new divisor D_1 that keeps I_1 unchanged, and the new weights.
- **Report**: D_0, I_1, D_1 and both sets of weights, with every intermediate product
- **Interpretation**
  - why must the divisor change at a reallocation?
  - what would an investor see without the change?
- Notebook: section A1

Back Exercise list

A1: Solution



- **1.** $M_i = P_{i,\mathbf{0}} Q_i$: 1 170, 360, 75 billion USD, total 1 605; weights $w_i = M_i / \sum_j M_j$: 72.9%, 22.4%, 4.7%; $D_{\mathbf{0}} = \sum_j M_j / 1000 = 1.605 \cdot 10^{\mathbf{9}}$
- **2.** $I_{\mathbf{1}} = \sum_i P_{i,\mathbf{1}} Q_i / D_{\mathbf{0}} = 1\,059.81$; returns +10%, −10%, +20%; check: $72.9\% \cdot 10\% - 22.4\% \cdot 10\% + 4.7\% \cdot 20\% = 5.98\%$
- **3.** $D_{\mathbf{1}} = \sum_i P_{i,\mathbf{1}} Q_i^{\text{new}} / I_{\mathbf{1}} = 1.6167 \cdot 10^{\mathbf{9}}$; new weights 75.2%, 18.9%, 5.9% (chart); the level stays 1 059.81
- **Interpretation**: the divisor absorbs basket changes that are not price moves
  - Without it the index would jump by the value of the added coins, a fake return

MFM_ch16_seminar

A2 [Solved]: The No-Arbitrage Band of a Pool with Fees — Task

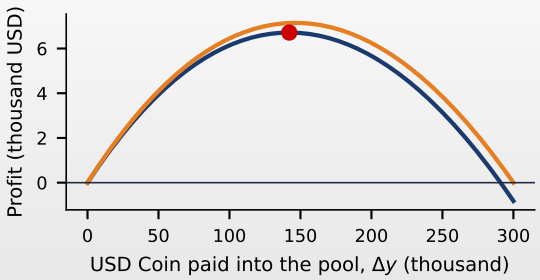
- **Questions**
  - when does it pay to trade against a pool?
  - how large is the optimal trade?
- **Data and inputs**
  - Constant-product pool with $x = 1\,000$ Ether and $y = 3\,000\,000$ USD Coin: $k = 3 \cdot 10^{9}$, pool price $p = 3\,000$
  - Fee $f = 0.3\%$ charged on the input and kept in the pool (Uniswap v2); outside price P; one trade, no transaction cost on the blockchain
- **Tasks**
  1. Derive the marginal cost of buying and the marginal revenue of selling Ether in the pool, and show that no trade pays while $\gamma P \leq p \leq P/\gamma$.
  2. For $P = 3\,300$, write the profit $\pi(\Delta y)$ of buying Ether with Δy USD Coin, solve $\pi'(\Delta y) = 0$ and compute Δy^*, the Ether bought and the new pool price.
  3. Compute the arbitrage profit with and without the fee.
- **Report**: the band, Δy^*, Δx, the new pool price and both profits
- **Interpretation**
  - where does arbitrage leave the pool price?
  - who pays the profit?
- Notebook: section A2

Back Exercise list

A2: Solution (1/2)

- **1.** Buying Δx Ether: $(x - \Delta x)(y + \gamma \Delta y) = k$ gives $\Delta y = \frac{1}{\gamma}\left(\frac{k}{x - \Delta x} - y\right)$; at $\Delta x = 0$ the marginal cost is $k/(\gamma x^2) = p/\gamma$
- Selling Ether: $(x + \gamma \Delta x)(y - \Delta y) = k$ gives the marginal revenue γp
- Buying pays iff $p/\gamma < P$, selling iff $\gamma p > P$; neither while $\gamma P \leq p \leq P/\gamma$, here $[3\,290.10, 3\,309.93]$ for $P = 3\,300$: the pool price $3\,000$ is below the band, so buying Ether pays
- For fixed external price P, the band is $\gamma P \leq p \leq P/\gamma$. For fixed pool price p, the same inequalities give $\gamma p \leq P \leq p/\gamma$: the variable being bounded changes.

A2: Solution (2/2)



- **2.** $\pi(\Delta y) = P\,[x - k/(y + \gamma\Delta y)] - \Delta y$; $\pi' = 0$ gives the effective reserve $y^e = y + \gamma\Delta y^* = \sqrt{\gamma P k} = 3\,141\,703$
- $\Delta y^* = (y^e - y)/\gamma = 142\,130$ USD Coin buys $\Delta x = 45.10$ Ether ($x' = k/y^e = 954.90$); the fee stays in the pool: $y' = y + \Delta y^*$, new price $y'/x' = 3\,290.55$
- **3.** Profit $P\Delta x - \Delta y^* = 6\,713$ USD; without the fee 7 147 USD (chart)
- **Interpretation**: arbitrage brings the effective pool price, $y^e/x' = \gamma P = 3\,290.10$, to the edge of the band, not to P; the retained fee makes the reserve ratio $y'/x' = 3\,290.55$
- The profit is paid by the liquidity providers, who sold Ether below the market price

MFM_ch16_seminar

B1: Calendars and Bootstrap Blocks

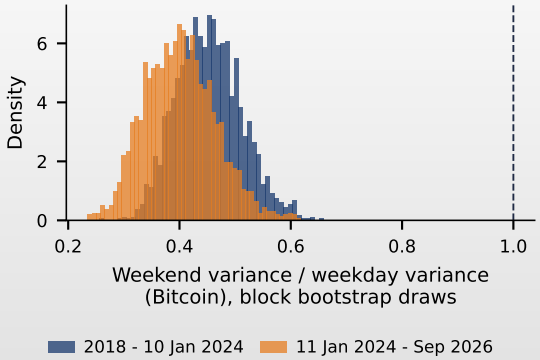
- Crypto returns use their seven-day calendar from January 2018; S&P 500 (Standard & Poor's) returns use its own trading calendar from 2 January 2018 to 18 September 2026. Annualisation uses 365 and 252 respectively.
- Use 2000 moving-block draws: 20 observations for volatility, seed 42. For the weekend/weekday variance ratio use 21 observations (three complete weeks), seed 42 before the ETF date and 43 after. Pair the two draw sequences for the change interval.
- The sampled rows retain their original weekday labels; truncating the last block can leave fewer than 21 observations. B7 repeats this protocol for Ether.

B1 [Solved]: Bitcoin's Volatility and the Weekend Effect — Task

- **Question**: is Bitcoin's weekend lull (lower variance on Saturdays and Sundays) a new effect of the spot ETFs?
- **Data and inputs**
  - Bitcoin daily log returns, 2018-01-02 – 18 Sep 2026, 3 182 days, 7 a week; S&P 500 on its own calendar
  - Split at the first spot ETF day, 11 January 2024: 2 200 days before, 982 after; moving-block bootstrap Künsch (1989), 2 000 resamples
- **Tasks**
  1. Compute the annualised volatility $s\sqrt{365}$ with a 95% block-bootstrap interval (blocks of 20 days) and its ratio to the S&P 500 volatility.
  2. Compute the ratio of weekend to weekday return variance before and after the ETFs, each with a block-bootstrap interval (blocks of 21 days).
  3. Compute a 95% interval for the change in the ratio from the difference of the two bootstrap distributions.
- **Report**: the volatility with its interval, both ratios with intervals, the change with its interval, the bootstrap histogram
- **Interpretation**
  - is the weekend lull new since the ETFs?
  - why resample blocks rather than single days?
- Notebook: section B1

Back Exercise list

B1: Solution



- **1.** Resample 2 000 series of 20-day blocks, recompute $s\sqrt{365}$, take the 2.5% and 97.5% percentiles: 64.1%, CI $[58.2, 70.4]$; S&P 500 19.2%, ratio 3.3
- **2.** Before: 0.45 $[0.35, 0.57]$; after: 0.39 $[0.29, 0.53]$ (chart: the two bootstrap distributions)
- **3.** Change -0.05, CI $[-0.22, 0.11]$: contains 0, so no significant change at 5%
- **Interpretation**: weekends carry less than half the weekday variance before and after 2024: the lull is not ETF-driven
  - Blocks keep the volatility clusters that single-day resampling destroys
- Caveat: the Hill tail-index estimate $\hat{\alpha}_H$ is near 2.7 (Chapter 16): the fourth moment Er^4 may be infinite
  - A variance bootstrap need not be consistent Athreya (1987): the intervals are indicative

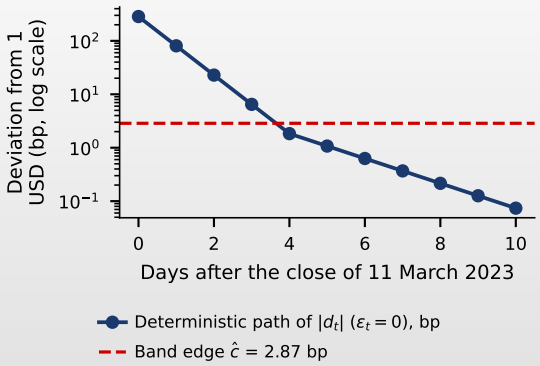
MFM_ch16_seminar

A3 [Solved]: Threshold AR Algebra — Task

- **Question**: how many days does USD Coin need to return inside its band after the March 2023 shock?
- **Data and inputs**
  - The threshold AR (threshold AR, What You Need for Today); these USD Coin estimates are supplied, so B5 need not be run first: $c = 2.8650$ bp, $\phi_{\text{in}} = 0.5855$, $\phi_{\text{out}} = 0.2832$
  - Deterministic path: all shocks set to zero, $\varepsilon_t = 0$; start at $d_0 = -285$ bp, the close of 11 March 2023
- **Tasks**
  1. Derive the half-life in each regime and compute it.
  2. Derive and compute the number of days until $|d_n| \leq c$.
  3. State what $\phi_{\text{in}} = 1$ would mean for the dynamics inside the band.
- **Report**: two half-lives in days, the number of days of task 2, one sentence on task 3
- **Interpretation**: what market mechanism could produce the reversion inside the band?
- Notebook: section A3

Back Exercise list

A3: Solution



- **1.** In one regime $d_h = \phi^h d_0$; $\phi^h = 0.5$ gives $h_{1/2} = \ln 0.5 / \ln \phi$: 1.29 days inside, 0.55 days outside
- **2.** $|d_0|\phi_{\text{out}}^n \le c \iff n \ge \ln(c/|d_0|)/\ln \phi_{\text{out}} = \ln(2.8650/285)/\ln 0.2832 = 3.65$: inside the band after 4 days, then the slower inside dynamics (chart)
- **3.** $\phi_{\text{in}} = 1$: a random walk inside the band (the band-TAR of Balke & Fomby (1997)), no pull towards par until the deviation leaves it
- **Interpretation**: $\phi_{\text{in}} < 1$ means reversion even inside the band
  - Market makers are one hypothesis, testable only with transaction-level data

MFM_ch16_seminar

B2–B4: Lag Choices for Inference

- Daily and weekly tracking regressions use four HAC lags; the Dimson regression and its same-sample OLS comparison use five. The log price-ratio trend uses 20 lags.
- The mean first difference of the log ratio and the sup-Wald statistic use $L = \lfloor 4(T/100)^{2/9} \rfloor$, with T the available regression observations. The weekly B3/B4 coefficient and interaction regressions use four lags.
- ADF uses AIC (Akaike Information Criterion) lag selection with the implementation default maximum; KPSS uses its automatic bandwidth. On the log ratio include a constant and trend; on its first difference ADF uses a constant. Record the selected lag/bandwidth in your output.

B2 [Solved]: IBIT: Tracking, Timing and the Fee Drift — Task

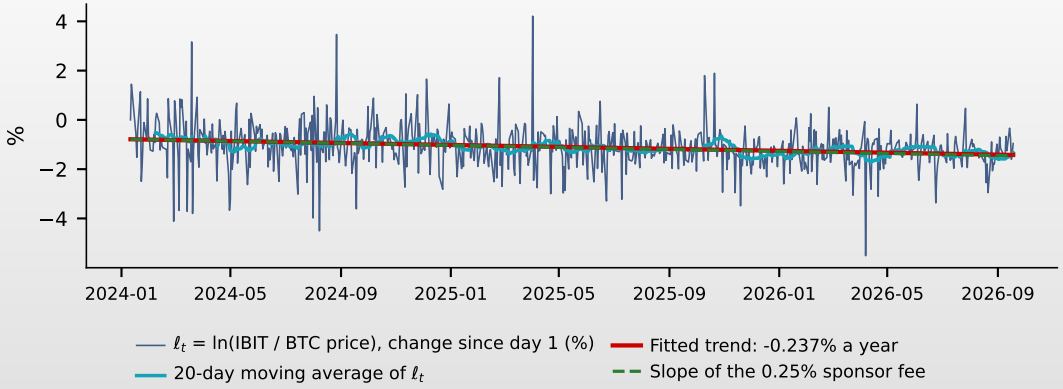
- **Questions**
  - does IBIT track Bitcoin one for one?
  - is its 0.25% fee visible in the prices?
- **Data and inputs**
  - IBIT (adjusted close) and Bitcoin on IBIT's trading days, 11 Jan 2024 – 18 Sep 2026: 673 days, 140 Friday-to-Friday weeks; prices aligned on common trading days first, then returns
  - $\ell_t = \ln(P_t^{\mathrm{IBIT}} / P_t^{\mathrm{BTC}})$; sponsor fee 0.25% a year, paid in Bitcoin BlackRock (2026); level 5%
- **Tasks**
  1. Estimate the daily and the weekly beta with Newey–West SE, and the Dimson sum beta $b_{-1} + b_0 + b_1$ with one lag and one lead of Bitcoin Dimson (1979).
  2. Run an ADF test Dickey & Fuller (1979) and a KPSS test Kwiatkowski et al. (1992), both with trend, on ℓ_t.
  3. Estimate the annual drift two ways, each with a 95% CI: the mean of $\Delta\ell_t$ (HAC) and the slope of ℓ_t on time in years (HAC, 20 lags).
- **Report**: three betas with SE, both test statistics with decisions, both drifts with CIs, the chart of ℓ_t
- **Interpretation**
  - why is the daily OLS beta below 1?
  - which drift estimator is precise?
- Notebook: section B2

Back Exercise list

B2: Solution (1/3)

- In the Dimson regression, the same-day coefficient 0.94 differs from the separate daily OLS beta 0.93. The lag/lead sum gives the Dimson exposure.
- **1.** Daily OLS $\hat{\beta} = 0.93$ (SE 0.022); weekly $\hat{\beta} = 1.007$ (SE 0.011), $R^2 = 0.98$
- Dimson: $\hat{b}_{-1} = 0.05$, $\hat{b}_0 = 0.94$, $\hat{b}_1 = 0.01$; sum 1.00 (SE 0.023), $t(\beta = 1) = -0.07$: not different from 1
- Why below 1: Bitcoin's daily close and the New York close are hours apart, part of each move reaches IBIT a day later; the lag term restores $\beta = 1$, as weekly data do
- **2.** ADF $= -11.07$, $p < 0.001$: reject a unit root; KPSS $= 0.041$, $p > 0.10$: do not reject trend stationarity. So ℓ_t = fee trend + stationary timing noise

B2: Solution (2/3)



MFM_ch16_seminar

B2: Solution (3/3)

- **3.** Mean of $\Delta\ell_t = \ell_t - \ell_{t-\mathbf{1}}$ per year: -0.36%, CI $[-9.28, 8.57]$: it equals $(\ell_T - \ell_{\mathbf{0}})/T$, because in the sum $\sum_t(\ell_t - \ell_{t-\mathbf{1}}) = \ell_T - \ell_{\mathbf{0}}$ the intermediate terms cancel ($\ell_{\mathbf{0}}$, ℓ_T: first and last values, T: sample length in years), so it uses two noisy days only
- Trend slope (red line): -0.237% a year, CI $[-0.315, -0.158]$: excludes 0, contains -0.25% (green line)
- Under a trend-stationary ℓ_t the trend slope with HAC SE is consistent and its CI valid; with a unit root the mean of differences would be the right tool (a pretest does not guarantee coverage)
- **Interpretation**: IBIT tracks Bitcoin one for one, and the 0.25% fee is visible as a slow downward trend in ℓ_t

MFM_ch16_seminar

B3 [Solved]: When Did Strategy's Bitcoin Beta Change? — Task

- **Questions**
  - did Strategy's exposure to Bitcoin change?
  - did the data choose the ETF launch as the date?
- **Data and inputs**
  - Weekly (Friday) log returns of Strategy (MSTR), Bitcoin and the S&P 500, 16 Apr 2021 – 18 Sep 2026, 284 weeks; prices aligned on common trading days first, then returns
  - $r_t^{\mathrm{MSTR}} = a + \beta_{\mathrm{BTC}} r_t^{\mathrm{BTC}} + \beta_{\mathrm{SPX}} r_t^{\mathrm{SPX}} + u_t$, Newey–West SE; level 5%; β_{BTC}, β_{SPX}: sensitivity of Strategy's return to Bitcoin and to the S&P 500; a: alpha, the return not explained by the two
- **Tasks**
  1. Estimate the model on the full period and on the two subperiods split at the ETF launch (to 5 Jan 2024; from 12 Jan 2024), and test $\beta_{\mathrm{BTC}} = 1$ in each.
  2. For every week in the middle 70% of the sample, compute the HAC Wald statistic $W(\tau)$ for a break in all three coefficients, and compare $\sup W$ with its 5% critical value Andrews (1993).
  3. Estimate the break date by least squares with its 95% CI Bai (1997), and the coefficients in each regime.
- **Report**: the table of task 1, $\sup W$ with the critical value and decision, the date with its CI, the chart of $W(\tau)$
- **Interpretation**: why is a split at an imposed date not a test of when the beta changed?
- Notebook: section B3

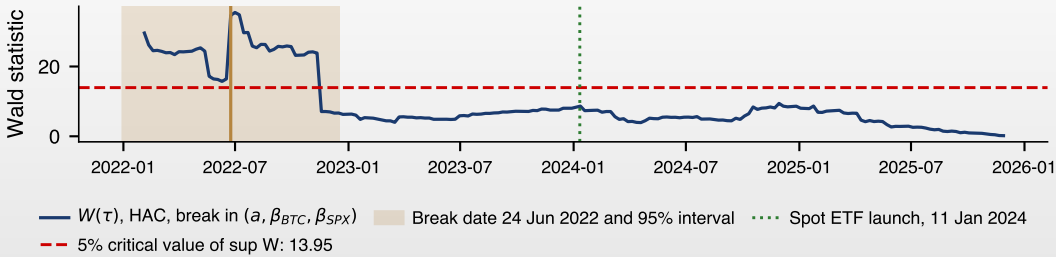
Back Exercise list

B3: Solution (1/2)

	n	$\hat{\beta}_{\mathrm{BTC}}$ (SE)	$t(\beta_{\mathrm{BTC}}=1)$	$\hat{\beta}_{\mathrm{SPX}}$ (SE)	alpha, % a year (t)
Apr 2021 – Sep 2026	284	1.18 (0.11)	1.64	1.02 (0.24)	−4.5 (−0.19)
Apr 2021 – 5 Jan 2024	143	1.01 (0.14)	0.08	1.40 (0.35)	−0.7 (−0.02)
12 Jan 2024 – Sep 2026	141	1.50 (0.12)	4.12	0.48 (0.29)	−9.3 (−0.29)

- **1.** $\hat{\beta}_{\mathrm{BTC}}$ rises from 1.01 to 1.50; $t = 4.12$ tests $\beta_{\mathrm{BTC}} = 1$ after 2024, not a change
- The change itself: a pooled regression with post-2024 interactions gives $\hat{\beta}_{\mathrm{post}} - \hat{\beta}_{\mathrm{pre}} = 0.49$ (HAC SE 0.18, $t = 2.69$); alpha is never significant

B3: Solution (2/2)



- **2.** sup $W = 35.49$ against a 5% critical value of 13.95, $p < 0.001$: reject stability at 5%
- **3.** Break 24 Jun 2022, 95% CI with regime-specific moments 1 Oct 2021 – 30 Sep 2022; W at the ETF launch only 8.64
- Regimes: Bitcoin beta 0.69 → 1.48, S&P 500 beta 1.25 → 1.00
- **Interpretation**: the data choose mid-2022, a period of crypto-wide stress, not the ETF launch
- A split at an imposed date finds a change whenever the true break lies before that date, so it cannot date the change

MFM_ch16_seminar

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

- **Prompt** sent to an AI assistant: "Compare Bitcoin with lending USD Coin (USDC) at 5% a year: compute Bitcoin's annualised Sharpe ratio on daily data for 2021–2025, with the USDC yield as the risk-free rate"
- **Answer**: the code below and the conclusion

```
import numpy as np, pandas as pd
btc = pd.read_csv('data/market/BTC-USD.CC.csv', index_col='date',
                  parse_dates=True)['close']
r = np.log(btc.loc['2021':'2025']).diff().dropna()   # daily log returns
mu = r.mean() * 252                                  # annual mean
sd = r.std() * np.sqrt(252)                          # annual volatility
rf = 0.05                                            # USDC lending: risk-free
sharpe = (mu - rf) / sd
print(f"mu = {mu:.1%}, vol = {sd:.1%}, Sharpe = {sharpe:.2f}")
```

- **Conclusion of the AI**: "mu = 15.1%, vol = 48.5%, Sharpe = 0.21. USDC is fully backed by dollars and has never traded below $0.99, so lending it is a risk-free investment"

Back Exercise list

C1 [Proposed]: Critique an AI Answer — Task

- **Model**: B1 (annualising crypto); the USD Coin daily close and low (stablecoins, What You Need for Today)
- **Data**: Bitcoin (daily close), 2021–2025, including the 31 December 2020 close for the first return; USD Coin (daily close and low); 3-month Treasury bill (use the period mean discount yield divided by 100 as the annual-rate approximation), FRED DTB3
- **Tasks**
  1. Find the five substantive errors, three in the code and two in the text. Count a repeated root cause once.
  2. For each, say how you would detect it: a count in the data, the data themselves, or a definition from What You Need for Today.
  3. Write the corrected code and report the annualised mean, volatility and Sharpe ratio.
  4. Log every error you found in the format of `AI_ERRORS.md`: request, answer, error, how found, correction.
- **Interpretation**
  - how many daily returns should 2021–2025 contain for Bitcoin?
  - if lending USDC (USD Coin, stablecoin issued by Circle) were risk-free, why would it pay more than a Treasury bill?
- AI code usually runs without an error message: the mistakes are in the finance, not in the syntax

B4 [Proposed]: Coinbase — Task

- **Questions**
  - is Coinbase a Bitcoin proxy or a US growth stock?
  - when did its Bitcoin beta change?
- **Model**: B3
- **Data and inputs**
  - Weekly log returns of Coinbase (listed April 2021), Bitcoin and the S&P 500, 16 Apr 2021 – 18 Sep 2026, 283 weeks
  - Same two-factor model, HAC settings, trimming and level as B3
- **Tasks**
  1. Estimate the model on the full period and the two subperiods, and test $\beta_{\mathrm{BTC}} = 1$ in each.
  2. Compute $W(\tau)$ for every candidate week, plot it, and compare $\sup W$ with the 5% critical value.
  3. Estimate the break date with its 95% CI and the Bitcoin beta in each regime.
- **Report**: the table of task 1, the chart of $W(\tau)$, the test decision, the date with its CI
- **Interpretation**
  - do Strategy and Coinbase break in the same week?
  - what common event could explain both breaks?
- Notebook: section B4

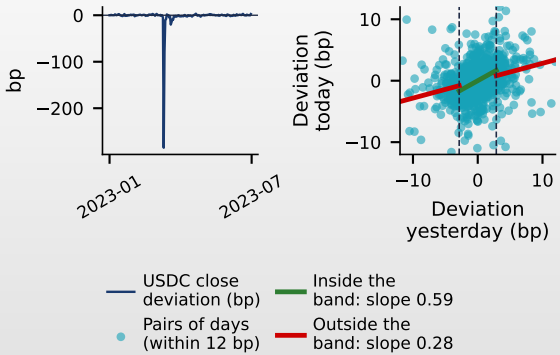
Back Exercise list

B5 [Solved]: Is the USD Coin Peg a Threshold Process? — Task

- **Question**: does USD Coin revert to 1 USD faster outside an arbitrage band than inside it?
- **Data and inputs**
  - USD Coin daily close, 1 Jan 2021 – 18 Sep 2026, 2 087 days; $d_t = 10^4(P_t - 1)$ in bp
  - Threshold AR without intercept Balke & Fomby (1997) (threshold AR, What You Need for Today): slope ϕ_{in} if $|d_{t-1}| \le c$, ϕ_{out} otherwise; level 5%
- **Tasks**
  1. Compute the standard deviation of d_t, the days with $|d_t| > 50$ bp, the lowest close and the lowest daily low.
  2. Estimate c by least squares over the values of $|d_{t-1}|$ that leave at least 15% of the days in each regime, and ϕ_{in}, ϕ_{out} with White standard errors.
  3. Test H_0: $\phi_{\text{in}} = \phi_{\text{out}}$ with $\sup_c W(c)$ and the fixed-regressor wild bootstrap of Hansen (1996) (999 draws); also report the χ^2_1 p-value.
  4. Repeat steps 2–3 without the pairs (d_{t-1}, d_t) that touch 9–16 March 2023.
- **Report**: $\hat{c}$, both slopes with SE and half-lives, $\sup W$, both p-values, with and without March 2023
- **Interpretation**
  - why is χ^2_1 wrong here?
  - does one stress week carry the evidence?
- Notebook: section B5

Back Exercise list

B5: Solution (1/2)



- Left: close deviation in early 2023; right: today against yesterday within ± 12 bp, slopes inside and outside $\pm \hat{c}$
- **1.** s.d. 7.5 bp (3.7 without March 2023)
  - 3 days beyond 50 bp; lowest close -285 bp, lowest low $-1\,226$ bp, both on 11 March 2023
- **2.** For each c in $[0.29, 2.93]$ fit OLS in each regime, keep the smallest total sum of squared residuals (SSR): $\hat{c} = 2.87$ bp (15.7% of days outside)
  - $\hat{\phi}_{\text{in}} = 0.59$ (SE 0.04), $\hat{\phi}_{\text{out}} = 0.28$ (SE 0.03)
  - Half-lives 1.29 and 0.55 days

MFM_ch16_seminar

B5: Solution (2/2)

- **3.** $W(c) = (\hat{\phi}_{\text{in}} - \hat{\phi}_{\text{out}})^2/(\hat{v}_{\text{in}} + \hat{v}_{\text{out}})$ for each c, with $\hat{v}_{\text{in}}$, $\hat{v}_{\text{out}}$ the squared White standard errors of the two slopes; $\sup W = 32.5$
- Bootstrap: $d_t^* = \hat{e}_t \eta_t$ ($\hat{e}_t$: residuals of the linear AR(1), $\eta_t \sim N(0,1)$), regress on the same fixed d_{t-1}, recompute $\sup W^*$ 999 times: $p = 0.001$ (χ_1^2: < 0.001): reject linearity at 5%
  - Multiplying each residual by η_t keeps its own variance: the sample excess kurtosis of d_t is 1 031 (a descriptor of extreme days, not an estimate), calm and stress days must not be mixed
- **4.** Without 9–16 March 2023: $\sup W = 5.5$, bootstrap $p = 0.291$: linearity no longer rejected; χ_1^2 would give 0.019, a rejection based on an inadequate null distribution
- Why χ_1^2 fails: c is not identified under H_0, and the supremum over many c exceeds any single $W(c)$ (the Davies problem)
- **Interpretation**: the evidence for a band rests on one stress week; without it the data do not reject a linear, fast-reverting peg (non-rejection, not proof of linearity)

MFM_ch16_seminar

B5: Serial Dependence of the Scores

- The bootstrap with independent η_t assumes serially uncorrelated scores $d_{t-1}\hat{e}_t$
  - their lag-1 autocorrelation: -0.23 (full sample), -0.28 (without 9–16 March 2023)
- Multipliers held constant on blocks of 13 consecutive days (dependent wild bootstrap, Shao (2010))
  - full sample: $p = 0.001$; without the stress week: $p = 0.191$
- Lagged differences Δd_{t-j} added until Ljung–Box (10 lags) no longer rejects at 5%, threshold re-estimated
  - full sample: 2 lags, $p = 0.001$; without the stress week: 9 lags, $p = 0.373$
- **Interpretation**: the full-sample rejection survives both corrections; without the stress week, no calibration rejects linearity

Extension Packs: What Each One Needs

Pack	Exercises	Prerequisites
1. Index divisor	A4 [Proposed]	A1
2. Weighted pools	A5 [Proposed], A6 [Solved], A7 [Proposed]	A2; Chapter 16: LVR
3. Stablecoin pegs	A8 [Proposed], B6 [Proposed]	B5, A3
4. Crypto volatility	B7 [Proposed]	B1
5. Spot ETFs	B8 [Proposed]	B2
6. Hedge or risk asset?	C2 [Proposed]	B1, B2; Forbes–Rigobon (What You Need for Today)
7. Crypto factors after 2020	C3 [Proposed], project seed	B2, B3; the Chapter 16 case study

- Each exercise has its notebook section; solutions of [Proposed] exercises are discussed in class

A4 [Proposed]: Replacing a Constituent — Task

- **Question**: how must the divisor change when a coin leaves the index and another enters?
- **Model**: A1
- **Data and inputs**
  - The index of A1 after the reallocation: level $I = 1\,059.81$, divisor $1.6167 \cdot 10^9$, month-end prices of A1
  - On the same day coin C is delisted and coin D enters with $P_D = 0.5$ USD and $Q_D = 40$ billion coins
- **Tasks**
  1. Compute the market value of the old and of the new basket on the switch day.
  2. Derive the divisor that keeps I continuous and compute it.
  3. Compute the level the day after if the divisor were not changed and no price moved.
- **Report**: both basket values, the new divisor, the level of task 3
- **Interpretation**
  - why does the divisor fall here?
  - who would be misled by the fake loss?

Back Exercise list

A5 [Proposed]: Arbitrage in a Weighted Pool — Task

- **Question**: does a pool that holds mostly Ether lose less to arbitrageurs than a 50/50 pool?
- **Model**: A2
- **Data and inputs**
  - Weighted pool (Balancer type) $x^w y^{1-w} = k$, w the value weight of Ether ($w = 1/2$: constant product), here $w = 0.8$, $x = 1\,000$ Ether, pool price $p = 3\,000$; fee 0.3% on the input, kept in the pool
  - The outside price rises to $P = 3\,300$; one arbitrage trade
- **Tasks**
  1. Differentiate the invariant to show that the pool price is $p = \frac{w}{1-w}\frac{y}{x}$, and compute y.
  2. Show that the no-arbitrage band $\gamma P \le p \le P/\gamma$ does not depend on w.
  3. Derive the reserves after the optimal trade, compute Δy^*, the profit, and the profit per unit of pre-trade pool value $px + y$; compare with A2 (6 million USD pool).
- **Report**: y, x', Δy^*, the profit, the profit per unit of value and its ratio to A2
- **Interpretation**: why does a pool with $w = 0.8$ lose less to arbitrageurs?

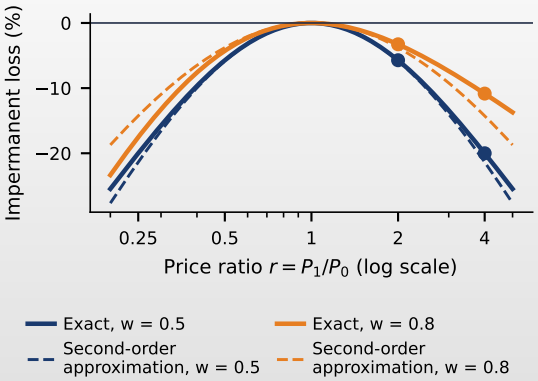
Back Exercise list

A6 [Solved]: Impermanent Loss of a Weighted Pool — Task

- **Question**: how does the impermanent loss depend on the weight w of Ether in the pool?
- **Data and inputs**
  - A liquidity provider owns a pool $x^w y^{\mathbf{1}-w} = k$ with pool price $p = \frac{w}{\mathbf{1}-w}\frac{y}{x}$; arbitrage keeps p at the market price P, no fees
  - The price moves from $P_{\mathbf{0}}$ to $P_{\mathbf{1}} = rP_{\mathbf{0}}$; the benchmark is holding the initial tokens
- **Tasks**
  1. Show that the value shares stay w and $1-w$ and that the pool value is $V(P) = cP^w$, with c a constant that depends on k and w.
  2. Derive the impermanent loss $\mathrm{IL}_w(r)$, pool value over holding value, minus 1: $\mathrm{IL}_w(r) = r^w/(wr + 1 - w) - 1$; check that $w = 1/2$ gives $2\sqrt{r}/(1+r) - 1$.
  3. Expand $\ln(1 + \mathrm{IL}_w)$ to second order in $u = \ln r$ and compute IL_w for $r = 2, 4$ and $w = 0.5, 0.8$.
- **Report**: the formula, the approximation and four values
- **Interpretation**
  - which pool loses less when the price doubles?
  - why does it lose less?

Back Exercise list

A6: Solution



- **1.** $p = P$ gives $Px = wV$, $y = (1-w)V$ with $V = Px + y$; in the invariant: $V = k\, w^{-w}(1-w)^{w-1} P^w = cP^w$
- **2.** Pool: $V_1/V_0 = r^w$; holding: $wr + 1 - w$; ratio minus 1 is IL_w; at $w = 1/2$: $2\sqrt{r}/(1+r) - 1$
- **3.** $\ln(1 + \mathrm{IL}_w) \approx -\frac{1}{2} w(1-w) u^2$ (dashed); $r = 2$: -5.72% ($w = 0.5$), -3.27% ($w = 0.8$); $r = 4$: -20.00%, -10.84%
- **Interpretation**: the loss scales with $w(1-w)$, the curvature of $V(P)$: a mostly-Ether pool rebalances less and sells less volatility

MFM_ch16_seminar

A7 [Proposed]: Loss Versus Rebalancing by Itô — Task

- **Question**: at what rate does a weighted pool lose against the rebalancing portfolio?
- **Model**: A6
- **Data and inputs**
  - Price $dP_t = \mu P_t\,dt + \sigma P_t\,dW_t$ (geometric Brownian motion: drift μ, volatility σ, W_t a standard Brownian motion); pool value $V(P) = cP^w$ (A6); frictionless arbitrage, no fees Milionis et al. (2022)
  - The rebalancing portfolio holds the pool's Ether $V'(P_t)$ at every instant and trades at market prices: $dR_t = V'(P_t)\,dP_t$, with R_t its value and V', V'' the first and second derivatives of V
- **Tasks**
  1. Compute V' and V'', apply Itô's lemma to $V(P_t)$ and show $d(R_t - V_t) = \frac{1}{2}w(1-w)\sigma^2 V_t\,dt$.
  2. Check that $w = 1/2$ gives the constant-product rate $\sigma^2/8$ Milionis et al. (2022).
  3. Compute the annual rate for $\sigma = 70\%$ and $w = 0.5,\ 0.8,\ 0.95$.
- **Report**: the derivation and three rates in % a year
- **Interpretation**: why does the drift μ not appear in the loss?

Back Exercise list

A8 [Proposed]: What Does a Linear AR(1) Estimate? — Task

- **Question**: what does a linear AR(1) fitted to a threshold process measure?
- **Model**: A3
- **Data and inputs**
  - The USD Coin deviations of B5; linear AR(1) $d_t = \phi d_{t-1} + e_t$ (e_t: error) by OLS without intercept
  - $\hat{\phi}_{\text{in}}$, $\hat{\phi}_{\text{out}}$: the regime-wise OLS slopes at $\hat{c}$ of B5
- **Tasks**
  1. Show that, in sample, $\hat{\phi}_{\text{lin}} = (1-s)\hat{\phi}_{\text{in}} + s\,\hat{\phi}_{\text{out}}$ with $s = \sum_{\text{out}} d_{t-1}^2 / \sum_t d_{t-1}^2$, the share of $\sum d_{t-1}^2$ from days outside the band.
  2. Compute s for USD Coin and check the identity numerically.
  3. Compute the linear half-life and say which regime it describes.
- **Report**: s, the weighted average, $\hat{\phi}_{\text{lin}}$, the linear half-life
- **Interpretation**: why can a linear AR(1) report fast reversion although most days revert slowly?

Back Exercise list

B6 [Proposed]: Tether and Dai — Task

- **Questions**
  - are the Tether and Dai pegs threshold processes?
  - which band is widest?
- **Model**: B5
- **Data and inputs**
  - Daily closes of Tether (USDT) and Dai (DAI), 1 Jan 2021 – 18 Sep 2026; deviations in bp
  - Same threshold AR, grid, bootstrap (999 draws) and level as B5
- **Tasks**
  1. Estimate $\hat{c}$, $\hat{\phi}_{\text{in}}$ and $\hat{\phi}_{\text{out}}$ with White SE for each coin.
  2. Test linearity with $\sup W$ and the bootstrap, with and without 9–16 March 2023.
  3. Plot both slopes with 95% intervals for Tether, USD Coin and Dai.
- **Report**: one table with the three coins and the chart
- **Interpretation**
  - which coin reverts fastest outside its band?
  - what would a test of two thresholds need?

Back Exercise list

B7 [Proposed]: Ether and Solana — Task

- **Questions**
  - are Ether and Solana riskier than Bitcoin relative to shares?
  - do they share its weekend lull?
- **Model**: B1
- **Data and inputs**
  - Daily log returns of Ether (2018–2026) and Solana (April 2020 – 2026); S&P 500 over the same periods
  - Same split date as B1; volatility blocks 20, seed 42; weekend-ratio blocks 21, seeds 42 before and 43 after; 2000 draws
- **Tasks**
  1. Compute each coin's annualised volatility with a block-bootstrap CI and its ratio to the S&P 500.
  2. Compute the weekend / weekday variance ratio before and after the ETFs, and a CI for the change.
- **Report**: one table with Ether, Solana and Bitcoin (B1) and a chart of the intervals
- **Interpretation**: which coin is the riskiest relative to shares?

Back Exercise list

B8 [Proposed]: ETHA and Ether — Task

- **Question**: does ETHA track Ether as well as IBIT tracks Bitcoin?
- **Model**: B2
- **Data and inputs**
  - ETHA (adjusted close) and Ether on ETHA's trading days, 23 Jul 2024 – 18 Sep 2026: 541 days, 112 weeks
  - Sponsor fee 0.25% a year; same tests and settings as B2
- **Tasks**
  1. Estimate the daily, weekly and Dimson betas.
  2. Run the ADF and KPSS tests on the log price ratio.
  3. Estimate both drifts with 95% CIs and plot the log price ratio with its trend.
- **Report**: betas with SE, test decisions, both drifts with CIs, the chart
- **Interpretation**: is ETHA a worse tracker than IBIT, or only noisier?

Back Exercise list

What Does a Tail-Regime Slope Mean?

- In C2 the three tail indicators are nested. A day below the 1% quantile is also below the 5% and 10% quantiles: all three extra coefficients are active.
- Each coefficient is an increment. For example, with baseline 1 and successive increments -0.2, -0.3, -0.1:

Market-return region	Active increments	Total slope
Above the 10% quantile	0	1
Between the 5% and 10% quantiles	1	0.8
Between the 1% and 5% quantiles	2	0.5
Below the 1% quantile	3	0.4

- For a confidence interval of that sum, use the full HAC covariance matrix, including the covariances between coefficients. Adding individual standard errors does not give the standard error of the sum.
- A quantile-regression slope answers a different question: it concerns a quantile of Bitcoin conditional on the market return, rather than conditioning on an extreme market day.

C2: Inference and the Hedge Verdict

- Use January 2021–10 January 2024 as the Forbes–Rigobon base. The variance change is post divided by pre minus one, including negative values. Recompute both variances and the adjusted correlation in each 20-day pairs-block bootstrap replication.
- For regime slopes use four HAC lags and the full coefficient covariance matrix. For quantile-regression intervals, resample paired return vectors in 20-day blocks and refit each quantile.
- A significantly positive all-day slope is evidence against a hedge; a significantly positive crisis-regime slope is evidence against a safe haven there. An interval spanning zero is inconclusive, not proof of protection.
- A regime slope of -0.2 with SE 0.2 gives an approximate interval $[-0.592, 0.192]$: both negative and positive relations remain plausible.

C2 [Proposed]: Alternative Asset, Risk Asset or Hedge? — Task

- **Questions**
  - is Bitcoin a hedge or a safe haven for equity investors?
  - did the ETFs change that?
- **Model**: B1, B2
- **Data and inputs**
  - Daily log returns of Bitcoin, S&P 500, QQQ (Invesco QQQ Trust, on the Nasdaq-100 index), GLD (SPDR Gold Shares) and TLT (iShares 20+ Year Treasury Bond ETF) on common days, Jan 2021 – 10 Jan 2024 against 11 Jan 2024 – 18 Sep 2026
  - Hedge: uncorrelated or negatively correlated on average; safe haven: on the worst market days Baur & Lucey (2010); earlier evidence Bouri et al. (2017)
- **Tasks**
  1. Compute the correlations with Bitcoin in both periods, with block-bootstrap CIs for each and for the change.
  2. Test the change of the S&P 500 correlation with the Forbes–Rigobon adjusted correlation Forbes & Rigobon (2002).
  3. Estimate $r_t^{\mathrm{BTC}} = a + b\, r_t^{\mathrm{S\&P}} + \sum_q c_q r_t^{\mathrm{S\&P}} 1\{r_t^{\mathrm{S\&P}} \le q\} + u_t$ (q: the 10%, 5%, 1% quantiles of $r_t^{\mathrm{S\&P}}$; $1\{\cdot\}$: 1 if the condition holds, else 0; c_q: the extra slope on days below q; HAC SE) and quantile-regression betas $\beta(\tau)$, the slope of the τ-quantile of r_t^{BTC}, at $\tau = 1\%, 5\%, 50\%$ Koenker & Bassett (1978).
- **Report**: a correlation table, the Forbes–Rigobon test, the regime slopes with intervals, a verdict
- **Interpretation**: is Bitcoin "alternative" in its distribution but "risk" in its co-movement?

Back Exercise list

C3 [Proposed]: Crypto Factors After July 2020 — Context and Data

- **Questions**
  - out of sample, does three-week momentum among large coins survive?
  - is the crypto market still segmented from equities?
- **Model**: B2 (HAC inference), B3 (break at an unknown date); a project seed from the Chapter 16 case study
- **Context**: Liu, Tsyvinski & Wu (2022), building on Liu & Tsyvinski (2021), price 1,827 coins with CMKT (Crypto Market factor, value-weighted market excess return), CSMB (Crypto Size factor, small minus big coins) and CMOM (Crypto Momentum factor, past winners minus losers), over 2014 – July 2020; three-week momentum earns 3.1% a week (Table 4)
- **Data**
  - Daily closes of Bitcoin, Ether, XRP (native token of the XRP Ledger), BNB (native token of BNB Chain, formerly Binance Coin), Cardano, Solana, Dogecoin, Litecoin, Chainlink; the five-coin market-value index of Chapter 16 as the CMKT proxy (from February 2018)
  - One-month bill rate and global five factors Fama & French (2015) from the Kenneth French Data Library; the Canadian, Singapore, Australian dollar, euro and pound against the USD from FRED
  - Weekly simple returns on the paper's calendar (week 52 ends on 31 December); in sample to July 2020, out of sample from August 2020

Back Exercise list

C3: Currency Quotes and Eligible Coins

- FRED exchange rates, for a USD investor; simple returns, compounded to the weekly calendar
  - DEXCAUS and DEXSIUS: Canadian dollars (CAD) and Singapore dollars (SGD) per USD; invert their price ratios
  - DEXUSAL, DEXUSEU and DEXUSUK: USD per Australian dollar (AUD), per euro (EUR) and per pound sterling (GBP)
- Factors SMB (Small Minus Big: size), HML (High Minus Low: value), RMW (Robust Minus Weak: profitability) and CMA (Conservative Minus Aggressive: investment)
  - $\prod_t (1 + r_{L,t} - r_{S,t}) - 1$: return of a strategy rebalanced daily, long (L) and short (S) legs equal to the capital at the start of each day
  - Not the difference between the weekly returns of the two portfolios
- The CMKT proxy starts in February 2018, so its in-sample factor regression covers only the common weeks from then through July 2020; descriptive coin statistics use the longer coin samples of the task (different samples)
- The replication ranks coins with three observed past weekly returns and an observed next-week return, keeping weeks with at least six coins; conditioning on future data and surviving coins makes it a descriptive illustration, not a proven investable strategy
- Report sample kurtosis only as a descriptor: the fourth moment may be infinite; use the Hill analysis to discuss that limitation

C3 [Proposed]: Crypto Factors After July 2020 — Task

- **Tasks**
  1. Build weekly returns for Bitcoin, XRP (from 2014) and Ether (from week 33 of 2015) to July 2020, and compare mean, median, standard deviation, skewness and kurtosis with Table 1, Panel B.
  2. Regress the weekly excess return of the CMKT proxy on the five global factors and the five currencies (footnote 16 of Liu, Tsyvinski & Wu (2022)), in and out of sample, with HAC t-statistics and R^2.
  3. Each week, rank the coins by their past three-week return and hold the top three minus the bottom three for one week (a descriptive illustration, not the paper's CMOM); report the mean and the crypto CAPM (Capital Asset Pricing Model) α, β, R^2, in and out of sample.
  4. Test a change in alpha after July 2020 with a pooled HAC regression, then a break in (α, β) at an unknown date with sup-Wald, as in B3.
- **Report**: the comparison with the paper, the in- and out-of-sample tables, the tests of task 4, the full list of tests run
- **Interpretation**: if the momentum alpha of these large coins stays significant, is it a premium or survivorship bias (the chosen coins survived to 2026)?

Acronyms used in this material

- **ADF**: Augmented Dickey–Fuller (test)
- **AI**: Artificial Intelligence
- **AIC**: Akaike Information Criterion
- **AR**: AutoRegressive (model)
- **AUD**: Australian Dollar (ISO 4217 code)
- **BNB**: BNB (native token of BNB Chain, formerly Binance Coin)
- **CAD**: Canadian Dollar (ISO 4217 code)
- **CAPM**: Capital Asset Pricing Model
- **CI**: Confidence Interval
- **CMKT**: Crypto Market factor (value-weighted market excess return)
- **CMOM**: Crypto Momentum factor (past winners minus losers)
- **CRIX**: CRypto IndeX
- **CSMB**: Crypto Size factor (small minus big coins)
- **DAI**: Dai (stablecoin of the Maker/Sky protocol)
- **DTB3**: 3-Month Treasury Bill: secondary market rate (FRED series code)
- **ETF**: Exchange-Traded Fund
- **ETHA**: iShares Ethereum Trust ETF (ticker)
- **FRED**: Federal Reserve Economic Data
- **HAC**: Heteroskedasticity and Autocorrelation Consistent (standard errors)
- **KPSS**: Kwiatkowski–Phillips–Schmidt–Shin test
- **LVR**: Loss-Versus-Rebalancing
- **MSTR**: Strategy Inc., formerly MicroStrategy (Nasdaq ticker)
- **OLS**: Ordinary Least Squares
- **SE**: Standard Error
- **SSR**: Sum of Squared Residuals
- **TAR**: Threshold AutoRegression
- **US**: United States
- **USDC**: USD Coin (stablecoin issued by Circle)
- **XRP**: XRP (the native token of the XRP Ledger)

References (1/2)

- Adams, H., Zinsmeister, N., Robinson, D. (2020). Uniswap v2 Core. Technical white paper.
- Andrews, D.W.K. (1993). Tests for parameter instability and structural change with unknown change point. *Econometrica*, 61(4), 821–856.
- Athreya, K.B. (1987). Bootstrap of the mean in the infinite variance case. *Annals of Statistics*, 15(2), 724–731.
- Bai, J. (1997). Estimation of a change point in multiple regression models. *Review of Economics and Statistics*, 79(4), 551–563.
- Balke, N.S., Fomby, T.B. (1997). Threshold cointegration. *International Economic Review*, 38(3), 627–645.
- Baur, D.G., Lucey, B.M. (2010). Is gold a hedge or a safe haven? An analysis of stocks, bonds and gold. *Financial Review*, 45(2), 217–229.
- BlackRock (2026). iShares Bitcoin Trust ETF (IBIT). Product page.
- Bouri, E., Molnár, P., Azzi, G., Roubaud, D., Hagfors, L.I. (2017). On the hedge and safe haven properties of Bitcoin: is it really more than a diversifier?. *Finance Research Letters*, 20, 192–198.
- Dickey, D.A., Fuller, W.A. (1979). Distribution of the estimators for autoregressive time series with a unit root. *Journal of the American Statistical Association*, 74(366), 427–431.
- 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. (2015). A five-factor asset pricing model. *Journal of Financial Economics*, 116(1), 1–22.

References (2/2)

- Forbes, K.J., Rigobon, R. (2002). No contagion, only interdependence: measuring stock market comovements. *Journal of Finance*, 57(5), 2223–2261.
- Hansen, B.E. (1996). Inference when a nuisance parameter is not identified under the null hypothesis. *Econometrica*, 64(2), 413–430.
- Koenker, R., Bassett, G. (1978). Regression quantiles. *Econometrica*, 46(1), 33–50.
- Künsch, H.R. (1989). The jackknife and the bootstrap for general stationary observations. *Annals of Statistics*, 17(3), 1217–1241.
- Kwiatkowski, D., Phillips, P.C.B., Schmidt, P., Shin, Y. (1992). Testing the null hypothesis of stationarity against the alternative of a unit root. *Journal of Econometrics*, 54(1–3), 159–178.
- Liu, Y., Tsyvinski, A. (2021). Risks and returns of cryptocurrency. *Review of Financial Studies*, 34(6), 2689–2727.
- Liu, Y., Tsyvinski, A., Wu, X. (2022). Common risk factors in cryptocurrency. *Journal of Finance*, 77(2), 1133–1177.
- Milionis, J., Moallemi, C.C., Roughgarden, T., Zhang, A.L. (2022). Automated market making and loss-versus-rebalancing. arXiv:2208.06046.
- Newey, W.K., West, K.D. (1987). A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix. *Econometrica*, 55(3), 703–708.
- Pele, D.T., Wesselhöfft, N., Härdle, W.K., Kolossiatis, M., Yatracos, Y.G. (2023). Are cryptos becoming alternative assets?. *European Journal of Finance*, 29(10), 1064–1105.
- Trimborn, S., Härdle, W.K. (2018). CRIX an index for cryptocurrencies. *Journal of Empirical Finance*, 49, 107–122.