








Modelling Financial Markets

Seminar 6: Multivariate Volatility and Dependence — 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

When Do Two Markets Move Together?

- Question of the seminar
  - **when do two markets move together, and how sure can we be?**
- Deliverables
  - the answers to the A exercises
  - one table and one chart per B exercise
  - the corrected lines of the artificial intelligence (AI) answer in C1
  - one interpretation sentence per exercise
- Assessment criteria
  - correct derivations in the A exercises
  - estimates reported with their uncertainty: standard errors, intervals or bootstrap p-values
  - one interpretation per result
  - each error in the AI answer of C1 explained, not only corrected
- Notebook for today: chapter6_seminar_notebook.ipynb (Colab)
- Seminar: Prof. Daniel Traian Pele, danpele@ase.ro; Antoaneta Amza, PhD candidate, antoaneta.amza@csie.ase.ro

Which Assets and Horizons Do We Compare?

- An Exchange-Traded Fund (ETF) is a fund traded on an exchange
- SPY (SPDR S&P 500 ETF Trust) tracks the S&P 500 (Standard & Poor's) stock index; TLT (iShares 20+ Year Treasury Bond ETF) holds long-term US (United States) Treasury bonds
- BET: Bucharest Exchange Trading index; Euro Stoxx 50 and S&P 500 close at different times
- VIX: the Cboe volatility index, inferred from option quotes; in C3, VIX is compared with global common volatility (COVOL)
- Prices are aligned on common dates before computing log returns
- Weekly observations use Friday closes; a two-day return spans two trading steps

Data Used

Exercise	Series, price used	Frequency, sample	Rows
B1, A5	SPY and TLT, adjusted closes	daily, 31 July 2002 – 18 September 2026	6 073
A2, B2, B6	S&P 500 and Euro Stoxx 50, closes	weekly (Friday), 14 January 2000 – 18 September 2026	1 393
B4	BET and Euro Stoxx 50, closes	daily, 5 January 2010 – 18 September 2026	4 154
B3, A1 ext.	S&P 500, Euro Stoxx 50 and BET, closes	2-day, 5 January 2006 – 30 December 2021	3 880
B5	Bitcoin (USD) and S&P 500, closes	daily, 18 September 2014 – 18 September 2026	3 018
C2	BET and Euro Stoxx 50, closes	weekly, 15 January 2010 – 18 September 2026	871

- Rows: return observations after aligning the prices on common trading days (the rule for every joint analysis); log returns; weekly returns from the last common close of each week
- Daily prices of the course data, one series per asset: ETFs use the adjusted close (adjusted for splits and distributions), indices and Bitcoin the close; C3 uses 20 US-listed ETFs and the VIX index (see the C3 slide)

Setup: Run and Verify

```
# notebook section "Data" defines
# joint_returns(names, start) and
# weekly_returns(names, start):
# prices aligned on common days, then log returns
R = joint_returns(['spy', 'tlt'])
print(len(R), R.index[0].date(),
      R.index[-1].date())
print((100 * R).head(3).round(3))
W = weekly_returns(['sp500', 'stoxx'],
                   start='2000-01-01')
print(len(W))
```

- Expected: 6 073 rows from 31 July 2002 to 18 September 2026, then

date	SPY (%)	TLT (%)
31 July 2002	0.242	1.231
1 August 2002	−2.646	0.568
2 August 2002	−2.267	1.019

- and 1 393 weekly pairs
- If a count differs: check the series, the price used (close or adjusted close) and that prices were aligned before computing returns

Exercise Map

Exercise	Question	Type	Worked example to follow
A1	how much of a crisis correlation is volatility?	Solved	–
A2	which copula parameter matches Kendall's τ?	Solved	–
A3	do extremes arrive together in the t copula?	Solved	–
B1	how does the SPY–TLT correlation move?	Solved	–
B2	Gaussian or t copula for two equity indices?	Solved	–
B3	contagion in 2008 and 2020? (tasks 3–4 after the lecture)	Proposed	A1
C1	what is wrong in an AI answer?	Proposed	A2, A3
B4, C2	joint losses; integration after 2020 (at home, after the lecture)	Proposed	B2, A2; A1, B1
Extensions	A1, A2, B1, B2 extended; A4–A6, B5, B6, C3 (optional)	both	on each task

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

Pearson Correlation

- **Pearson correlation**: the covariance of two returns divided by the product of their standard deviations

$$\rho = \frac{\mathrm{Cov}(X_1, X_2)}{\sigma_1 \sigma_2}, \qquad \hat{\rho} = \frac{\sum_{i=1}^{n}(x_{1i} - \bar{x}_1)(x_{2i} - \bar{x}_2)}{\sqrt{\sum_{i=1}^{n}(x_{1i} - \bar{x}_1)^2 \sum_{i=1}^{n}(x_{2i} - \bar{x}_2)^2}}$$

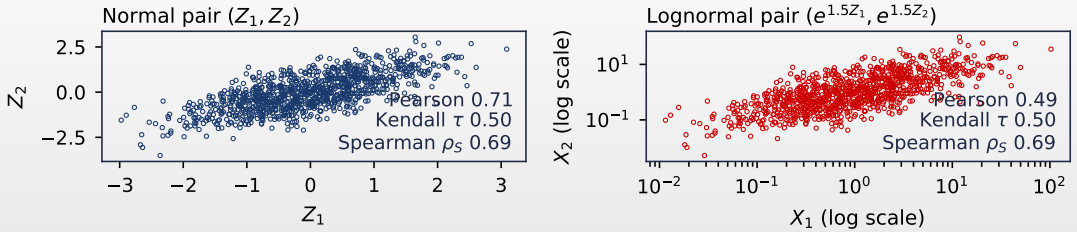
  - X_1, X_2: the two returns; σ_1, σ_2: their standard deviations; x_{1i}, x_{2i}: observation i of n; $\bar{x}_1$, $\bar{x}_2$: sample means; the hat marks an estimate
  - $\rho = 0$: no *linear* dependence, which does not imply independence
  - Needs finite variances and depends on the margins: an increasing transform of X_1 (e.g. e^{X_1}) changes ρ (next slide)

- **Beta** of X_2 on X_1 (C2): the slope of the regression of X_2 on X_1

$$\beta = \frac{\mathrm{Cov}(X_1, X_2)}{\mathrm{Var}(X_1)} = \rho \frac{\sigma_2}{\sigma_1}$$

  - $\beta = 0.5$: a 1% move of X_1 comes with a 0.5% move of X_2 on average

Pearson Against Rank Measures



- Left: 1 000 simulated pairs from a bivariate Normal distribution with $\rho = 0.7$; right: the same pairs after $x \mapsto e^{\mathbf{1.5}x}$ on both coordinates
- The dependence is the same, since the transform is increasing and keeps the ranks
  - Pearson falls from 0.71 to 0.49: it mixes the dependence with the shape of the margins
  - Kendall's τ (0.50) and Spearman's ρ_S (0.69) do not change: they depend only on the ranks, i.e. on the copula (next slides)
- Consequence (A6): with fixed margins, not every ρ in $[-1, 1]$ can be reached

Copulas and Sklar's Theorem

- **Sklar's theorem** Nelsen (2006): a joint distribution = its margins + a copula that holds all the dependence

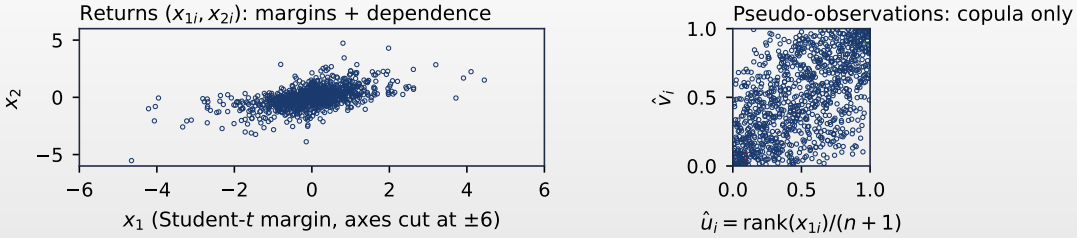
$$F(x_1, x_2) = C\big(F_1(x_1), F_2(x_2)\big), \qquad C(u, v) = P(U \le u,\ V \le v)$$

  - F: joint distribution function of (X_1, X_2); F_1, F_2: marginal distribution functions
  - $U = F_1(X_1)$, $V = F_2(X_2)$: the probability-integral transforms, both uniform on $[0, 1]$; the copula C is their joint distribution function
  - So the margins (heavy tails, volatility) and the dependence C can be modelled separately
- **Copula density** and the joint density

$$c(u, v) = \frac{\partial^2 C(u, v)}{\partial u\, \partial v}, \qquad f(x_1, x_2) = c\big(F_1(x_1), F_2(x_2)\big)\, f_1(x_1)\, f_2(x_2)$$

  - f_1, f_2: marginal densities; c is the factor by which the joint density differs from independence ($c \equiv 1$)
- Three reference copulas; every copula lies between the **Fréchet–Hoeffding bounds** $W \le C \le M$
  - $\Pi(u, v) = uv$: independence; $M(u, v) = \min(u, v)$: $V = U$ (comonotone); $W(u, v) = \max(u + v - 1, 0)$: $V = 1 - U$ (countermonotone)

Pseudo-Observations and the Empirical Copula



- In data the F_i are unknown: the **pseudo-observations** Genest, Ghoudi & Rivest (1995) replace them by ranks
  - $\hat{u}_i = \text{rank}(x_{1i})/(n+1)$, $\hat{v}_i = \text{rank}(x_{2i})/(n+1)$; dividing by $n+1$ keeps every value strictly inside $(0,1)$; example, $n=4$: $x = (0.3, -1.2, 2.0, 0.1)$ has ranks $(3,1,4,2)$, so $\hat{u} = (0.6, 0.2, 0.8, 0.4)$
  - **Empirical copula**: $C_n(u,v) = \frac{1}{n}\sum_{i=1}^{n} 1\{\hat{u}_i \le u,\ \hat{v}_i \le v\}$, with $1\{\cdot\} = 1$ if the condition holds and 0 otherwise
- Chart: simulated pairs with heavy-tailed margins (left) and their pseudo-observations (right); the corners are the joint extremes

Rank Correlations

- **Kendall's** τ: the probability of concordance minus the probability of discordance

$$\tau = P\big[(X_1 - X_1')(X_2 - X_2') > 0\big] - P\big[(X_1 - X_1')(X_2 - X_2') < 0\big], \qquad \hat{\tau} = \frac{n_c - n_d}{n(n-1)/2}$$

  - (X_1, X_2), (X_1', X_2'): two independent draws of the pair; concordant: both coordinates move in the same direction
  - n_c, n_d: numbers of concordant and discordant pairs among the $n(n-1)/2$ pairs of observations
- **Spearman's** ρ_S: the Pearson correlation of the ranks

$$\rho_S = \mathrm{corr}(U, V) = 12\,\mathbb{E}[UV] - 3, \qquad \hat{\rho}_S = \mathrm{corr}(\hat{u}_i, \hat{v}_i)$$

  - $\mathbb{E}[UV] = 1/4$ under independence, so $\rho_S = 0$; both τ and ρ_S lie in $[-1, 1]$ and depend only on C
- For the Gaussian copula with parameter ρ (next slide) both are known functions of ρ

$$\tau = \frac{2}{\pi} \arcsin \rho, \qquad \rho_S = \frac{6}{\pi} \arcsin \frac{\rho}{2}$$

  - Example: $\rho = 0.5$ gives $\tau = 0.333$ and $\rho_S = 0.483$; the first formula also holds for the t copula

Elliptical Copulas

- **Gaussian copula**: the dependence of a bivariate Normal distribution

$$C_\rho^{Ga}(u, v) = \Phi_\rho(\Phi^{-1}(u),\ \Phi^{-1}(v))$$

  - Φ: distribution function of $N(0, 1)$, Φ^{-1} its quantile function; Φ_ρ: distribution function of the bivariate Normal with unit variances and correlation $\rho \in (-1, 1)$

- **t copula**: the dependence of a bivariate Student t distribution

$$C_{\rho,\nu}^{t}(u, v) = t_{\rho,\nu}(t_\nu^{-1}(u),\ t_\nu^{-1}(v))$$

  - t_ν: Student t distribution function with ν degrees of freedom, t_ν^{-1} its quantile function; $t_{\rho,\nu}$: bivariate t distribution function with correlation parameter ρ
  - ρ: strength of the dependence; ν: how often extremes arrive together, smaller ν = more joint extremes; $\nu \to \infty$ gives back the Gaussian copula

- Both are symmetric: joint lows and joint highs are equally likely
  - Kendall link for both: $\rho = \sin(\pi\tau/2)$, e.g. $\tau = 0.5$ gives $\rho = 0.707$

Archimedean Copulas

- **Clayton** ($\theta > 0$): dependence concentrated in joint *lows*

$$C_\theta^{Cl}(u,v) = \left(u^{-\theta} + v^{-\theta} - 1\right)^{-\mathbf{1}/\theta}, \qquad \tau = \frac{\theta}{\theta+2} \iff \theta = \frac{2\tau}{1-\tau}$$

- **Gumbel** ($\theta \geq 1$): dependence concentrated in joint *highs*

$$C_\theta^{Gu}(u,v) = \exp\Big(-\left[(-\ln u)^\theta + (-\ln v)^\theta\right]^{\mathbf{1}/\theta}\Big), \qquad \tau = 1 - \frac{1}{\theta} \iff \theta = \frac{1}{1-\tau}$$

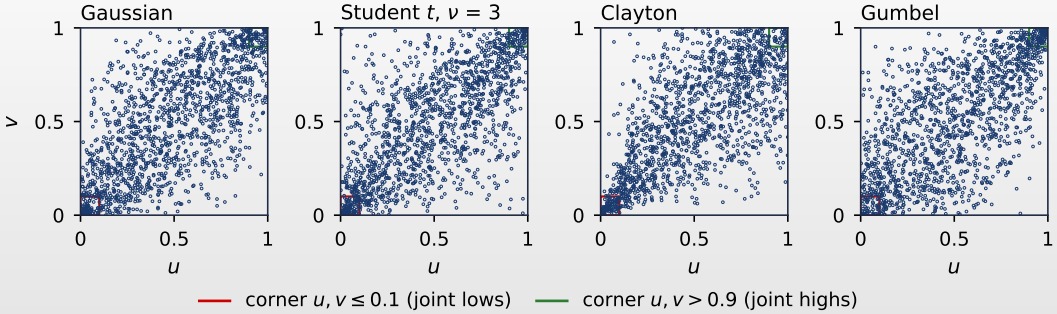
- **Frank** ($\theta \neq 0$): symmetric, no joint extremes

$$C_\theta^{Fr}(u,v) = -\frac{1}{\theta}\ln\left(1 + \frac{(e^{-\theta u}-1)(e^{-\theta v}-1)}{e^{-\theta}-1}\right), \qquad \tau = 1 - \frac{4}{\theta}\left[1 - D_{\mathbf{1}}(\theta)\right]$$

  - $D_{\mathbf{1}}(\theta) = \frac{\mathbf{1}}{\theta}\int_{\mathbf{0}}^{\theta} \frac{t}{e^t - \mathbf{1}}\,dt$: the Debye function; $\tau(\theta)$ has no closed-form inverse, so θ is found numerically (A2)

- θ: the dependence parameter; a larger θ means stronger dependence; Clayton $\theta \to 0$ and Gumbel $\theta = 1$ give independence uv

Four Families with the Same Kendall Tau



- 1 500 simulated pairs from each family, all with $\tau = 0.5$: Gaussian $\rho = 0.707$, t $\rho = 0.707$ and $\nu = 3$, Clayton $\theta = 2$, Gumbel $\theta = 2$
  - The clouds look alike in the middle; they differ in the corners: Clayton crowds the lower corner, Gumbel the upper one, t both
- **Survival (rotated) copula**: the copula of $(1 - U, 1 - V)$, used for losses $L = -X$ (B4)
  - $\hat{C}(u, v) = u + v - 1 + C(1 - u, 1 - v)$: it swaps the two corners, so a rotated Clayton has its dependence in joint highs

Tail Dependence

- **Tail dependence** Demarta & McNeil (2005): the chance that one variable is extreme, given that the other one is extreme

$$\lambda_L = \lim_{q \downarrow \mathbf{0}} P(V \le q \mid U \le q) = \lim_{q \downarrow \mathbf{0}} \frac{C(q,q)}{q}, \qquad \lambda_U = \lim_{q \downarrow \mathbf{0}} P(V > 1-q \mid U > 1-q)$$

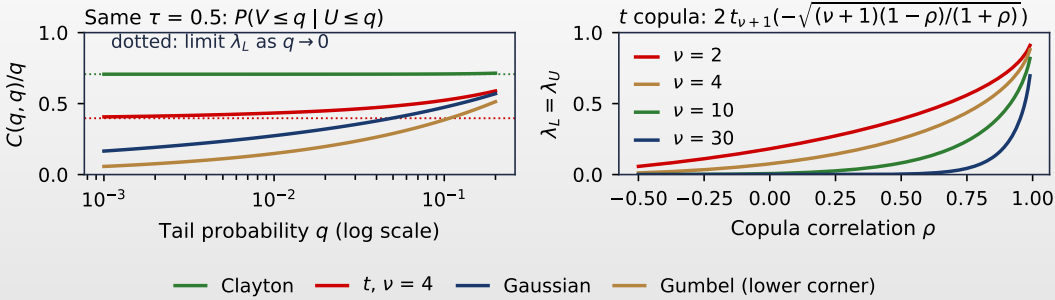
  - q: tail probability (e.g. $q = 5\%$: the worst 5% of days); λ_L: lower tail (joint crashes), λ_U: upper tail (joint rallies); both in $[0, 1]$
  - $\lambda > 0$: extremes arrive together however extreme; $\lambda = 0$: asymptotic independence
  - At a finite q, $C(q,q)/q$ is a probability we can estimate (B2: $C_n(q,q)/q$); λ is only its limit

- Closed forms

$$\lambda_L^t = \lambda_U^t = 2\, t_{\nu+\mathbf{1}}\Big(-\sqrt{\tfrac{(\nu+\mathbf{1})(\mathbf{1}-\rho)}{\mathbf{1}+\rho}} \Big), \qquad \lambda_L^{Cl} = 2^{-\mathbf{1}/\theta}, \qquad \lambda_U^{Gu} = 2 - 2^{\mathbf{1}/\theta}$$

  - $t_{\nu+\mathbf{1}}$: Student t distribution function with $\nu+1$ degrees of freedom (A3 derives it); Gaussian: $\lambda = 0$ for $|\rho| < 1$; Clayton $\lambda_U = 0$, Gumbel $\lambda_L = 0$, Frank $\lambda_L = \lambda_U = 0$
  - Example, $\tau = 0.5$: Gaussian $\lambda = 0$; Clayton $\theta = 2$, $\lambda_L = 2^{-\mathbf{1}/\mathbf{2}} = 0.707$; Gumbel $\theta = 2$, $\lambda_U = 2 - \sqrt{2} = 0.586$: the same τ, very different joint crashes

Joint Crashes at a Finite Level



- Left: $C(q,q)/q$ for four copulas with $\tau = 0.5$, from the exact distribution functions; dotted: the limits λ_L
  - At $q = 10\%$ the four values are close; at $q = 0.1\%$ Clayton keeps 0.71, t 0.41, the Gaussian falls to 0.17 and keeps falling
- Right: the t-copula limit grows with ρ and shrinks as ν grows; it is positive even for $\rho < 0$

Fitting a Copula: Likelihood and AIC

- **Canonical maximum likelihood** (CML) Genest, Ghoudi & Rivest (1995): maximise the copula log-likelihood on the pseudo-observations; the margins stay unmodelled

$$\ell(\theta) = \sum_{i=\mathbf{1}}^{n} \ln c_\theta(\hat{u}_i, \hat{v}_i), \qquad \hat{\theta} = \arg\max_\theta \ell(\theta)$$

  - c_θ: copula density with parameter vector θ (e.g. (ρ, ν) for the t copula); a higher ℓ = the pairs are more likely under the model

- **AIC** (Akaike information criterion): fit penalised by the number of parameters

$$\mathrm{AIC} = -2\,\ell(\hat{\theta}) + 2k$$

  - k: number of copula parameters; the family with the *lower* AIC is preferred; a difference of more than 10 points is strong support
  - Families with the same k (Clayton, Gumbel, Frank, Gaussian: $k = 1$) are ranked simply by ℓ

The Likelihood-Ratio Test

- **Likelihood ratio** (LR) for a restricted model H_0 inside a larger model H_1

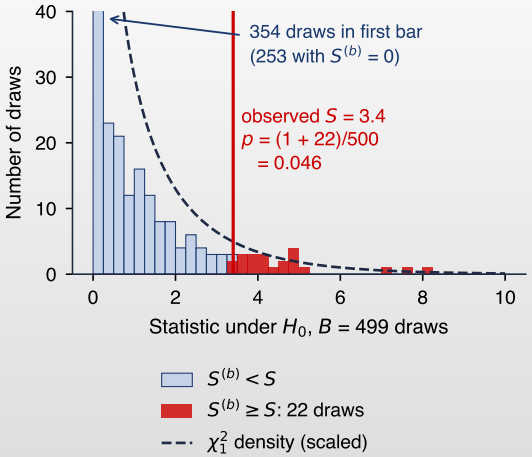
$$LR = 2\big(\ell_1(\hat{\theta}_1) - \ell_0(\hat{\theta}_0)\big) \geq 0$$

  - ℓ_1, ℓ_0: maximised log-likelihoods of the two models; large LR = the restriction costs much fit = evidence against H_0
- Reference distribution when the H_0 value is *inside* the parameter space

$$LR \xrightarrow{d} \chi^2_r, \qquad p = P(\chi^2_r \geq LR)$$

  - r: number of restrictions; $\xrightarrow{d}$: convergence in distribution as $n \to \infty$; e.g. $r = 1$: reject at 5% when $LR > 3.84$
  - On the *boundary* the χ^2_r reference fails: the Gaussian copula is the t copula with $\nu = \infty$
  - Then the p-value comes from a parametric bootstrap (next slide)

Parametric Bootstrap p-Value



- When the reference distribution is unknown, simulate it under H_0
  - fit the H_0 model; simulate B samples of size n from it
  - on each sample: re-rank, refit both models, recompute the statistic $S^{(b)}$, $b = 1, \ldots, B$
- p-value: the share of simulated statistics at least as large as the observed S

$$p = \frac{1 + \#\{b : \; S^{(b)} \geq S\}}{B + 1}$$

  - $\#\{\cdot\}$: number of draws; with $B = 499$ the smallest possible p is $1/500 = 0.002$
  - Chart (simulated): half of the null draws are exactly 0, so the χ^2_1 tail ($p = 0.065$) is too large here; the bootstrap gives $p = 0.046$

Goodness of Fit: the Rosenblatt Transform

- **Rosenblatt transform**: if the copula is right, it turns the pairs into *independent* uniforms

$$e_{1i} = \hat{u}_i, \qquad e_{2i} = C_{\hat{\theta}}(\hat{v}_i \mid \hat{u}_i) = \frac{\partial C_{\hat{\theta}}(u,v)}{\partial u}\Big|_{(\hat{u}_i,\hat{v}_i)}$$

  - $C(v \mid u) = P(V \le v \mid U = u)$: the conditional distribution function of V given $U = u$
  - Gaussian: $C(v \mid u) = \Phi\big((y - \rho x)/\sqrt{1-\rho^2}\big)$, $x = \Phi^{-1}(u)$, $y = \Phi^{-1}(v)$
  - t: $C(v \mid u) = t_{\nu+1}\Big((y - \rho x)/\sqrt{(\nu + x^2)(1-\rho^2)/(\nu+1)}\Big)$, $x = t_\nu^{-1}(u)$, $y = t_\nu^{-1}(v)$ (the result used in A3)

- **Test statistic** Genest, Rémillard & Beaudoin (2009): the Cramér–von Mises distance of (e_{1i}, e_{2i}) from independence

$$S_n = \sum_{i=1}^{n} \big[D_n(e_{1i}, e_{2i}) - e_{1i}\, e_{2i}\big]^2, \qquad D_n(a,b) = \frac{1}{n}\sum_{j=1}^{n} 1\{e_{1j} \le a,\ e_{2j} \le b\}$$

  - D_n: empirical distribution function of the transformed pairs; $e_1 e_2$: its value under independence; large S_n = misfit; p-value by the parametric bootstrap of the previous slide

- B2 and B4 use the integral over the unit square, $S_n^{(B)} = n \int_0^1 \int_0^1 [D_n(a,b) - ab]^2\, da\, db$, bootstrapping the same statistic; its values are not comparable with S_n above

The Rosenblatt Transform on Simulated Data



- 2 000 pairs simulated from a t copula ($\rho = 0.5$, $\nu = 2$), transformed with the right t copula (left) and with a Gaussian copula with the same τ (right)
- Left: a uniform cloud, $S_n = 0.08$; right: points pile up near the edges, where the Gaussian copula cannot explain the joint extremes, $S_n = 0.52$

GARCH(1,1) and Standardised Residuals

- **GARCH(1,1)** (Chapter 5): today's variance is a weighted sum of a constant, yesterday's squared shock and yesterday's variance

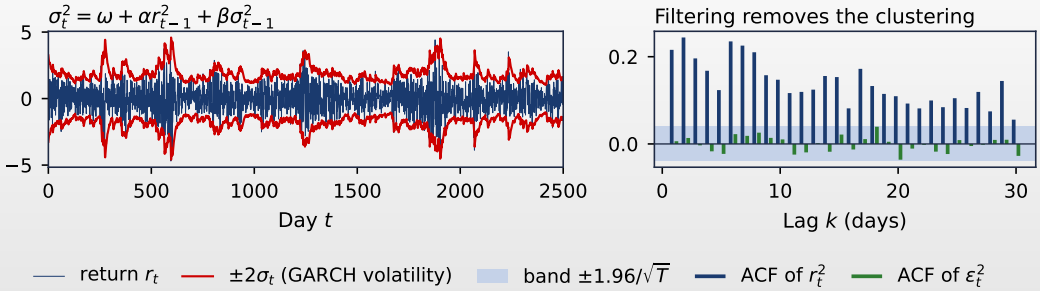
$$r_{it} = \mu_i + \sigma_{it}\varepsilon_{it}, \qquad \sigma^2_{it} = \omega_i + \alpha_i \left(r_{i,t-1} - \mu_i\right)^2 + \beta_i \, \sigma^2_{i,t-1}$$

  - r_{it}: return of asset i on day t; μ_i: constant mean; σ_{it}: conditional volatility; ε_{it}: independent and identically distributed (i.i.d.) shock with mean 0, variance 1
  - $\omega_i > 0$; $\alpha_i \geq 0$: reaction to yesterday's shock; $\beta_i \geq 0$: memory; $\alpha_i + \beta_i$ measures persistence; $\alpha_i + \beta_i < 1$ gives covariance stationarity, with long-run variance $\omega_i/(1 - \alpha_i - \beta_i)$
  - Annualised volatility from a daily return variance σ^2 in %: $\sqrt{252\,\sigma^2}$ %
- Two markets can both become more volatile while their standardised shocks remain equally correlated: a higher raw crisis correlation is not by itself evidence of contagion
- **Standardised residuals**: the returns divided by their fitted volatility

$$\hat{\varepsilon}_{it} = \frac{r_{it} - \hat{\mu}_i}{\hat{\sigma}_{it}}$$

  - If the model is right they have no volatility clustering: the autocorrelation of $\hat{\varepsilon}^2_{it}$ stays inside $\pm 1.96/\sqrt{T}$ (B1, task 1)
  - Parameters by quasi-maximum likelihood (QML): the Normal likelihood, valid even when ε_{it} is not Normal

GARCH Filtering on Simulated Data



- Left: 2 500 days simulated from a GARCH(1,1) with $\omega = 0.02$, $\alpha = 0.10$, $\beta = 0.88$, and the band $\pm 2\sigma_t$: calm and turbulent periods alternate
- Right: the ACF (autocorrelation function) of r_t^2 is about 0.2 at every lag up to 30; after dividing by σ_t, the ACF of ε_t^2 stays inside the band
- Why it matters for correlation: without filtering, a correlation estimate mixes the dependence with the volatility regimes

Two-Step DCC: the Recursion

- Step 1: a GARCH(1,1) per series gives $\varepsilon_t = (\varepsilon_{1t}, \varepsilon_{2t})^\top$; step 2, the DCC (Dynamic Conditional Correlation) model Engle (2002): a GARCH-type recursion for a 2×2 matrix Q_t, rescaled to a correlation matrix R_t

$$Q_t = (1 - a - b)\,\bar{Q} + a\,\varepsilon_{t-1}\varepsilon_{t-1}^\top + b\,Q_{t-1}, \qquad R_t = Q_t^{*-1/2}\,Q_t\,Q_t^{*-1/2}, \qquad \rho_{12,t} = \frac{Q_{12,t}}{\sqrt{Q_{11,t}Q_{22,t}}}$$

- Notation
  - $a \geq 0$: reaction to yesterday's joint shock $\varepsilon_{t-1}\varepsilon_{t-1}^\top$; $b \geq 0$: memory; $a + b < 1$: mean reversion towards $\bar{Q}$
  - $\bar{Q}$: fixed target parameter, replaced in estimation by the sample covariance of ε_t (correlation targeting); $Q_t^* = \mathrm{diag}(Q_t)$: the diagonal of Q_t; R_t: conditional correlation, $\mathbb{E}(\varepsilon_t\varepsilon_t^\top \mid \mathcal{F}_{t-1}) = R_t$, $\mathcal{F}_{t-1}$: information up to day $t-1$
  - CCC (constant conditional correlation): $a = b = 0$, so $R_t = \bar{Q}$ every day when the target has unit diagonal; otherwise normalise the target by its diagonal
- Example: $\bar{Q}$ with unit diagonal and $\bar{Q}_{12} = 0.3$, $Q_{t-1} = \bar{Q}$, $a = 0.05$, $b = 0.90$, $\varepsilon_{t-1} = (2, -1)$:
  $Q_{12} = 0.015 - 0.100 + 0.270 = 0.185$, $Q_{11} = 1.15$, $Q_{22} = 1.00$, $\rho_{12,t} = 0.185/\sqrt{1.15} = 0.173$

Two-Step DCC: Estimation and Persistence

- **Correlation part of the quasi-likelihood**, maximised over (a, b) with ε_t and $\bar{Q}$ fixed

$$\ell_C(a,b) = -\frac{1}{2}\sum_{t=\mathbf{1}}^{T}\left(\ln|R_t| + \varepsilon_t^\top R_t^{-\mathbf{1}}\varepsilon_t - \varepsilon_t^\top \varepsilon_t\right)$$

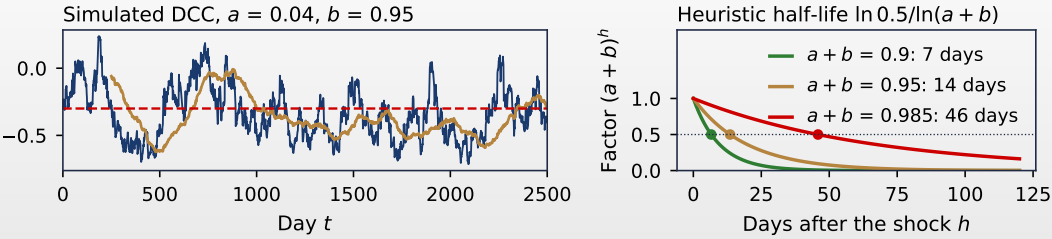
  - $|R_t|$: determinant of R_t; $\varepsilon_t^\top R_t^{-\mathbf{1}}\varepsilon_t$: how surprising the joint shock is under R_t; the last term does not depend on (a, b)
  - Standard errors from the inverse Hessian (matrix of second derivatives of $-\ell_C$) treat step 1 as known; a bootstrap of both steps Newey & McFadden (1994) adds the step-1 uncertainty

- **Persistence**: the heuristic measure $h_{\mathbf{1}/\mathbf{2}} = \ln 0.5 / \ln(a+b)$ summarises how slowly Q_t returns to $\bar{Q}$

$$\mathbb{E}_t[Q_{t+\mathbf{2}}] - \bar{Q} = a\,(R_{t+\mathbf{1}} - \bar{Q}) + b\,(Q_{t+\mathbf{1}} - \bar{Q})$$

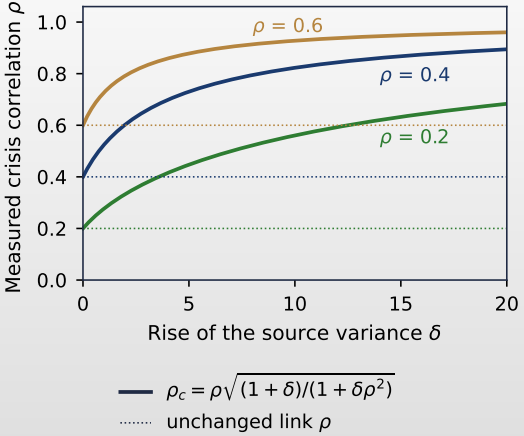
  - $\mathbb{E}_t[\varepsilon_{t+\mathbf{1}}\varepsilon_{t+\mathbf{1}}^\top] = R_{t+\mathbf{1}} \neq Q_{t+\mathbf{1}}$, so the deviation does not shrink exactly by the factor $a+b$: $h_{\mathbf{1}/\mathbf{2}}$ is not an exact half-life; the actual response of the correlation comes from a conditional simulation (A5)
  - With the same future shocks ε, the difference between two paths of Q shrinks by the factor b: a different experiment from the one described by $a+b$
  - Test of CCC: $LR = 2(\ell_C^{DCC} - \ell_C^{CCC})$; under $a = 0$ the parameter b is not identified, so the p-value comes from a bootstrap under CCC

DCC on Simulated Data



- Left: 2 500 days simulated from a DCC with $\bar{Q}_{\mathbf{12}} = -0.3$, $a = 0.04$, $b = 0.95$; the daily $\rho_{\mathbf{12},t}$ moves between -0.71 and 0.24 around the target
- The rolling 250-day correlation follows the same swings with a delay and smooths them; a constant correlation would be the dashed line
- Right: the factor $(a+b)^h$ and the heuristic measure $h_{\mathbf{1/2}}$; persistence close to 1 means slow mean reversion

The Forbes–Rigobon Adjustment



- Model Forbes & Rigobon (2002): $y_t = \alpha + \beta x_t + \varepsilon_t$, $\mathrm{Cov}(x_t, \varepsilon_t) = 0$; in the crisis only the variance of the source x rises

$$\rho_c = \rho\sqrt{\frac{1+\delta}{1+\delta\rho^2}}, \qquad \delta = \frac{\sigma^2_{x,c}}{\sigma^2_{x,calm}} - 1$$

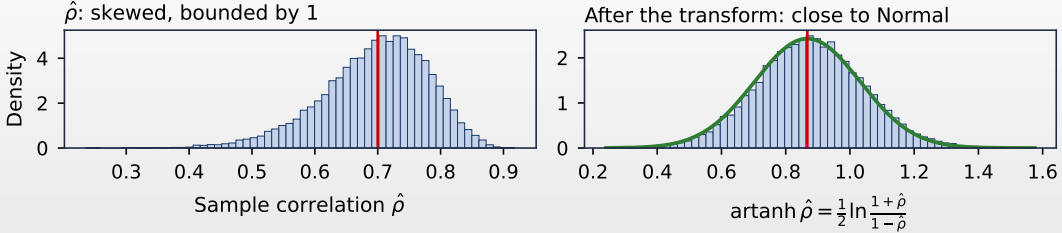
  - ρ: calm-period correlation; ρ_c: crisis correlation measured with the *same* link β; δ: relative rise of the source variance
  - Chart: $\rho_c > \rho$ whenever $\delta > 0$: higher volatility alone raises the measured correlation

- Adjusted crisis correlation: the crisis correlation net of the volatility effect

$$\rho^* = \frac{\rho_c}{\sqrt{1+\delta(1-\rho_c^2)}}$$

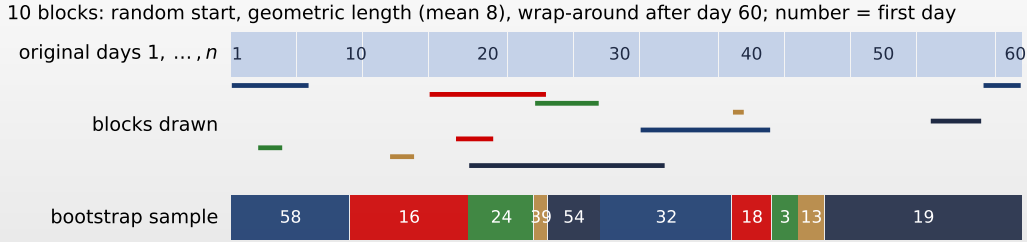
  - Contagion: ρ^* significantly above ρ_{calm}; otherwise interdependence

Comparing Two Correlations: the Fisher z Test



- Fisher's transform $\operatorname{artanh}\hat{\rho} = \frac{1}{2}\ln\frac{1+\hat{\rho}}{1-\hat{\rho}}$ is close to Normal with variance $1/(n-3)$ (chart, $n = 40$, $\rho = 0.7$)
  - $z = \dfrac{\operatorname{artanh}\hat{\rho}_1 - \operatorname{artanh}\hat{\rho}_0}{\sqrt{1/(n_1-3) + 1/(n_0-3)}} \approx N(0,1)$ under $H_0 : \rho_1 = \rho_0$; $\hat{\rho}_1$, $\hat{\rho}_0$: correlations from n_1 and n_0 independent observations
  - One-sided test of $H_0 : \rho_1 \le \rho_0$: $p = 1 - \Phi(z)$; reject at 5% when $z > 1.645$; assumes i.i.d. Normal pairs and a test date fixed before seeing the data
- **Fisher F** (used in A1 extension): under i.i.d. normality the ratio of the crisis to the calm sample variance, each divided by its true variance, has the distribution $F(n_c - 1, n_0 - 1)$

Multi-Day Returns and the Stationary Bootstrap



- 2-day log return (B3): $r_t^{(2)} = \ln P_t - \ln P_{t-2}$, P_t: closing level on day t
  - Consecutive 2-day returns share one day (they overlap); so we use half their count, rounded down, as the effective sample size: a heuristic, not an exact identity
- **Stationary bootstrap** Politis & Romano (1994): resample blocks of consecutive days with random starts and geometric lengths of mean $\bar{L}$, so the serial dependence inside a block is kept
  - Recompute the statistic D on each of B resamples; 90% percentile interval: the 5% and 95% quantiles of the $D^{(b)}$; one-sided p for $H_0 : D \leq 0$: the share of $D^{(b)} \leq 0$, an approximation by inverting the percentile interval, not a simulation of the statistic under H_0

A1 [Solved]: How Much of a Crisis Correlation Is Volatility?

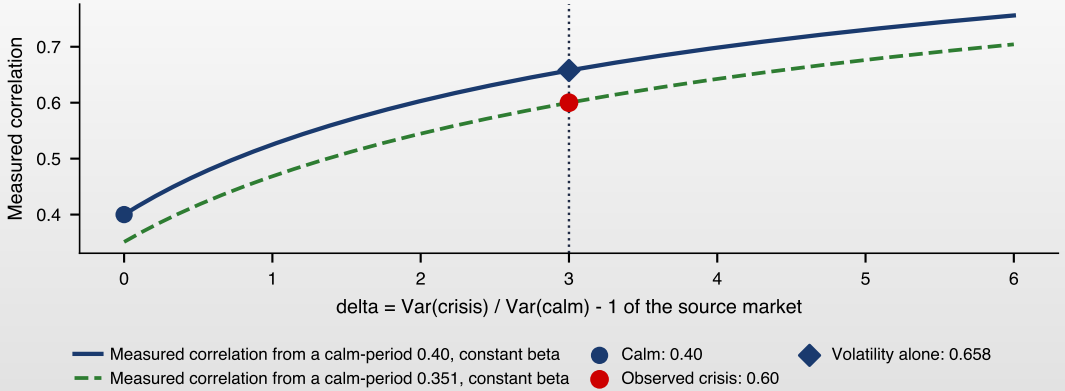
- **Question**: a correlation rises from 0.40 in a calm year to 0.60 in a crisis while the source market's variance quadruples; is this contagion?
- **Model**: $y_t = \alpha + \beta x_t + \varepsilon_t$, $\mathrm{Cov}(x_t, \varepsilon_t) = 0$, β and σ_ε^2 constant; in the crisis $\sigma_{x,c}^2 = (1+\delta)\sigma_x^2$
  - Start from the covariance: $\mathrm{Cov}(x, y) = \beta\sigma_x^2$, $\mathrm{Var}(y) = \beta^2\sigma_x^2 + \sigma_\varepsilon^2$, hence $\rho^2 = \beta^2\sigma_x^2/(\beta^2\sigma_x^2 + \sigma_\varepsilon^2)$, which gives ρ_c (slide «The Forbes–Rigobon Adjustment»)
- **Tasks**

1. Compute δ and the adjusted crisis correlation ρ^* (slide «The Forbes–Rigobon Adjustment»).
2. Compute the crisis correlation we would measure from the volatility rise alone, with a constant β and a calm-period correlation of 0.40.
3. State whether the rise is evidence of contagion under the maintained assumptions, and name the assumption you trust least.

- **Report**: each number with its formula and substituted values, and one sentence of interpretation; notebook: section A1

Back Exercise list

A1 [Solved]: Solution (1/2)



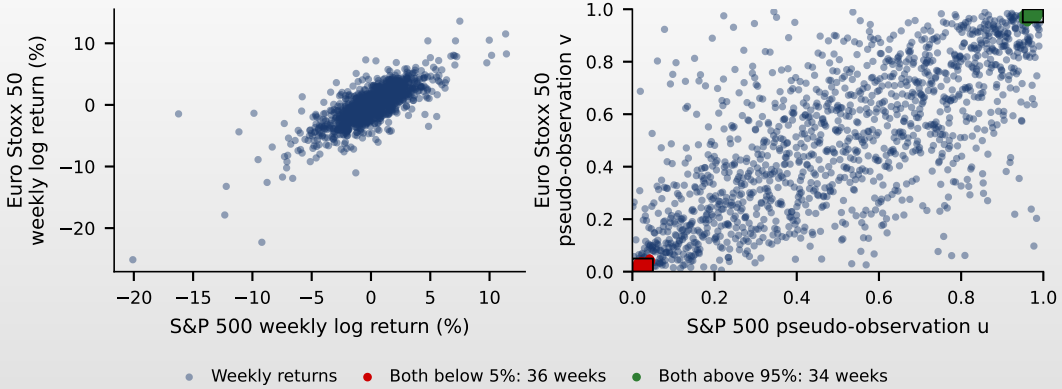
MFM_ch6_seminar

A1 [Solved]: Solution (2/2)

- (1) Variance ratio 4, so $\delta = 3$; $\rho^* = 0.60/\sqrt{1 + 3(1 - 0.36)} = 0.60/\sqrt{2.92} = 0.60/1.709 = 0.351$
- (2) $\rho_c = 0.40\sqrt{4/(1 + 3 \times 0.16)} = 0.40\sqrt{4/1.48} = 0.40 \times 1.644 = 0.658$
- (3) The adjusted 0.351 is below the calm 0.40, and volatility alone would lift 0.40 to 0.658, above the observed 0.60: no evidence of an increase after this adjustment
- Least credible: no common shock; a global shock raises σ_ε too, and the adjustment then over-corrects (A4, lecture)
- The observed 0.60 lies below the curve of a calm-period correlation of 0.40 with a constant β
- It lies on the curve of a calm-period correlation of 0.351

MFM_ch6_seminar

A2 [Solved]: Kendall's Tau and Copula Parameters — Task (1/2)



MFM_ch6_seminar

A2 [Solved]: Kendall's Tau and Copula Parameters — Task (2/2)

- **Questions**
  - which copula parameter reproduces the rank dependence of two equity indices?
  - does it fix the tail risk?
- **Inputs**: weekly S&P 500 and Euro Stoxx 50 log returns (Data Used), $n = 1\,393$; $\hat{\tau} = 0.5642$, $\hat{\rho}_S = 0.7462$, Pearson 0.781; pairs treated as i.i.d. (an unconditional weekly working model)
- **Tasks**

1. Invert $\hat{\tau}$ for the Gaussian (slide «Elliptical Copulas»), Clayton, Gumbel and Frank copulas (slide «Archimedean Copulas»); for Frank solve $\tau(\theta) = \hat{\tau}$ numerically (helper `theta_from_tau`).
2. Compute the tail dependence implied by each fitted copula.
3. Compute the Spearman's ρ_S implied by the Gaussian copula (slide «Rank Correlations») and compare it with $\hat{\rho}_S$.

- **Report**: a table of the four parameters and their λ_L, λ_U, and the gap of task 3

- Left: weekly returns; right: pseudo-observations, boxes at the 5% corners; notebook: section A2

MFM_ch6_seminar

Back Exercise list

A2 [Solved]: Solution

Copula	Parameter	λ_L	λ_U
Gaussian	$\rho = 0.775$	0	0
Clayton	$\theta = 2.589$	0.765	0
Gumbel	$\theta = 2.295$	0	0.647
Frank	$\theta = 7.044$	0	0

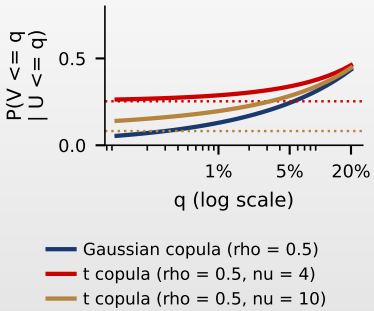
- Gaussian: $\sin(\pi \times 0.564/2) = 0.775$; Clayton: $2 \times 0.564/0.436 = 2.589$; Gumbel: $1/0.436 = 2.295$; Frank: numerical root
- Clayton $\lambda_L = 2^{-1/\theta}$, Gumbel $\lambda_U = 2 - 2^{1/\theta}$
- Task 3: $\frac{6}{\pi} \arcsin(0.775/2) = 0.760$ against $\hat{\rho}_S = 0.746$: a gap of 0.0134 (its precision: extension A2)
- The same $\hat{\tau}$ gives tail dependence between 0 and 0.77, depending only on the family chosen
- Interpretation: $\hat{\tau}$ identifies a parameter within a family, never the family; the joint tail must be estimated (B2), not read off τ

A3 [Solved]: Tail Dependence of the t Copula

- **Question**: in the chart, as q falls the Gaussian probability heads to 0 while the t curves level off; what happens as $q \to 0$?
- **Inputs**: $\rho = 0.5$, $\nu \in \{3, 4, 10, 30\}$; the conditional t distribution $C(v \mid u)$ (slide «Goodness of Fit: the Rosenblatt Transform»)
- **Tasks**

1. Use L'Hôpital's rule and exchangeability to show that $\lambda_L = 2 \lim_{q \downarrow 0} P(V \le q \mid U = q)$.
2. Substitute the conditional t distribution, let $x = t_\nu^{-1}(q) \to -\infty$ and derive $\lambda = 2\, t_{\nu+1}\big(-\sqrt{(\nu+1)(1-\rho)/(1+\rho)}\big)$.
3. Repeat the argument for the Gaussian copula and show that $\lambda = 0$ for $|\rho| < 1$.
4. Evaluate λ for each ν, and $C(q,q)/q$ at $q = 5\%$ and 0.1% for the Gaussian and the t copula with $\nu = 4$, same ρ (notebook helper `joint_tail_prob`).

- **Report**: the derivation as a chain of equalities and a table of λ and of the four probabilities



- Exact $P(V \le q \mid U \le q)$ from the bivariate distribution functions; dotted: the limits λ; notebook: section A3

MFM_ch6_seminar

Back Exercise list

A3 [Solved]: Solution (1/2) — Derivation

- (1) $C(q,q) \to 0$ and $q \to 0$, so by L'Hôpital $\lambda_L = \lim \frac{d}{dq} C(q,q) = \lim[\partial_1 C(q,q) + \partial_2 C(q,q)]$, with $\partial_1 C(u,v) = P(V \le v \mid U = u)$; exchangeability makes the two terms equal: $\lambda_L = 2 \lim P(V \le q \mid U = q)$
- (2) With $x = t_\nu^{-1}(q)$: $P(V \le q \mid U = q) = P(X_2 \le x \mid X_1 = x) = t_{\nu+1}\Big(\frac{(1-\rho)\,x}{\sqrt{(\nu + x^2)(1-\rho^2)/(\nu+1)}}\Big)$
- As $x \to -\infty$, $x/\sqrt{\nu + x^2} \to -1$; since $(1-\rho)/\sqrt{1-\rho^2} = \sqrt{(1-\rho)/(1+\rho)}$, the argument tends to $-\sqrt{(\nu+1)(1-\rho)/(1+\rho)}$, which gives the formula
- (3) Gaussian: $P(V \le q \mid U = q) = \Phi\big(x\sqrt{(1-\rho)/(1+\rho)}\big)$ with $x = \Phi^{-1}(q) \to -\infty$, so the limit is 0
- Why: in the t case the conditional scale grows with $|x|$, so the standardised argument stays bounded; in the Gaussian case the scale is fixed and the argument diverges

A3 [Solved]: Solution (2/2) — Numbers

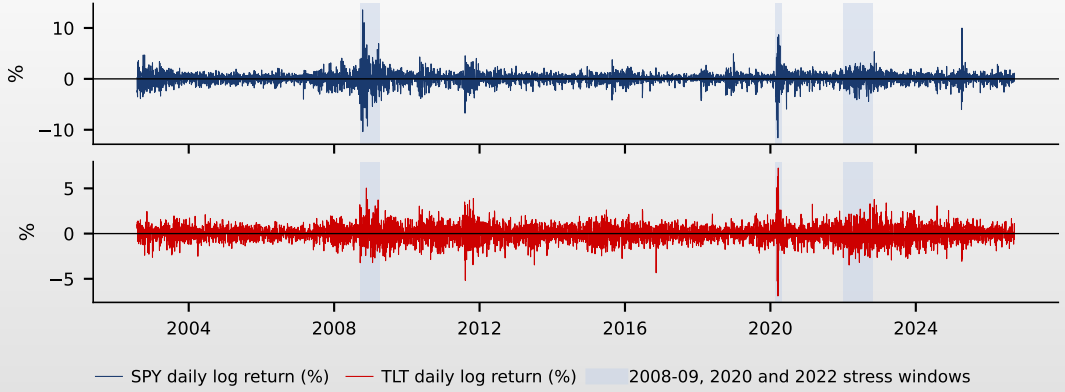
ν	$a_\nu = \sqrt{(\nu+1)/3}$	$\lambda = 2t_{\nu+1}(-a_\nu)$
3	1.155	0.313
4	1.291	0.253
10	1.915	0.082
30	3.215	0.003

- $\nu = 4$ step by step: $t_5(-1.291) = 0.1266$, so $\lambda = 2 \times 0.1266 = 0.253$
- Tail dependence vanishes as ν grows

	$q = 5\%$	$q = 0.1\%$
Gaussian, $\rho = 0.5$	0.244	0.054
t, $\rho = 0.5$, $\nu = 4$	0.339	0.263

- $C(q,q)/q = P(V \le q \mid U \le q)$ with the same ρ in both models
- Interpretation: at 5% the two models look alike, but at 0.1% the Gaussian probability is a fraction of the t one; a model judged on moderate losses cannot reveal joint crashes

B1 [Solved]: DCC for SPY and TLT — Task (1/2)



MFM_ch6_seminar

B1 [Solved]: DCC for SPY and TLT — Task (2/2)

- **Questions**
  - how has the stock–bond correlation moved since 2002?
  - is it constant?
- **Data**: SPY and TLT adjusted closes, aligned on common days, log returns in %, $T = 6\,073$
- **Tasks**

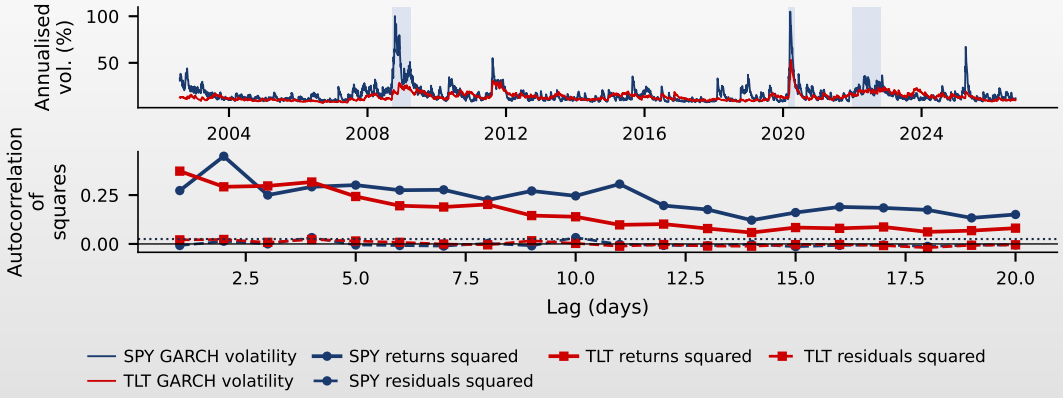
1. Fit a GARCH(1,1) with constant mean to each series, form $\varepsilon_{it} = (r_{it} - \hat{\mu}_i)/\hat{\sigma}_{it}$ and compare the autocorrelation of $r_{it}^{\mathbf{2}}$ and of $\varepsilon_{it}^{\mathbf{2}}$.
2. Set $\bar{Q}$ to the sample covariance of ε_t, $Q_{\mathbf{1}} = \bar{Q}$, compute $R_{\mathbf{2}}$ from $\varepsilon_{\mathbf{1}}$, then maximise $\ell_C(a, b)$ (slide «Two-Step DCC: Estimation and Persistence») subject to $a, b \geq 0$, $a + b < 1$.
3. Plot $\rho_{\mathbf{12},t}$ against the 252-day rolling correlation and the target, and re-estimate on 2002–2019 and 2020–2026.
4. Test CCC ($a = b = 0$) with the likelihood ratio, calibrated by a bootstrap of both steps under CCC (199 paths).

- **Report**: GARCH and DCC tables, $R_{\mathbf{2}}$, the path chart, the LR with its bootstrap p, one interpretation; notebook: section B1

MFM_ch6_seminar

Back Exercise list

B1 [Solved]: GARCH Volatility Paths



MFM_ch6_seminar

B1 [Solved]: GARCH Parameter Estimates

	ω	α	β	mean ann. vol.
SPY	0.0283	0.127	0.848	16.0%
TLT	0.0057	0.049	0.943	13.6%

- Returns in %; volatility peaks: SPY 105% (17 March 2020), TLT 53% (24 March 2020)
- Checkpoint: lag-1 autocorrelation of squares, SPY 0.273 $\rightarrow -0.008$, TLT 0.373 $\rightarrow$ 0.021; mean over lags 1–20 SPY 0.233 $\rightarrow -0.002$
- Filtering removes the volatility clustering: the correlation of ε_t is no longer driven by volatility regimes

MFM_ch6_seminar

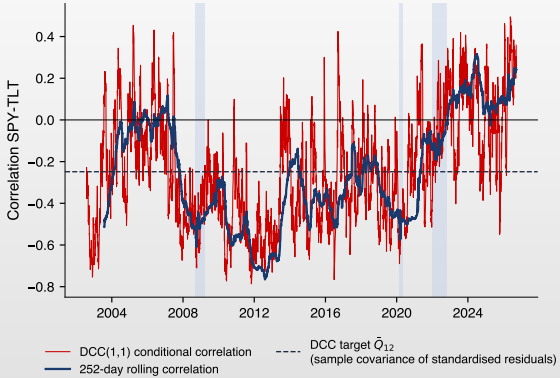
B1 [Solved]: Step 2, the DCC Estimates

- One recursion step: $\bar{Q} = \begin{pmatrix} 0.998373 & -0.248487 \\ -0.248487 & 1.000091 \end{pmatrix}$, $\varepsilon_{\mathbf{1}} = (0.079960,\ 1.497987)$ on 31 July 2002
- $Q_{\mathbf{2}} = (1-a-b)\bar{Q} + a\,\varepsilon_{\mathbf{1}}\varepsilon_{\mathbf{1}}^{\top} + b\bar{Q}$: $Q_{\mathbf{2},\mathbf{11}} = 0.946560$, $Q_{\mathbf{2},\mathbf{22}} = 1.065060$, $Q_{\mathbf{2},\mathbf{12}} = -0.229252$
- $R_{\mathbf{2},\mathbf{12}} = -0.229252/\sqrt{0.946560 \times 1.065060} = -0.2283$: a positive TLT shock with a flat SPY day moves the correlation towards zero

	estimate	se (standard error, Hessian)
a	0.052232	0.0059
b	0.933518	0.0084
$a+b$	0.985749	–
$\bar{Q}_{12}$	−0.248	–

- Heuristic persistence measure $\ln 0.5/\ln(a+b) = 48.3$ days, not the exact half-life of Q_t, since $R_t \neq Q_t$
- Hessian se treats step 1 as known; both-step bootstrap: extension B1

B1 [Solved]: Path and Subsamples

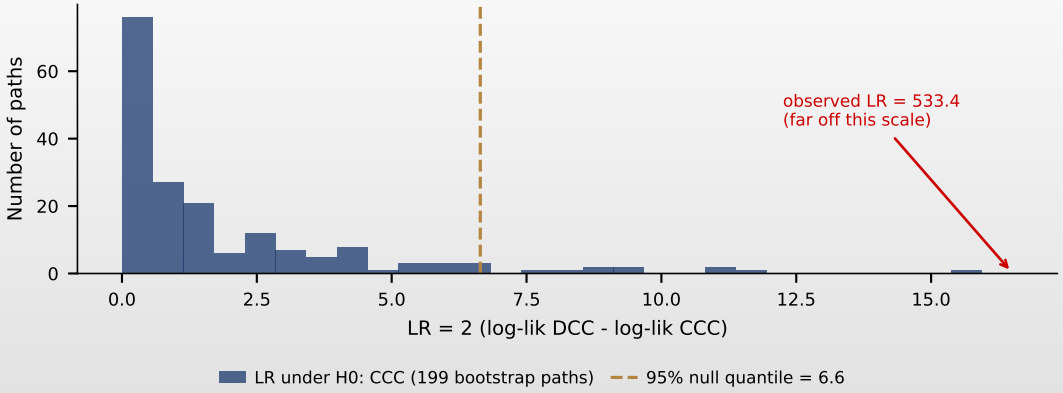


	2002–2019	2020–2026
a	0.0565	0.0341
b	0.9199	0.9527
$\bar{Q}_{12}$	−0.350	0.003
T	4 386	1 687

- Full sample: $\rho_{\mathbf{12},t}$ from -0.79 to 0.49; above the target on 92% of days since 2022
- Persistence is similar in both halves, but the target moves from about -0.35 to about 0
- Interpretation: a full-sample DCC pulls today's correlation towards a target that no longer applies

MFM_ch6_dcc_spy_tlt

B1 [Solved]: Bootstrap Distribution of LR



MFM_ch6_seminar

B1 [Solved]: Bootstrap Test of Constant Correlation

- $LR = 2(\ell_C^{DCC} - \ell_C^{CCC}) = 2(460.2593 - 193.5447) = 533.4$
- χ^2_2 is not the null law: under $a = 0$, b is not identified; the bootstrap under CCC

1. decorrelate ε_t with the constant correlation matrix and resample the pairs;
2. rebuild CCC innovations and GARCH returns with the fitted parameters;
3. re-estimate both steps and the LR on each of 199 paths;
4. $p = (1 + \#\{LR^{(b)} \geq LR\})/200$.

- Largest null value 15.7, 95% null quantile 6.6: $p = 0.005$, the smallest possible; reject constant correlation

- This bootstrap simulates under CCC to test dynamics; the both-step bootstrap of the extension simulates under the fitted DCC to measure uncertainty

MFM_ch6_seminar

B1ext [Proposed]: Correlation Uncertainty and EWMA

- **Data**: the same SPY–TLT adjusted-price daily log returns as B1, restricted to 2002–2021; align prices before differencing
- **Tasks**

1. Compute the correlation and three standard errors: Fisher, delta method with no lags ($L = 0$), and HAC (Heteroskedasticity and Autocorrelation Consistent) with a Bartlett kernel and $L = \lfloor 4(T/100)^{2/9} \rfloor$ (delta-method derivation on the moment vector: course appendix).
2. Compute an EWMA (Exponentially Weighted Moving Average) correlation path with decay 0.94: initialise the covariance from the first 100 returns, then forecast each day using only the previous return.
3. Compare the EWMA path with a 252-day rolling correlation.
4. Explain why a larger standard error changes the strength of inference without changing the estimated full-sample correlation.

- **Check values**: standard errors about 0.0120, 0.0217, 0.0227
- **Report**: the sample size, the lag, the table of the three errors, the HAC/Fisher ratio and the plot of both paths; notebook: section B1ext (not "B1 Extended", the break date)

B2 [Solved]: Gaussian or t Copula for the S&P 500 and the Euro Stoxx 50?

- **Question**: do weekly joint crashes of US and euro-area equities need a copula with tail dependence?
- **Data and model**: the pseudo-observations of A2 ($n = 1\,393$); an unconditional i.i.d. weekly working model (no GARCH filter: the copula mixes volatility regimes with dependence); canonical maximum likelihood (CML) on the ranks; the Gaussian copula is the t copula with $\nu = \infty$, a boundary point
- **Tasks**

1. Fit the Gaussian and the t copula by CML and compare them by AIC.
2. Compute the likelihood ratio for $H_0 : \nu = \infty$; its p-value comes from a parametric bootstrap under H_0 (simulate, re-rank, refit both; 499 draws).
3. Run the Rosenblatt goodness-of-fit test (Genest, Rémillard & Beaudoin (2009)) for both copulas using the integral statistic $S_n^{(B)}$ (simulate from the fitted copula, re-rank, refit; 499 draws).
4. Compare the empirical $C_n(q, q)/q$ ($C_n(q, q)$: share of weeks with both pseudo-observations $\leq q$) with the fitted $C_\theta(q, q)/q$ at the same $q = 5\%$, and report the limit λ of the t copula separately.

- **Report**: a table ($\hat{\rho}$, $\hat{\nu}$, ℓ, AIC, fit p), the LR with its p, the finite-threshold chart; notebook: section B2

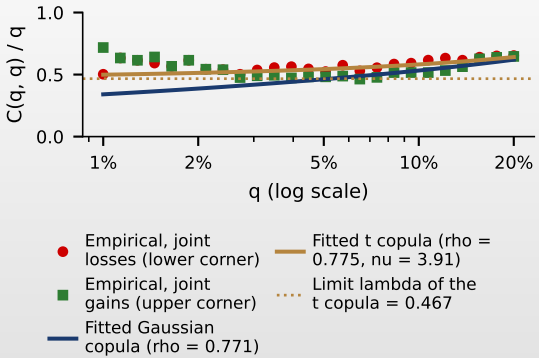
Back Exercise list

B2 [Solved]: Solution (1/2) — Fit and Tests

	$\hat{\rho}$	$\hat{\nu}$	ℓ	AIC	fit p
Gaussian	0.771	–	623.3	−1 244.6	0.002
Student t	0.775	3.91	674.2	−1 344.4	0.874

- (1) The t copula has the lower AIC by 99.8 points: the extra parameter ν is strongly supported
- (2) $LR = 2(674.2 - 623.3) = 101.8$; the 95% quantile of the 499 null draws is 3.2 and no simulated value exceeds 101.8, so $p = 0.002$, the smallest possible (1/500): reject the Gaussian copula
- χ^2_1 would be wrong: $\nu = \infty$ lies on the boundary, and the $\frac{1}{2}\chi^2_0 + \frac{1}{2}\chi^2_1$ mixture is not established for a rank-based pseudo-likelihood
- (3) Fit test: Gaussian $p = 0.002$, rejected at 5%; t $p = 0.874$, not rejected: absence of evidence against the t copula, not proof

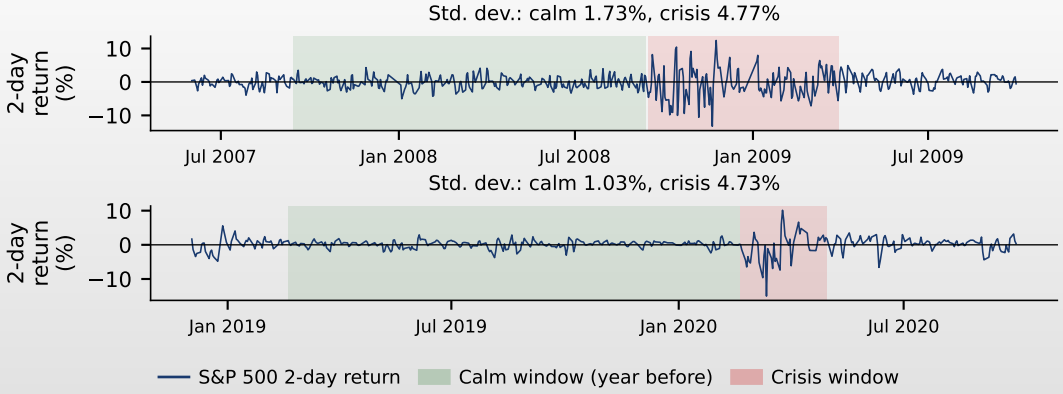
B2 [Solved]: Solution (2/2) — The Tails at a Finite Threshold



- (4) At $q = 5\%$: empirical lower 0.517 (36 weeks in the corner), upper 0.488 (34 weeks)
- Fitted at the same q: Gaussian 0.462, t 0.544; the limit of the t copula is $\lambda = 0.467$, a different object
- The t curve tracks the empirical points down to small q; the Gaussian curve falls away, as A3 predicts
- Interpretation: the unconditional weekly model shows joint extremes; B4 illustrates filtering on a different daily pair; survival after filtering is not tested for this weekly pair; the empirical upper corner at small q rests on a handful of weeks

MFM_ch6_seminar

B3 [Proposed]: Contagion in 2008 and 2020? — Data and Windows



MFM_ch6_seminar

B3 [Proposed]: Contagion in 2008 and 2020? — Test Conventions

- **Questions**
  - did the links between the S&P 500 and the Euro Stoxx 50 or BET strengthen in the two crises?
  - do they remain stronger once the higher US volatility is accounted for?
- **Data**: 2-day log returns $\ln P_t - \ln P_{t-\mathbf{2}}$ (P_t: closing level on day t) on common trading days; the S&P 500 is the source; a return belongs to the window of its end date

	calm	crisis
2008	14 Sep 2007 – 12 Sep 2008	15 Sep 2008 – 31 Mar 2009
2020	19 Feb 2019 – 19 Feb 2020	20 Feb – 30 Apr 2020
2008: n (2-day)	242	130
2020: n (2-day)	241	47

MFM_ch6_seminar

Back Exercise list

B3 [Proposed]: Contagion in 2008 and 2020? — In-Class Tasks

- **Worked example to follow**: A1; **Tasks in class**

1. Work 2008 with BET fully: compute ρ_{calm}, ρ_c, δ and ρ^*, then repeat for the other three cases.
2. As a historical comparison, run the one-sided Fisher z tests (slide «Comparing Two Correlations: the Fisher z Test») of $H_0 : \rho_c \leq \rho_{calm}$ and $H_0 : \rho^* \leq \rho_{calm}$, with n the floor of half the overlapping 2-day return count, a heuristic correction. The 2020 crisis has only 23 effective pairs by this rule: treat Fisher as indicative and compare with the bootstrap.

- **Report**: a table by crisis and market (δ, ρ_{calm}, ρ_c, p_{raw}, ρ^*, p_{adj}); notebook: section B3
- **Interpretation**
  - why is δ so much larger in 2020?
  - which assumption is least credible for a pandemic?
- Tasks 3–4 (bootstrap with an estimated δ, shorter windows) come after the lecture, in the homework section

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

- **Prompt** sent to an AI assistant: "The pandas table e holds the GARCH-standardised residuals of JPM (JPMorgan Chase) and BAC (Bank of America), two US bank stocks. Calibrate a t copula with $\nu = 4$ from Kendall's τ and give me its tail dependence."
- **Answer**: the code below and the conclusion

```
import numpy as np
from scipy import stats
tau = stats.kendalltau(e['JPM'], e['BAC'])[0]        # Kendall's tau
rho = 2 * np.sin(np.pi * tau / 6)                    # copula correlation
u = e.rank() / len(e)                                # pseudo-observations
z = stats.t.ppf(u, df=4)                             # input of the t-copula likelihood
nu = 4
x = np.sqrt((nu + 1) * (1 - rho) / (1 + rho))
lam = 2 * stats.t.cdf(-x, df=nu)                     # tail dependence
```

- **Conclusion of the AI**: "With $\hat{\tau} = 0.50$: $\rho = 0.518$ and $\lambda_L = \lambda_U = 0.276$."

C1 [Proposed]: Critique an AI Answer

- **Worked example to follow**: A2 and A3 [Solved]; **inputs**: the code above with $\hat{\tau} = 0.50$ and $\nu = 4$; for the data check, the JPM (JPMorgan Chase & Co., NYSE ticker)–BAC residuals of the notebook (section C1)
- The answer contains **three** mistakes typical of copula modelling; everything else is correct
- **Tasks**

1. Find the three wrong lines and explain in one sentence why each is wrong.
2. For each, give the evidence that exposes it: a formula from A2 or A3, a check of the largest pseudo-observation, or a small simulation.
3. Write the three corrected lines and recompute ρ and λ for $\hat{\tau} = 0.50$, $\nu = 4$.
4. Say which errors change the displayed λ and which one only breaks the likelihood inputs.

- **Report**: a three-row table (error, evidence, correction) and the corrected ρ and λ; log in `AI_ERRORS.md` the errors you actually found

B3 [Proposed]: Contagion in 2008 and 2020? — Tasks 3–4, after the Lecture

Solution discussed in class

- **Worked example to follow**: A1 and your in-class table of B3; **Tasks**

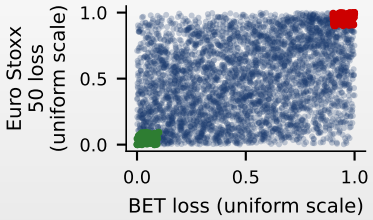
3. Test $D = \rho^* - \rho_{calm}$ by bootstrap with δ estimated in each draw:
   - resample the joint vectors (S&P 500, target) within each window with a stationary bootstrap Politis & Romano (1994), mean block 10 days, 1 999 draws;
   - recompute $\hat{\delta}$, ρ^* and ρ_{calm}; report the 90% interval of D and $p =$ the share of draws with $D \leq 0$.
4. Sensitivity: end the crisis windows on 31 December 2008 and on 31 March 2020.

- **Report**: the in-class table with a column for the 90% interval of D, and one chart; notebook: section B3, the bootstrap cell
- **Interpretation**: does estimating δ change any verdict of task 2?

B4 [Proposed]: Joint Losses: BET and Euro Stoxx 50

- **Worked example to follow**: B2 and A2; **data**: daily BET and Euro Stoxx 50 closes from 5 January 2010, aligned on common days (almost synchronous closes), $T = 4\,154$; notebook: section B4
- Filter each series with constant-mean GARCH(1,1), Gaussian QML as in B1; losses $L = -\varepsilon$ (minus the standardised returns) and their pseudo-observations: joint crashes sit in the *upper-right* corner
- **Tasks**

1. **What do you think**: which family puts its tail dependence in the corner of joint crashes, Clayton or Gumbel on losses?
2. Fit Clayton and Gumbel by CML (Gaussian and t as benchmarks); report AIC, λ_L, λ_U and the Rosenblatt fit p (200 draws; 100 for t).
3. Compute the empirical probabilities in both corners (joint large losses and joint large gains) at $q = 5\%$ and 10%, and compare them with the fitted curves.
4. Explain what the 180° rotated (survival) Clayton on losses represents.



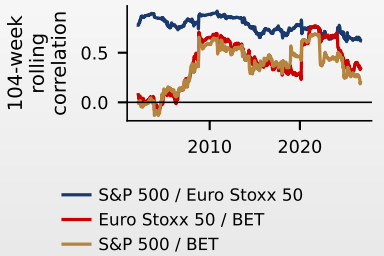
MFM_ch6_seminar

Back Exercise list

C2 [Proposed]: Did Romanian and Euro-Area Equities Become More Integrated after 2020?

- **Worked example to follow**: A1, B1; **data**: weekly BET and Euro Stoxx 50 returns on common days, 2010–2019 ($n = 520$) against 2020–2026 ($n = 351$); notebook: section C2
- 2020 brought COVID-19 and the FTSE Russell (Financial Times Stock Exchange) upgrade of Romania (FTSE Russell (2019)): a descriptive period comparison, not a causal test of the upgrade
- **Tasks**

1. Estimate the change in weekly correlation with a 95% stationary-bootstrap interval (Politis & Romano (1994), resample within each period, mean block 8 weeks, 2 000 draws) and the one-sided p for H_0: no increase, the share of draws ≤ 0.
2. Run two separate checks against the same baseline: the Forbes–Rigobon adjustment for the change in Euro Stoxx 50 volatility, and the correlation without 20 February – 30 April 2020.
3. Optional: Kendall's τ, the beta $\mathrm{Cov}(\mathrm{BET}, \mathrm{Stoxx})/\mathrm{Var}(\mathrm{Stoxx})$, the $q = 10\%$ lower-corner probability and a weekly DCC.



- Context from the lecture: 104-week rolling correlations
- **Report**: the table of measures by period, the verdict and its caveats

MFM_ch6_integration

Back Exercise list

Extension Toolkit: Standard Errors

- **Delta method**: the standard error of a smooth function g of an estimate

$$\operatorname{Var}\big(g(\hat\theta)\big) \approx g'(\theta)^{\mathbf{2}} \operatorname{Var}(\hat\theta), \qquad \operatorname{Var}\big(g(\hat\theta_{\mathbf{1}}, \hat\theta_{\mathbf{2}})\big) \approx \nabla g^{\top}\, \Sigma_{\hat\theta}\, \nabla g$$

  - g': derivative; ∇g: gradient; $\Sigma_{\hat\theta}$: covariance matrix of the estimates; e.g. Clayton $\theta = 2\tau/(1-\tau)$ has $g'(\tau) = 2/(1-\tau)^{\mathbf{2}}$

- **Standard error of Kendall's** $\hat\tau$, a U-statistic (an average over all pairs of observations)

$$\operatorname{se}(\hat\tau) \approx \frac{4\,\mathrm{sd}(W)}{\sqrt{n}}, \qquad W_i = 2\,C_n(\hat u_i, \hat v_i) - \hat u_i - \hat v_i$$

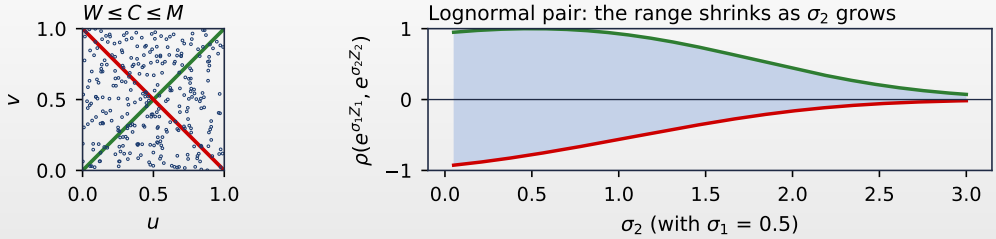
  - W_i: the Hoeffding projection, the contribution of observation i to $\hat\tau$; sd: sample standard deviation

- **Sandwich** covariance of a CML estimate

$$\widehat{\operatorname{Var}}(\hat\theta) = \frac{1}{n}\, A^{-\mathbf{1}}\, \Sigma\, A^{-\mathbf{1}}$$

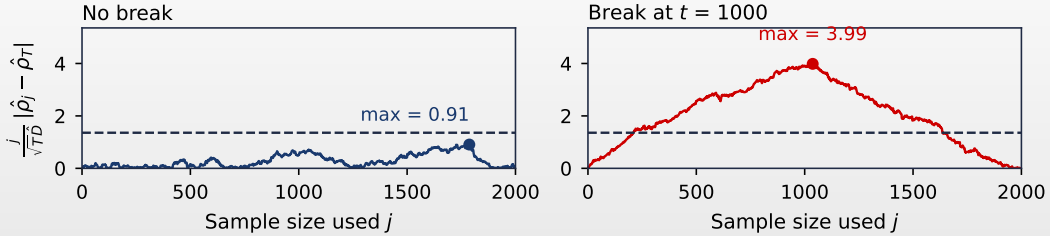
  - A: Fisher information (minus the expected second derivative of $\ln c_\theta$); Σ: variance of the score plus the rank corrections; $A^{-\mathbf{1}}/n$ alone ignores that the ranks are estimated

Extension Toolkit: Attainable Correlations



- **Hoeffding's identity**: the covariance is the integrated gap between the joint distribution and independence
  - $\mathrm{Cov}(X, Y) = \iint \left[F(x, y) - F_1(x)F_2(y)\right] dx\, dy$; with $F = C(F_1, F_2)$ and $W \le C \le M$, the margins fix the range of ρ
  - Right: two lognormals $e^{\sigma_1 Z_1}$, $e^{\sigma_2 Z_2}$, $Z_i \sim N(0,1)$, $\mathbb{E}(e^{aZ}) = e^{a^2/2}$; the range of Pearson correlations shrinks as the margins differ more

Extension Toolkit: Testing a Constant Correlation



- **Fluctuation test** Wied, Krämer & Dehling (2012): compare the correlation of the first j pairs with the full-sample one, for every j
  - $Q_T = \max_j \frac{j}{\sqrt{T}\hat{D}}|\hat{\rho}_j - \hat{\rho}_T|$; $\hat{D}$: long-run standard deviation of the correlation (delta method with the gradient of the five moments, Bartlett HAC, width $\lfloor T^{\mathbf{1/4}} \rfloor$); the j of the maximum dates the break
  - Under H_0 (constant correlation) Q_T tends to the supremum of |Brownian bridge|; 5% critical value 1.358; chart: simulated, without and with a break

Extension Toolkit: Dynamic Models

- **Corrected DCC** Aielli (2013): the shock is rescaled by $Q^{*\,1/2}_{t-1}$, so that $\mathbb{E}[Q_t] = \bar{Q}^*$ holds; the starred fixed matrix is the target of the corrected recursion (normalised to unit diagonal), not the time-varying diagonal

$$Q_t = (1 - a - b)\,\bar{Q}^* + a\,Q^{*\,1/2}_{t-1}\varepsilon_{t-1}\varepsilon_{t-1}^{\top}Q^{*\,1/2}_{t-1} + b\,Q_{t-1}$$

- **Decorrelation** (B5): $\eta_t = L_t^{-1}\varepsilon_t$, with L_t the Cholesky factor of R_t ($R_t = L_t L_t^{\top}$, L_t lower triangular): η_t has identity covariance and can be resampled
- **Generalised Autoregressive Score (GAS) t copula** Creal, Koopman & Lucas (2013): the correlation moves with the score of yesterday's observation

$$\rho_t = \tanh f_t, \qquad f_{t+1} = \omega + A\,s_t + \beta_{GAS} f_t, \qquad s_t = (1 - \rho_t^2)\,\frac{\partial \ln c(\hat{u}_t, \hat{v}_t; \rho_t, \nu)}{\partial \rho}$$

  - tanh keeps ρ_t in $(-1, 1)$; A: reaction to the score s_t; β_{GAS}: persistence; $A = 0$ gives the static copula
- **Common volatility** (COVOL, C3) Engle & Campos-Martins (2023): squared GARCH residuals share a global factor x_t

$$e_{it}^2 = (s_i\,x_t + 1 - s_i)\,\varepsilon_{it}^2$$

  - e_{it}: standardised GARCH residual of series i; $s_i \in [0, 1]$: loading of series i on the common factor; ε_{it}: idiosyncratic shock with variance 1

Extension Packs: What Each One Needs

Pack	Exercises	Prerequisites
1. Copula inference	A2 Extended [Solved], B2 Extended [Solved]	A2, B2
2. Contagion inference	A4 [Proposed], A1 Extended [Solved]	A1, B3
3. DCC dynamics	A5 [Proposed], B1 Extended [Solved], B5 [Proposed]	B1
4. Margins and dynamic copulas	A6 [Proposed], B6 [Proposed]	A2, B2
5. Project brief	C3 [Proposed]: global common volatility	B1, C2, lecture case study

- Each pack has its own notebook section
- C3 is a project, not a seminar exercise: its reference outputs are partial, so it suits a team that wants a replication

A2 Extended [Solved]: How Precise Are Kendall's τ and the Copula Parameters?

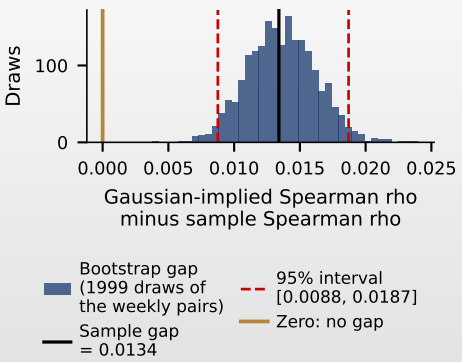
- **Inputs**: the pseudo-observations of A2, $n = 1\,393$, $\hat{\tau} = 0.5642$; the empirical copula C_n (slide «Pseudo-Observations and the Empirical Copula»)
- **Tasks**

1. Using the Hoeffding projection W_i (slide «Extension Toolkit: Standard Errors»; its sample standard deviation is 0.1239), show that $\mathrm{se}(\hat{\tau}) = 4\,\mathrm{sd}(W)/\sqrt{n}$ and compute it.
2. Propagate it to $\hat{\theta}_{Clayton}$ and $\hat{\rho}_{Gauss}$ by the delta method.
3. Bootstrap the weekly pairs jointly (1 999 draws), recompute both rank statistics in each draw and report the 95% percentile interval of the gap $g(\hat{\tau}) - \hat{\rho}_S$ of A2.

- **Report**: three standard errors, the interval and a decision at 5%; the i.i.d. bootstrap assumes no serial dependence in the weekly pairs

A2 Extended [Solved]: Solution

- (1) $\hat{\tau}$ is a U-statistic with $E[h \mid Z_i] = 2W_i + 1$, so $\mathrm{Var}(\hat{\tau}) \approx 16\mathrm{Var}(W)/n$: $\mathrm{se} = 4 \times 0.1239/\sqrt{1\,393} = 0.0133$
- (2) Clayton: $d\theta/d\tau = 2/(1-\tau)^2 = 10.53$, se 0.140; Gaussian: $\frac{\pi}{2}\cos(\pi\tau/2) = 0.993$, se 0.0132 (Genest & Rivest (1993))
- (3) Gap 0.0134, se 0.0026, 95% interval [0.0088, 0.0187]: zero is outside, so the gap is significant at 5%
- Why so precise: $\hat{\tau}$ and $\hat{\rho}_S$ move together across samples, so their difference has a much smaller se than each
- Interpretation: the pair $(\hat{\tau}, \hat{\rho}_S)$ already rejects the Gaussian copula, consistently with B2



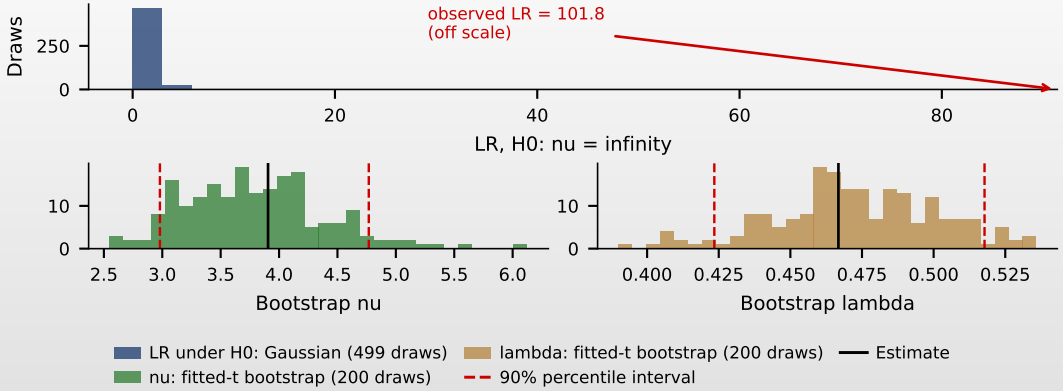
MFM_ch6_seminar

B2 Extended [Solved]: Bootstrap Distributions and Sandwich Standard Errors

- **Inputs**: the B2 fits; the ranks are estimated, so the inverse Hessian understates the uncertainty of CML estimates
- **Tasks**

1. Plot the null distribution of the LR of B2 against the observed value.
2. Build 90% intervals for ν and λ from 200 draws of the fitted t copula, re-ranked and refitted.
3. Compute the se of $\hat{\rho}$ and $\hat{\nu}$ from the inverse Hessian A^{-1}/n and from the Genest, Ghoudi & Rivest (1995) sandwich $A^{-1}\Sigma A^{-1}/n$, $\Sigma = \mathrm{Var}[\dot{\ell}_\theta + W_1(U) + W_2(V)]$; A: Fisher information of the copula, $\dot{\ell}_\theta$: score of $\ln c_\theta$, W_1, W_2: corrections for the margins estimated by ranks.

B2 Extended [Solved]: Solution (1/2)



MFM_ch6_seminar

B2 Extended [Solved]: Solution (2/2)

- (1) 499 draws under H_0: the 95% quantile is 3.2; no simulated value reaches the observed 101.8
- (2) $\hat{\nu} = 3.91$, 90% interval [2.98, 4.77]; $\hat{\lambda} = 0.467$, [0.423, 0.518]
- (3) Gaussian: se($\hat{\rho}$) 0.0086 from the inverse Hessian, 0.0138 from the sandwich
- t, same order: se($\hat{\rho}$) 0.0115 versus 0.0131; se($\hat{\nu}$) 0.56 versus 0.58
- Interpretation: ranking adds variance to $\hat{\rho}$; ν is driven by the tails, where ranking errors matter little

MFM_ch6_seminar

A4 [Proposed]: Deriving the Forbes–Rigobon Bias

- **Worked example to follow**: A1; **model**: $y_t = \alpha + \beta x_t + \varepsilon_t$ with $\mathrm{Cov}(x_t, \varepsilon_t) = 0$ maintained; β and σ_ε^2 equal in calm and crisis; only $\sigma_{x,c}^2 = (1+\delta)\sigma_x^2$ changes
- **Tasks**

1. Write ρ^2 as a function of $k = \beta^2\sigma_x^2/\sigma_\varepsilon^2$, using $\mathrm{Cov}(x, y)$ and $\mathrm{Var}(y)$.
2. Derive $\rho_c = \rho\sqrt{(1+\delta)/(1+\delta\rho^2)}$ and show that $|\rho_c| > |\rho|$ for $0 < |\rho| < 1$.
3. Invert the relation to obtain $\rho^* = \rho_c/\sqrt{1+\delta(1-\rho_c^2)}$.
4. Suppose that in the crisis σ_ε^2 also rises, to $(1+\gamma)\sigma_\varepsilon^2$ with $0 < \gamma < \delta$: compute the true crisis correlation and, using the calm correlation as the benchmark, say in which direction ρ^* is then biased.

- **Report**: each step with one sentence of justification and the answer to task 4 in economic terms

Back Exercise list

A1 Extended [Solved]: The Forbes–Rigobon Adjustment with an Estimated δ

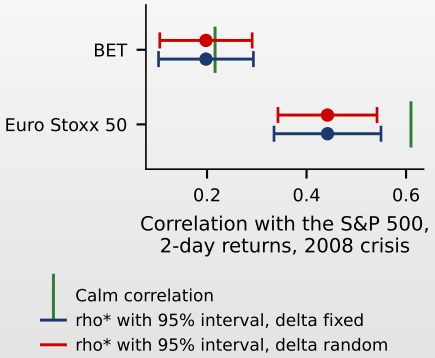
- **Inputs** (2008 windows of B3): $\hat{\delta} = 6.54$; effective $n_c = 65$, $n_0 = 121$, the floor of half the overlapping 2-day return count (a heuristic, not an exact effective sample size)
- **Tasks**

1. Under i.i.d. Normality, $(1+\hat{\delta})/(1+\delta) \sim F(n_c - 1, n_0 - 1)$, the Fisher F distribution: derive $\mathrm{Var}(\hat{\delta})$ to first order.
2. Write the gradient (g_r, g_δ) of ρ^* in $(\hat{\rho}_c, \hat{\delta})$ and their covariance (same crisis window).
3. Compute $\mathrm{se}(\rho^*)$ with δ fixed and random for both markets, and the z of $\rho^* - \hat{\rho}_{calm}$ including the uncertainty of $\hat{\rho}_{calm}$.

- These Normal, effective-sample formulas are illustrative; the stationary bootstrap of B3 is the primary inference
- **Report**: $\mathrm{se}(\hat{\delta})$, the two se of ρ^*, their percentage difference, the two z values

A1 Extended [Solved]: Solution

- (1) $\operatorname{Var}\ln(1+\hat{\delta}) \approx \frac{2}{n_c - 1} + \frac{2}{n_0 - 1}$, so $\operatorname{se}(\hat{\delta}) = 7.54\sqrt{2(1/64 + 1/120)} = 1.65$
- (2) $k = 1 + \delta(1-\rho_c^2)$, $g_r = (1+\delta)k^{-3/2}$, $g_\delta = -\rho_c(1-\rho_c^2)k^{-3/2}/2$; $\operatorname{Cov}(\hat{\rho}_c, \hat{\delta}) \approx (1+\delta)\rho_c(1-\rho_c^2)/n_c$
- $\operatorname{Var}(\rho^*) \approx g_r^2 \operatorname{Var}\hat{\rho}_c + g_\delta^2 \operatorname{Var}\hat{\delta} + 2 g_r g_\delta \operatorname{Cov}(\hat{\rho}_c, \hat{\delta})$
- (3) Euro Stoxx 50: se $0.055 \to 0.051$ (-7%), $z = -2.44$; BET: $0.049 \to 0.047$ (-3%), $z = -0.19$
- A random δ need not widen the interval: a volatile crisis window raises $\hat{\delta}$ and $\hat{\rho}_c$ together, and their errors partly cancel; still no evidence of an increase



MFM_ch6_seminar

Why the Adjusted Difference Needs a Covariance

- Let $r_0 = \hat{\rho}_{calm}$. Although the two windows are independent, ρ^* and r_0 share the calm-window data through $\hat{\delta}$.
- Under the i.i.d. Normal approximation, $s_0 = (1 - r_0^2)/\sqrt{n_0}$ and $c_{*0} = \mathrm{Cov}(\rho^*, r_0) \approx -g_\delta(1+\delta)r_0(1-r_0^2)/n_0$.
- The statistic reported above uses $z_{random} = (\rho^* - r_0)/\sqrt{s_{random}^2 + s_0^2 - 2c_{*0}}$, where s_{random} is the standard error from the previous page.
- With fixed δ, use $s_{fixed}^2 + s_0^2$ in the denominator variance. This test on the correlation scale differs from the earlier Fisher transformation.

A5 [Proposed]: DCC Algebra and a Shock

- **Worked example to follow**: B1; target $\bar{Q} = \begin{pmatrix} 1 & -0.248487 \\ -0.248487 & 1 \end{pmatrix}$, $a = 0.052232$, $b = 0.933518$, $Q_{t-1} = \bar{Q}$, a shock $\varepsilon_{t-1} = (-2.0,\ 1.0)$; notation: $Q_t^* = \operatorname{diag}(Q_t)$, a time-varying matrix, not the fixed target
- **Tasks**

1. Show by induction that Q_t is positive definite for every t when $Q_0 = \bar{Q}$ is.
2. Compute Q_t and $R_{t,12}$ after the shock, and the heuristic persistence measure $\ln 0.5 / \ln(a + b)$.
3. Show that $\mathbb{E}[Q_t] = \bar{Q}$ requires $\mathbb{E}[R_t] = \bar{Q}$, which fails in general; for the corrected DCC of Aielli (2013) (slide «Extension Toolkit: Dynamic Models»), show $\mathbb{E}[Q_t] = \bar{Q}^*$.

- **Report**: the proofs, $R_{t,12}$ and the persistence measure; is it a rule for how often to rebalance a hedge?

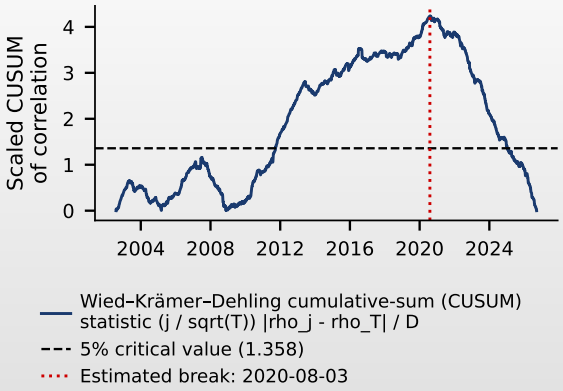
Back Exercise list

B1 Extended [Solved]: Break Date and DCC Precision

- **Questions**
  - when did the correlation break?
  - how precise is DCC?
- **Inputs**: the standardised residuals and the DCC fit of B1
- **Tasks**

1. Compute the Wied, Krämer & Dehling (2012) statistic Q_T (slide «Extension Toolkit: Testing a Constant Correlation»), test at 5% (1.358) and date the break.
2. Explain why a Fisher test at January 2022, a date read off the lecture chart, is not valid.
3. Simulate 199 paths under the fitted DCC (resampled decorrelated residuals); compute the se of (a, b) re-estimating step 2 only and both steps (Newey & McFadden (1994)).

B1 Extended [Solved]: Solution (1/2)



MFM_ch6_corr_inference

B1 Extended [Solved]: Solution (2/2)

se of	$\hat{a}$	$\hat{b}$
Hessian	0.0059	0.0084
step 2 only	0.0051	0.0077
both steps	0.0052	0.0077

- (1) $Q_T = 4.25 > 1.358$ ($p < 0.001$); break 3 August 2020; residual correlation -0.361 before, 0.066 after; a second break on 9 July 2007
- The limit law is proved for observed series; on estimated GARCH residuals, the p-value is only approximate
- (2) Fisher at January 2022 gives $z = 14.9$, but the date was chosen after seeing the data and Fisher assumes i.i.d. Normal pairs
- (3) In this bootstrap experiment step 1 adds 2% and 1% to the dispersion, below the Monte Carlo error of about 5% of a bootstrap standard deviation with $B = 199$; the result is not general: the step-1 contribution can stay first-order as T grows
- The Hessian se($\hat{b}$) exceeds the bootstrap standard deviation, which does not prove conservative coverage

MFM_ch6_corr_inference

B5 [Proposed]: Bitcoin and the S&P 500: DCC with Bootstrap Bands

- **Worked example to follow**: B1; **data**: Bitcoin (USD) and the S&P 500 index (the index, not SPY: the lecture's SPY result is a separate benchmark), common trading days since 18 September 2014, $T = 3\,018$; weekend Bitcoin moves enter the Monday return
- **Tasks**

1. Fit constant-mean GARCH(1,1) by Gaussian QML and DCC(1,1); compute the average conditional correlation before and after 1 January 2020.
2. Bootstrap the parameter uncertainty, 200 draws:
   - keep the GARCH residuals and the target $\bar{Q}$ fixed; decorrelate $\eta_t = L_t^{-1}\varepsilon_t$ (slide «Extension Toolkit: Dynamic Models»); resample the η_t jointly;
   - simulate a DCC with $(\hat{a}, \hat{b})$, re-estimate (a, b), then refilter the *observed* residuals with each bootstrap (a, b).
3. Draw the 90% pointwise band of ρ_t and a 90% interval for the change in the mean correlation after 2020.

- **Interpretation**
  - which uncertainty does the band ignore?
  - why is it not a prediction interval?

Back Exercise list

A6 [Proposed]: Attainable Correlations

- **Worked example to follow**: A2; $X = e^{\sigma_1 Z_1}$, $Y = e^{\sigma_2 Z_2}$, $Z_i \sim N(0,1)$: hypothetical lognormal variables with fixed margins and free dependence
- **Tools**: the Fréchet–Hoeffding bounds $W \leq C \leq M$ (slide «Copulas and Sklar's Theorem»); the Hoeffding identity and $\mathbb{E}(e^{aZ}) = e^{a^2/2}$ (slide «Extension Toolkit: Attainable Correlations»)
- **Tasks**

1. Show that $\rho(X, Y)$ is largest for $Z_2 = Z_1$ and smallest for $Z_2 = -Z_1$.
2. Compute in turn $\mathbb{E}X$, $\mathbb{E}Y$, $\mathbb{E}XY$, $\mathrm{Var}X$, $\mathrm{Var}Y$ and derive $\rho_{\max}$ and $\rho_{\min}$ in closed form.
3. Evaluate them for $\sigma_1 = 1$, $\sigma_2 = 2$ and interpret $\rho = 0.3$ relative to this interval.

Back Exercise list

B6 [Proposed]: A Dynamic t Copula for the S&P 500 and the Euro Stoxx 50

- **Worked example to follow**: B2 and the lecture's dynamic t copula; data: the weekly pseudo-observations of A2 ($n = 1\,393$); benchmark: the static t copula of B2
- The weekly ranks are not GARCH-filtered: the dynamic correlation can also absorb common volatility regimes. Compare with filtered residuals before interpreting it as changing dependence.
- GAS t copula (Creal, Koopman & Lucas (2013), unit scaling) (slide «Extension Toolkit: Dynamic Models»), started at $f_1 = \omega/(1-\beta_{GAS})$; with x, y the t_ν quantiles and $Q = x^2 + y^2 - 2\rho xy$:
  $\ln c = \text{const} - \frac{1}{2}\ln(1-\rho^2) - \frac{\nu+2}{2}\ln\left(1 + \frac{Q}{\nu(1-\rho^2)}\right) + \frac{\nu+1}{2}\sum \ln(1 + x^2/\nu)$
- **Tasks**

1. Derive $s_t = (1-\rho_t^2)\,\partial \ln c/\partial\rho$ by the chain rule and check it numerically against a finite difference.
2. Estimate $(\omega, A, \beta_{GAS}, \nu)$ under $0 \le A \le 1$, $0 \le \beta_{GAS} < 1$, $\nu > 2$, starting from the static fit.
3. Compute the LR against the static t copula and its p by a bootstrap under the static copula ($B_{boot} = 199$, simulate, re-rank, refit both); plot ρ_t against the static ρ.

Back Exercise list

Which Funds Enter the Global Volatility Project?

- Sector funds: XLB (Materials Select Sector SPDR Fund) materials; XLC communication services; XLE (Energy Select Sector SPDR Fund) energy; XLF (Financial Select Sector SPDR Fund) financials; XLI (Industrial Select Sector SPDR Fund) industrials; XLK (Technology Select Sector SPDR Fund) technology.
- Other sectors: XLP consumer staples; XLRE real estate; XLU (Utilities Select Sector SPDR Fund) utilities; XLV (Health Care Select Sector SPDR Fund) health care; XLY consumer discretionary.
- Styles and geography: IWF (iShares Russell 1000 Growth ETF) Russell 1000 Growth; IWD (iShares Russell 1000 Value ETF) Russell 1000 Value; IWM (iShares Russell 2000 ETF) Russell 2000 small-cap stocks; EFA (iShares MSCI EAFE ETF) developed markets outside the US and Canada; EEM (iShares MSCI Emerging Markets ETF) emerging markets.
- Other assets: GLD (SPDR Gold Shares) gold; TLT long-term US Treasuries; LQD (iShares iBoxx Investment Grade Corporate Bond ETF) investment-grade corporate bonds; USO (United States Oil Fund) oil futures exposure.

C3 [Proposed]: Global Common Volatility, a Project Brief

- **Worked example to follow**: B1, C2 and the lecture case study (Engle & Campos-Martins (2023)); the COVOL model (slide «Extension Toolkit: Dynamic Models»), x_t the squared global COVOL
- **Data**: adjusted closes of 20 US-listed ETFs from the course market data (11 sector ETFs; Russell 1000 Growth and Value; Russell 2000; developed and emerging markets outside the US; gold; long Treasuries; corporate bonds; crude oil): XLB, XLC, XLE, XLF, XLI, XLK, XLP, XLRE, XLU, XLV, XLY, IWF, IWD, IWM, EFA, EEM, GLD, TLT, LQD, USO, each from its first day to 18 September 2026 (unbalanced panel); the VIX index, daily closes
- **Tasks**

1. Filter: daily log returns; autoregressive AR(1) mean with the first Principal Component (PC1) of returns as factor (state how you form PC1 on the unbalanced panel); GARCH(1,1); keep the standardised residuals of the ETFs and of PC1.
2. Equal loadings: $\hat{x}_t = \frac{\mathbf{1}}{N} \sum_i e_{it}^{\mathbf{2}}$ over the series observed that day; list its 20 largest days before and after 1 March 2021 with mean return and an event label.
3. Compare the average correlation of the residuals with that of their squares; then, as a team project, estimate unequal loadings by the alternating algorithm (Section 5.3 of the paper), the test ξ and the VIX regression.

- **Report**: the two top-20 tables, the correlation comparison and, for the project, loadings, ξ and the regression with standard errors; interpretation: why can a day with a positive mean return rank high?

Back Exercise list

Acronyms used in this material

- **ACF**: AutoCorrelation Function
- **AI**: Artificial Intelligence
- **AIC**: Akaike Information Criterion
- **AR**: Autoregressive model
- **BAC**: Bank of America Corporation (NYSE ticker)
- **BET**: Bucharest Exchange Trading (index)
- **CBOE**: Chicago Board Options Exchange
- **CCC**: Constant Conditional Correlation (model)
- **CML**: Canonical Maximum Likelihood
- **COVID-19**: COronaVIrus Disease 2019
- **COVOL**: COmmon VOLatility (global COVOL: the size of the volatility shock common to all assets)
- **DCC**: Dynamic Conditional Correlation
- **ETF**: Exchange-Traded Fund
- **EWMA**: Exponentially Weighted Moving Average
- **FTSE**: Financial Times Stock Exchange (FTSE Russell, index provider)
- **GARCH**: Generalised AutoRegressive Conditional Heteroskedasticity
- **GAS**: Generalised Autoregressive Score (model)
- **HAC**: Heteroskedasticity and Autocorrelation Consistent (standard errors)
- **JPM**: JPMorgan Chase & Co. (NYSE ticker)
- **LQD**: iShares iBoxx Investment Grade Corporate Bond ETF
- **LR**: Likelihood Ratio (test)
- **NYSE**: New York Stock Exchange
- **PC1**: first Principal Component
- **QML**: Quasi-Maximum Likelihood
- **US**: United States
- **USO**: United States Oil Fund (ETF ticker)
- **VIX**: CBOE Volatility IndeX

References (1/2)

- Aielli, G.P. (2013). Dynamic conditional correlation: on properties and estimation. *Journal of Business & Economic Statistics*, 31(3), 282–299.
- Creal, D., Koopman, S.J., Lucas, A. (2013). Generalized autoregressive score models with applications. *Journal of Applied Econometrics*, 28(5), 777–795.
- Demarta, S., McNeil, A.J. (2005). The t copula and related copulas. *International Statistical Review*, 73(1), 111–129.
- Engle, R.F. (2002). Dynamic conditional correlation: a simple class of multivariate generalized autoregressive conditional heteroskedasticity models. *Journal of Business & Economic Statistics*, 20(3), 339–350.
- Engle, R.F., Campos-Martins, S. (2023). What are the events that shake our world? Measuring and hedging global COVOL. *Journal of Financial Economics*, 147(1), 221–242.
- Forbes, K.J., Rigobon, R. (2002). No contagion, only interdependence: measuring stock market comovements. *Journal of Finance*, 57(5), 2223–2261.
- Franke, J., Härdle, W.K., Hafner, C.M. (2019). Statistics of Financial Markets: An Introduction, 5th ed. Springer.
- FTSE Russell (2019). FTSE Russell announces results of country classification review (press release, 26 September 2019).

References (2/2)

- Genest, C., Ghoudi, K., Rivest, L.-P. (1995). A semiparametric estimation procedure of dependence parameters in multivariate families of distributions. *Biometrika*, 82(3), 543–552.
- Genest, C., Rivest, L.-P. (1993). Statistical inference procedures for bivariate Archimedean copulas. *Journal of the American Statistical Association*, 88(423), 1034–1043.
- Genest, C., Rémillard, B., Beaudoin, D. (2009). Goodness-of-fit tests for copulas: a review and a power study. *Insurance: Mathematics and Economics*, 44(2), 199–213.
- Nelsen, R.B. (2006). An Introduction to Copulas, 2nd ed. Springer.
- Newey, W.K., McFadden, D. (1994). Large sample estimation and hypothesis testing. In R.F. Engle, D.L. McFadden (eds.), *Handbook of Econometrics*, Vol. 4, 2111–2245. Elsevier.
- Politis, D.N., Romano, J.P. (1994). The stationary bootstrap. *Journal of the American Statistical Association*, 89(428), 1303–1313.
- Wied, D., Krämer, W., Dehling, H. (2012). Testing for a change in correlation at an unknown point in time using an extended functional delta method. *Econometric Theory*, 28(3), 570–589.