

Modelling Financial Markets

Chapter 6: Multivariate Volatility and Dependence



Daniel Traian PELE

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

When Do Markets Move Together?

> *"This paper shows that correlation coefficients are conditional on market volatility. Under certain assumptions, it is possible to adjust for this bias."*
>
> — *Forbes & Rigobon (2002)*

- A 60/40 portfolio of two exchange-traded funds (ETFs): SPY (SPDR S&P 500 ETF Trust), on the S&P 500 (Standard & Poor's 500) index, and TLT (iShares 20+ Year Treasury Bond ETF), on US (United States) Treasury bonds with maturities above 20 years
  - volatility 10.4% with the 2002–2021 correlation (-0.40), 13.2% with the 2022–2026 one (0.11)
- The 252-day SPY–TLT correlation was positive on 6% of days before 2022 and on 81% of days since January 2022
- 2008: the variance of 2-day S&P 500 returns was 7.5 times its calm-year level
  - with a constant β, a calm-period correlation of 0.4 is then measured as 0.77
  - after the Forbes–Rigobon adjustment, neither the Euro Stoxx 50 nor the BET (Bucharest Exchange Trading, index) crisis rise remains significant (stationary bootstrap: $p = 0.98$ and 0.62)
- JPM (JPMorgan Chase & Co., NYSE ticker) and BAC (Bank of America Corporation, NYSE ticker): when JPM has one of its 5% worst days, BAC is in its own worst 5% about 61% of the time

▸ Central chart: Stocks and Bonds: A Correlation That Changed Sign (1/2)

When Do Markets Move Together, and How Sure Can We Be? (1/2)

What we will explain

- Part I: why dependence matters and covariance uncertainty
  - multivariate GARCH (generalised autoregressive conditional heteroskedasticity)
  - DCC (dynamic conditional correlation)
  - crisis correlations and the Forbes–Rigobon bias
- Part II: copulas and tail dependence
  - European and Romanian markets, banks
  - common volatility, DCC in vast dimensions and research with artificial intelligence (AI)

What you should be able to do

- compute a time-varying covariance
  - rolling window, EWMA (exponentially weighted moving average), DCC
- derive the delta-method standard error of a correlation
- derive the volatility bias of crisis correlations and separate it from contagion (a genuine rise in cross-market linkages)
- compute copula fits and goodness-of-fit tests
- interpret tail dependence
  - at a finite threshold and in the limit
- Further Topics and Appendix: reading for familiarity

When Do Markets Move Together, and How Sure Can We Be? (2/2)

- Builds on
  - Chapter 1 (fat tails, volatility clustering)
  - Chapter 4 (Σ in portfolio choice)
  - Chapter 5 (univariate GARCH)
- Leads to
  - Chapter 7 (portfolio Value at Risk, VaR, and Expected Shortfall, ES)
  - Chapter 8 (backtesting)
  - Chapter 18 (systemic risk)
- Detailed examples and advanced topics are in Further Topics, after the Summary
  - the footer button leads there; the Back button returns to the starting slide

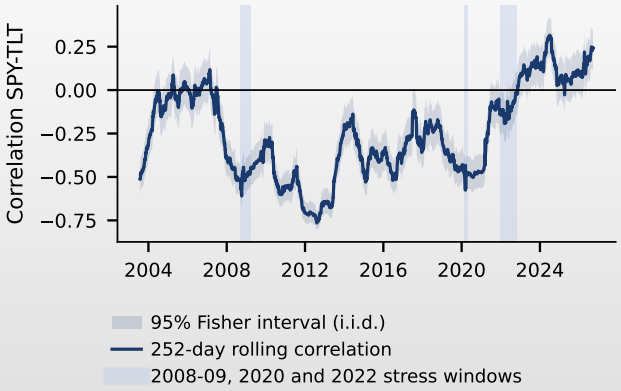
Why Dependence Matters (1/2): Portfolio Diversification

- **Diversification** (Markowitz (1952))
  - The portfolio variance depends on the correlation through a single cross term
  - $\sigma_p^2 = w^\top \Sigma w = w^2\sigma_1^2 + (1-w)^2\sigma_2^2 + 2w(1-w)\rho\,\sigma_1\sigma_2$
  - w: weight of asset 1 ($1-w$ in asset 2)
  - Σ: covariance matrix
  - σ_1, σ_2: volatilities
  - ρ: correlation
  - Portfolio 60/40 of SPY (S&P 500 exchange-traded fund) and TLT (20+ year US Treasury bond fund)
  - Annualised volatilities over July 2002 – September 2026: $\sigma_1 = 18.8\%$ (SPY), $\sigma_2 = 14.3\%$ (TLT); weights rebalanced daily to 60/40
  - We split at January 2022, the start of the inflation shock; the break test in Further Topics dates the change earlier, to August 2020
  - Portfolio volatility 10.4% with the 2002–2021 correlation (-0.40), 13.2% with the 2022–2026 one (0.11)

Why Dependence Matters (2/2): Hedges, Contagion and Moving Covariances

- **Hedging**
  - How much value of the hedge F to sell per unit of value held in the position S to minimise the variance
  - Minimum-variance hedge ratio $h^* = \rho\,\sigma_S/\sigma_F$: the slope of the regression of r_S on r_F. Estimated on returns, it is a ratio of values: units of F per unit of S are $h^* P_S/P_F$ (P_S, P_F: prices, contract size included)
  - r_S, r_F: returns of the position and of the hedge
  - σ_S, σ_F: their volatilities
  - Effectiveness ρ^2: the share of the variance removed
  - If the correlation moves, yesterday's hedge is wrong today: a dynamic hedge needs a dynamic ρ_t
- **Contagion and interdependence** (Forbes & Rigobon (2002))
  - Contagion: a *significant increase* in cross-market linkages after a shock
  - Interdependence: links present in all states (trade, common investors, common macro factors)
  - Correlations tend to rise exactly when diversification is needed most
- New in this chapter: Σ_t, the covariance matrix at date t, moves

Stocks and Bonds: A Correlation That Changed Sign (1/2)



MFM_ch6_spy_tlt_rolling ▶ Interactive

Stocks and Bonds: A Correlation That Changed Sign (2/2)

- Daily SPY and TLT returns, 252-day window
  - band: 95% Fisher interval (defined in the next section) assuming i.i.d. (independent and identically distributed) data
  - shaded: autumn 2008 – March 2009, February–April 2020, January–October 2022
- Values of the correlation
  - minimum: -0.77 (8 August 2012)
  - maximum: 0.32 (1 July 2024)
  - last: 0.24 (18 September 2026)
- Positive correlation
  - on 6% of days before 2022
  - on 81% of days since January 2022
- The sign depends on the type of shock
  - growth shocks move stocks and bond prices in opposite directions (negative correlation)
  - inflation shocks hit both (positive, as in 2022)
- The Fisher band is too narrow here
  - it ignores fat tails and volatility clustering (next section)
- Type of evidence: description (it shows what happened in the data, without testing a cause)

MFM_ch6_spy_tlt_rolling

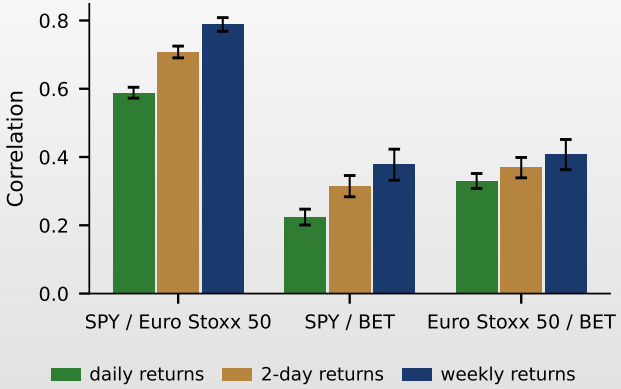
How Precise Is an Estimated Correlation?

- The question of this section: how do we estimate Σ, and how uncertain is each estimated correlation?
- First trap: closing times differ, so the data frequency changes the correlation
- Second trap: fat tails and volatility clustering can change the uncertainty substantially, and the Fisher interval ignores them
- Third trap: the window length decides how fast we see a regime change

Asynchronous Trading and the Bias of Daily Correlations

- Closing times differ
  - the Bucharest Stock Exchange and the euro-area exchanges close hours before New York
  - US news after the European close enters European prices only the next day
- Consequence, under positive dependence (as between equity markets)
  - the same-day correlation is biased towards zero
  - the lagged correlation is positive
- Remedies
  - Lower frequency: 2-day or weekly returns on common trading days
  - Add the lead–lag term: $\rho_0 + \rho_{+1}$ (ρ_0: same-day correlation, ρ_{+1}: US yesterday with Europe today)
  - The sum $\rho_0 + \rho_{+1}$ is a rough approximation of the synchronous correlation, valid only for a one-day lag, no autocorrelation and equal variances
- Rule used throughout
  - first align the *prices* on common trading days
  - then compute returns

Daily, 2-Day and Weekly Correlations, 2000–2026 (1/2)



MFM_ch6_async ▶ Interactive

Daily, 2-Day and Weekly Correlations, 2000–2026 (2/2)

- SPY, Euro Stoxx 50 and BET on common trading days
  - BET: the main index of the Bucharest Stock Exchange
  - vertical lines on the bars: 95% Fisher intervals
- SPY–Euro Stoxx 50 correlation
  - 0.59 daily, 0.71 on 2-day returns, 0.79 weekly
  - SPY yesterday with Euro Stoxx 50 today: 0.20
- SPY–BET correlation
  - 0.22 daily versus 0.38 weekly (lagged 0.17)
- Euro Stoxx 50–BET, almost synchronous closes
  - 0.33 versus 0.41
- Weekly data are cleaner
  - but give 1 367 instead of 6 446 observations: wider intervals
- Type of evidence: comparison of procedures (it shows which method works better on the same data, not why)

MFM_ch6_async

Sampling Frequencies: Daily, Two-Day and Weekly Returns

- Returns are log returns, $r_t = \ln P_t - \ln P_{t-1}$
  - computed after aligning the prices on common trading days
- Daily: DCC, HAC intervals and hedging, where the reaction speed matters
  - HAC: standard errors robust to heteroskedasticity and autocorrelation, defined below
- 2-day: the Forbes–Rigobon tests, to absorb different closing times
  - in these tests 2-day returns overlap: two consecutive returns share one day
  - in the three-frequency chart they do not overlap
- Weekly (Friday closes): exceedance correlations, market integration and bank networks
  - applications of the copula and applications sections, which mix time zones

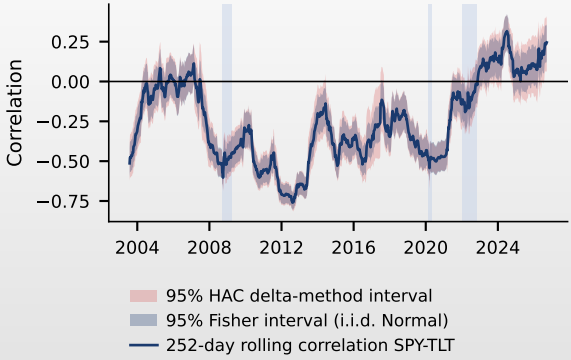
Sample Covariance and Inference on a Correlation (1/2): Estimator and Fisher Interval

- Sample covariance matrix and Pearson correlation of N assets over T days
  - $\hat{\Sigma} = \frac{1}{T-1} \sum_t (r_t - \bar{r})(r_t - \bar{r})^\top, \quad \hat{\rho}_{ij} = \hat{\Sigma}_{ij} / \sqrt{\hat{\Sigma}_{ii} \hat{\Sigma}_{jj}}$
  - $r_t \in \mathbb{R}^N$: vector of the N returns on day t, $t = 1, \ldots, T$
  - $\bar{r}$: vector of sample means
  - $\hat{\Sigma}_{ij}$: element (i, j)
  - $N(N+1)/2$ distinct elements to estimate: 125 250 for $N = 500$
- **Fisher transform**: $z = \operatorname{artanh} \hat{\rho} = \frac{1}{2} \ln \frac{1+\hat{\rho}}{1-\hat{\rho}}$
  - approximately Normal with a variance that does not depend on ρ
  - $z \approx N\big(\operatorname{artanh} \rho,\, 1/(T-3)\big)$ for i.i.d. pairs with the bivariate Normal distribution
  - 95% interval $\tanh(z \pm 1.96/\sqrt{T-3})$
  - $T = 252$, $\rho = 0.5$: half-width about 0.09

Sample Covariance and Inference on a Correlation (2/2): Delta Method and HAC

- Delta method: linearise $\hat{\rho}$, a smooth function of sample moments, and write the error $\hat{\rho} - \rho$ as a mean of daily contributions ψ_t
  - $\hat{\rho} - \rho \approx \frac{1}{T} \sum_t \psi_t, \quad \psi_t = \tilde{x}_t \tilde{y}_t - \frac{\rho}{2}(\tilde{x}_t^2 + \tilde{y}_t^2)$
  - $\tilde{x}_t$, $\tilde{y}_t$: the two returns standardised to mean 0 and variance 1
  - ψ_t: the contribution of day t to the error
  - volatility clustering does not change this form, only the variance of ψ_t
- Standard error (SE) from the long-run variance of ψ_t: its variance plus its autocovariances
  - HAC (heteroskedasticity and autocorrelation consistent): autocovariances up to lag L with declining weights (Bartlett kernel, Newey & West (1987))
  - Needs stationarity, finite fourth moments (questioned in Chapter 1) and weak dependence
  - Lags by the usual rule of thumb $L = \lfloor 4(T/100)^{2/9} \rfloor$: $L = 9$ for $T = 4\,891$, $L = 6$ for $T = 1\,182$, $L = 4$ for a 252-day window
  - i.i.d. Normal pairs give $\mathrm{Var}(\psi_t) = (1 - \rho^2)^2$, the Fisher variance
  - Fat tails and GARCH can change it substantially, in either direction (derivation in the Appendix)
- SPY–TLT daily, 2002–2021 ($T = 4\,891$, $\hat{\rho} = -0.403$): the SE of $\hat{\rho}$
  - 0.0120 (Fisher)
  - 0.0217 (fourth moments, $L = 0$)
  - 0.0227 (HAC, $L = 9$)

Fisher versus HAC Intervals: SPY and TLT (1/2)



MFM_ch6_corr_inference ▶ Interactive

Fisher versus HAC Intervals: SPY and TLT (2/2)

- 252-day rolling correlation, every 5th day
  - the 95% Fisher interval
  - the 95% HAC delta-method interval (Bartlett, $L = 4$)
- HAC is wider than Fisher in 91% of windows and narrower in 9%
  - width ratio: median 1.24, 90% quantile 1.54
- Correlation significantly different from zero
  - 73% of windows by Fisher, 65% by HAC
  - the verdict flips in 8.2%
  - pointwise results, window by window, not a simultaneous test
- Whole periods: HAC SE $= 1.9\times$ Fisher in 2002–2021, $1.4\times$ since 2022
  - fourth moments alone give $1.8\times$ in 2002–2021 ($1.28\times$ since 2022)
  - fat tails matter more than autocorrelation
- Rule for daily returns: report delta-method HAC intervals
  - valid under stationarity and weak dependence
  - and only if fourth moments are finite, which the Hill estimates of Chapter 1 put in doubt
  - HAC handles the dependence but does not compensate for missing moments
- Type of evidence: comparison of procedures

MFM_ch6_corr_inference

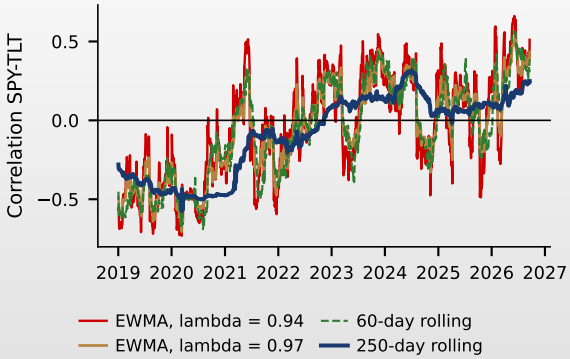
Rolling Windows

- **Rolling window** of length m: $\hat{\Sigma}_t = \frac{1}{m} \sum_{s=t-m}^{t-1} r_s r_s^\top$
  - mean set to zero for daily data
  - r_s: return vector of day s
  - m: number of past days used
- Positive semidefinite
  - for any weight vector v, $v^\top \hat{\Sigma}_t v = \frac{1}{m} \sum_s (v^\top r_s)^2 \geq 0$
  - so every portfolio variance is ≥ 0
- Singular if $m < N$
  - positive definite only if the m return vectors span $\mathbb{R}^N$
  - a singular matrix cannot be inverted, so it gives no minimum-variance weights
- **Ghost effect**: an extreme day enters with weight $1/m$ and leaves abruptly m days later
  - the estimate then jumps although the market has not changed
- The trade-off
  - a short window reacts fast but is noisy
  - a long one is stable but detects a regime change late

EWMA (RiskMetrics)

- **EWMA** (exponentially weighted moving average, J.P. Morgan/Reuters (1996))
  - $\Sigma_t = \lambda \Sigma_{t-1} + (1-\lambda) r_{t-1} r_{t-1}^\top = (1-\lambda) \sum_{k \geq 1} \lambda^{k-1} r_{t-k} r_{t-k}^\top$
  - $\lambda \in (0,1)$: decay factor
  - k: age of the observation in periods, weight $(1-\lambda)\lambda^{k-1}$
- Usual values: $\lambda = 0.94$ for daily and $\lambda = 0.97$ for monthly data
- Half-life $\ln 0.5 / \ln \lambda$: the number of periods after which an observation's weight halves
  - 11.2 periods for $\lambda = 0.94$ and 22.8 for $\lambda = 0.97$
- Effective sample $(1+\lambda)/(1-\lambda)$
  - about 32 days for $\lambda = 0.94$
- Positive semidefinite, no ghost effect
- Limitation: one λ for all elements and no mean reversion (forecasts do not converge to a long-run level)
  - i.e. an integrated GARCH(1,1) without a constant (Chapter 5)
- λ plays the role of m: the same trade-off between reaction speed and noise

EWMA versus Rolling Windows, SPY–TLT since 2019



- Two EWMA estimators ($\lambda = 0.94$ and 0.97) and two rolling windows (60 and 250 days)
- Noise: standard deviation of the daily change of the estimate
  - 0.060 (EWMA 0.94)
  - 0.032 (EWMA 0.97)
  - 0.0088 (250-day window)
- Regime signal: 20 positive days in a row
  - EWMA ($\lambda = 0.94$): began on 12 April 2022, confirmed on 10 May
  - 250-day window: began on 10 November, confirmed only on 8 December 2022
- With EWMA and rolling windows, λ and m are chosen by the analyst, not estimated
- DCC (next section) estimates the reaction speed and the memory of the correlation from the data
- Type of evidence: comparison of procedures

MFM_ch6_ewma_rolling ▶ Interactive

From Univariate to Multivariate Volatility (1/2): the Model

- The question of this section: how do we model a covariance matrix that changes every day?
  - the matrix must stay positive definite and have few parameters
- Chapter 5: one asset, $r_t = \sigma_t z_t$, $\sigma_t^2 = \omega + \alpha r_{t-1}^2 + \beta \sigma_{t-1}^2$
- Now N assets and a covariance matrix
  - $r_t = H_t^{1/2} z_t, \quad \mathbb{E}(z_t \mid \mathcal{F}_{t-1}) = 0, \quad \text{Var}(z_t \mid \mathcal{F}_{t-1}) = I_N, \quad \text{so Var}(r_t \mid \mathcal{F}_{t-1}) = H_t$
  - r_t, z_t: vectors of N returns and innovations
  - H_t: $N \times N$ conditional covariance matrix, known at $t-1$ (Σ_t in the previous section)
  - $H_t^{1/2}$: a matrix square root, $H_t^{1/2}(H_t^{1/2})^\top = H_t$
  - I_N: identity matrix
  - $\mathcal{F}_{t-1}$: information up to $t-1$
- Any model for H_t must deliver
  - positive definiteness at every t (every portfolio has a positive variance)
  - Few parameters
  - Interpretable dynamics
- Gaussian quasi-log-likelihood, maximised over the parameters of H_t
  - $\ell = -\frac{1}{2} \sum_t \left(N \ln 2\pi + \ln |H_t| + r_t^\top H_t^{-1} r_t \right)$
  - $|H_t|$: determinant of H_t
  - Quasi: the estimates stay consistent if z_t is not Normal but H_t is right

From Univariate to Multivariate Volatility (2/2): the Curse of Dimension

- **VEC** (Bollerslev, Engle & Wooldridge (1988)):
  $\operatorname{vech}(H_t) = c + A \operatorname{vech}(r_{t-1} r_{t-1}^\top) + B \operatorname{vech}(H_{t-1})$
  - vech: the vector of the $N(N+1)/2$ distinct elements of a symmetric matrix; A, B: square matrices of that size
  - Positive definiteness not guaranteed
- **BEKK** (Baba–Engle–Kraft–Kroner, Engle & Kroner (1995)):
  $H_t = CC^\top + A^\top r_{t-1} r_{t-1}^\top A + B^\top H_{t-1} B$
  - C: lower triangular; A, B: $N \times N$ matrices
  - Positive definite by construction (a sum of quadratic forms), but the elements of A, B have no direct interpretation (Appendix)
- **DCC**: $3N$ univariate GARCH parameters plus a, b
  - the target $\bar{Q}$ is estimated separately (next slide)
- The number of parameters grows like N^4 (VEC) and N^2 (BEKK)
  - infeasible for large portfolios

Number of parameters

N	VEC	full BEKK	DCC
2	21	11	$6+2$
5	465	65	$15+2$
10	6 105	255	$30+2$
50	3 252 525	6 275	$150+2$

The Two-Asset Case (1/2): the VEC Model

- With $N = 2$ assets, the covariance matrix has three distinct elements
  - $H_t = \begin{pmatrix} h_{11,t} & h_{12,t} \\ h_{12,t} & h_{22,t} \end{pmatrix}$, $\quad \text{vech}(H_t) = (h_{11,t}, h_{12,t}, h_{22,t})^\top$
  - $h_{11,t}$, $h_{22,t}$: the conditional variances of the two assets; $h_{12,t}$: their conditional covariance
- **Diagonal VEC**: each element has its own GARCH(1,1)-type equation
  - $h_{11,t} = c_1 + a_1 r^2_{1,t-1} + b_1 h_{11,t-1}$
  - $h_{12,t} = c_2 + a_2\, r_{1,t-1} r_{2,t-1} + b_2 h_{12,t-1}$
  - $h_{22,t} = c_3 + a_3 r^2_{2,t-1} + b_3 h_{22,t-1}$
  - the covariance rises after a day on which the two returns have the same sign and falls after one with opposite signs
- **Full VEC**: each element depends on all three, for example
  - $h_{11,t} = c_1 + a_{11} r^2_{1,t-1} + a_{12} r_{1,t-1} r_{2,t-1} + a_{13} r^2_{2,t-1} + b_{11} h_{11,t-1} + b_{12} h_{12,t-1} + b_{13} h_{22,t-1}$
  - parameters: 3 in c, $3 \times 3 = 9$ in A and 9 in B, so 21 (the table on the previous slide)
  - the general form $\text{vech}(H_t) = c + A\,\text{vech}(r_{t-1} r_{t-1}^\top) + B\,\text{vech}(H_{t-1})$ is exactly this system, written in matrix form
- The problem: nothing prevents $h^2_{12,t} > h_{11,t} h_{22,t}$, that is, a "correlation" above 1 in absolute value
  - hence no guarantee of positive definiteness

BEKK, Intuitively

- The idea: a variance cannot be negative if we write it as a square
  - univariate: $\alpha r_{t-1}^{2} = (a\, r_{t-1})^{2}$ with $a = \sqrt{\alpha}$, so the news term is always ≥ 0
  - BEKK does the same with matrices: $A^{\top} r_{t-1} r_{t-1}^{\top} A = (A^{\top} r_{t-1})(A^{\top} r_{t-1})^{\top}$, the "square" of a vector
- What $A^{\top} r_{t-1}$ is: yesterday's shocks, mixed
  - asset 1 receives $a_{11} r_{1,t-1} + a_{21} r_{2,t-1}$: its own shock plus part of the other asset's shock
  - with a diagonal A there is no mixing: each asset reacts only to its own shock
- The square of a vector is always a valid covariance matrix
  - the diagonal holds squares (≥ 0), and the "correlation" of this term is exactly $+1$ or -1, depending on the sign of $r_{1,t-1} r_{2,t-1}$
  - the same holds for $B^{\top} H_{t-1} B$ (a transformation of a valid covariance) and for $CC^{\top}$; a sum of valid covariances is valid
  - this is why H_t stays positive definite without any restriction on the parameters
- What we get: H_t = a baseline level + the reaction to yesterday's news + the memory of H_{t-1}
  - the correlation in H_t combines the correlations of the three terms, weighted by each term's contribution to the variances: the news brings ± 1, so days with same-sign shocks push it up
- The price: A and B have N^2 elements each, and one element does not describe a single variance or covariance
  - this is why the diagonal or scalar BEKK ($A = a\, I_N$, $B = b\, I_N$), close in spirit to DCC, is used in practice

The Two-Asset Case (2/2): the BEKK Model

- The general form: $H_t = CC^\top + A^\top r_{t-1} r_{t-1}^\top A + B^\top H_{t-1} B$
- **Diagonal BEKK** with $N = 2$: $A = \mathrm{diag}(a_{11}, a_{22})$, $B = \mathrm{diag}(b_{11}, b_{22})$
  - $h_{11,t} = (CC^\top)_{11} + a_{11}^2 r_{1,t-1}^2 + b_{11}^2 h_{11,t-1}$
  - $h_{22,t} = (CC^\top)_{22} + a_{22}^2 r_{2,t-1}^2 + b_{22}^2 h_{22,t-1}$
  - $h_{12,t} = (CC^\top)_{12} + a_{11} a_{22}\, r_{1,t-1} r_{2,t-1} + b_{11} b_{22}\, h_{12,t-1}$
- Compared with the diagonal VEC: the same form, but the covariance coefficients are no longer free
  - they are products of the coefficients of the two variance equations: $a_{11} a_{22}$ instead of a_2, $b_{11} b_{22}$ instead of b_2
  - $A^\top r_{t-1} r_{t-1}^\top A$ and $B^\top H_{t-1} B$ are quadratic forms, hence positive semidefinite; with $CC^\top$ positive definite, H_t stays positive definite at every t
  - 7 parameters: 3 in C (lower triangular), 2 in A, 2 in B
- **Full BEKK**: A, B arbitrary 2×2 matrices, so $3 + 4 + 4 = 11$ parameters
  - for example $h_{11,t}$ contains $(a_{11} r_{1,t-1} + a_{21} r_{2,t-1})^2$: it also depends on $r_{2,t-1}^2$ and on $r_{1,t-1} r_{2,t-1}$
  - this is why the elements of A and B have no direct interpretation

Constant and Dynamic Conditional Correlation

- Decomposition: covariance = volatilities $\times$ correlations $\times$ volatilities
  - $H_t = D_t R_t D_t, \quad D_t = \text{diag}(\sigma_{1t}, \ldots, \sigma_{Nt})$
  - σ_{it}: GARCH volatility of asset i; R_t: conditional correlation matrix; $\varepsilon_{it} = r_{it}/\sigma_{it}$: standardised residuals
- **CCC** (constant conditional correlation, Bollerslev (1990)): $R_t = R$ constant
- **DCC** (Engle (2002)): a GARCH-type recursion for a proxy matrix Q_t
  - $Q_t = (1 - a - b)\,\bar{Q} + a\,\varepsilon_{t-1}\varepsilon_{t-1}^\top + b\,Q_{t-1}, \quad a, b \geq 0,\ a + b < 1$
  - $R_t = \text{diag}(Q_t)^{-1/2}\, Q_t\, \text{diag}(Q_t)^{-1/2}$
  - $\bar{Q}$: the intercept matrix of the recursion, with $N(N-1)/2$ free elements. Correlation targeting replaces it by the sample covariance $\hat{\bar{Q}} = \frac{1}{T}\sum_t \varepsilon_t \varepsilon_t^\top$ (when this is consistent: slide "Two-Step Estimation (2/2)")
  - a: reaction to yesterday's product of shocks; b: memory; $a + b$: persistence, i.e. how slowly the correlation returns to the target
  - The second line rescales Q_t so that its diagonal is 1: R_t is a valid correlation matrix
- Two parameters govern all $N(N-1)/2$ correlations
  - parsimonious, but every pair has the same persistence
- Q_t is a covariance proxy, not a correlation: $\mathbb{E}_{t-1}[\varepsilon_t \varepsilon_t^\top] = R_t \neq Q_t$

The Two-Asset Case: CCC and DCC

- The volatility step: two univariate GARCH(1,1) models, as in Chapter 5
  - $\sigma^2_{i,t} = \omega_i + \alpha_i r^2_{i,t-1} + \beta_i \sigma^2_{i,t-1}$, $i = 1, 2$; $\quad \varepsilon_{i,t} = r_{i,t}/\sigma_{i,t}$
- The covariance is the correlation times the two volatilities: $h_{12,t} = \rho_{12,t}\, \sigma_{1,t}\sigma_{2,t}$
  - CCC: $\rho_{12,t} = \rho_{12}$ constant, so the covariance moves only through the volatilities
- DCC: three recursions, all with the same a and b
  - $q_{11,t} = (1 - a - b)\, \bar{q}_{11} + a\, \varepsilon^2_{1,t-1} + b\, q_{11,t-1}$, and likewise for $q_{22,t}$
  - $q_{12,t} = (1 - a - b)\, \bar{q}_{12} + a\, \varepsilon_{1,t-1}\varepsilon_{2,t-1} + b\, q_{12,t-1}$
  - $\rho_{12,t} = q_{12,t}/\sqrt{q_{11,t}\, q_{22,t}}$: the division brings $\rho_{12,t}$ between -1 and 1
- Example, SPY and TLT: $a = 0.052$, $b = 0.934$, $\bar{q}_{12} = -0.248$ (the estimates on the following slides)
  - $q_{12,t} = 0.014 \cdot (-0.248) + 0.052\, \varepsilon_{1,t-1}\varepsilon_{2,t-1} + 0.934\, q_{12,t-1}$
  - a day on which both assets fall (or rise) sharply pushes the correlation up, a day with opposite signs pushes it down; with a large b, the effect fades slowly
- In general (previous slide): the same recursions for all $N(N-1)/2$ pairs, written as one matrix equation in Q_t

Two-Step Estimation (1/2): Splitting the Likelihood

- The Gaussian log-likelihood splits into a volatility part and a correlation part
  - $\ell = \underbrace{-\tfrac{1}{2}\sum_t \left(N \ln 2\pi + 2\ln|D_t| + r_t^\top D_t^{-2} r_t\right)}_{\ell_V} \underbrace{-\tfrac{1}{2}\sum_t \left(\ln|R_t| + \varepsilon_t^\top R_t^{-1}\varepsilon_t - \varepsilon_t^\top \varepsilon_t\right)}_{\ell_C}$
  - ℓ_V: depends only on the GARCH parameters (volatility part)
  - ℓ_C: depends on (a, b) given the volatilities (correlation part)
  - $|D_t|$, $|R_t|$: determinants
- Where the split comes from
  - substitute $H_t = D_t R_t D_t$ and $\varepsilon_t = D_t^{-1} r_t$ into the quasi-log-likelihood of the previous section
  - Determinant: $|H_t| = |D_t|^2 |R_t|$, so $\ln|H_t| = 2\ln|D_t| + \ln|R_t|$
  - Quadratic form: $r_t^\top H_t^{-1} r_t = r_t^\top D_t^{-1} R_t^{-1} D_t^{-1} r_t = \varepsilon_t^\top R_t^{-1} \varepsilon_t$
  - Add and subtract $\varepsilon_t^\top \varepsilon_t = r_t^\top D_t^{-2} r_t = \sum_i r_{it}^2/\sigma_{it}^2$: then $\ell_V = \sum_i \left[-\tfrac{1}{2}\sum_t (\ln 2\pi + \ln\sigma_{it}^2 + r_{it}^2/\sigma_{it}^2)\right]$, the sum of the N univariate Gaussian GARCH log-likelihoods
- Step 1: fit N univariate GARCH(1,1) models, keep σ_{it} and ε_{it}
- Step 2: fix $\bar{Q}$ at the sample covariance of ε_t and maximise ℓ_C over (a, b)
- The gain: instead of one joint maximisation, N univariate ones and one with two parameters; the cost: harder inference (next slide)

Two-Step Estimation (2/2): Targeting and Inference

- Targeting with the sample covariance makes the standard DCC estimator inconsistent in general
  - why: take expectations of the recursion, using $\mathbb{E}[\varepsilon_{t-\mathbf{1}}\varepsilon_{t-\mathbf{1}}^{\top}] = \mathbb{E}[R_{t-\mathbf{1}}]$
  - Stationarity gives $(1-b)\,\mathbb{E}[Q_t] = (1-a-b)\,\bar{Q} + a\,\mathbb{E}[R_t]$
  - $\hat{\bar{Q}} = \frac{\mathbf{1}}{T}\sum_t \varepsilon_t\varepsilon_t^{\top}$ estimates $\mathbb{E}[R_t]$, which equals the intercept $\bar{Q}$ only if $\mathbb{E}[Q_t] = \mathbb{E}[R_t]$
  - That fails in general: R_t rescales Q_t by its own random diagonal
  - The corrected DCC of Aielli (2013) uses $\mathrm{diag}(Q_{t-\mathbf{1}})^{\mathbf{1}/\mathbf{2}}\varepsilon_{t-\mathbf{1}}$ in place of $\varepsilon_{t-\mathbf{1}}$, so that $\mathbb{E}[Q_t] = \bar{Q}$ and targeting remains consistent
- Step-2 standard errors treat the GARCH estimates and $\bar{Q}$ as known
  - propagate their error with the "sandwich" covariance matrix of both steps stacked together (Newey & McFadden (1994))
  - or by bootstrapping both steps: resample the data, re-estimate both steps and use the spread of the estimates (▸ Appendix)
- CCC against DCC: $H_{\mathbf{0}} : a = 0$
  - under $H_{\mathbf{0}}$, Q_t no longer depends on b, so b is not identified and the $\chi^{\mathbf{2}}$ reference is only indicative
  - the p-value comes from a bootstrap: simulate data under CCC and re-estimate

DCC Estimates for SPY and TLT (1/2)

Step 1: GARCH(1,1), returns in %

	ω	α	β	$\alpha+\beta$
SPY	0.0283	0.127	0.848	0.975
TLT	0.0057	0.049	0.943	0.993

Step 2: DCC(1,1) and three standard errors

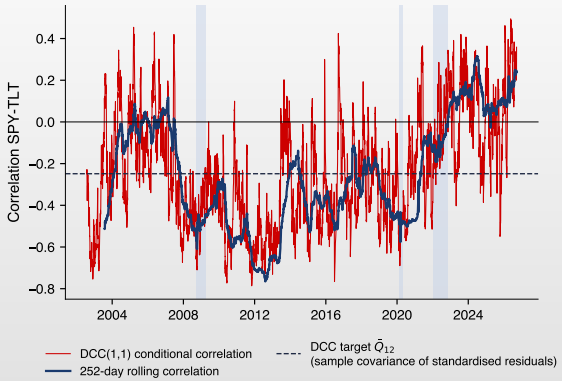
	estimate	Hessian	boot. step 2	boot. both
a	0.0522	0.0059	0.0051	0.0052
b	0.9335	0.0084	0.0077	0.0077
$a+b$	0.9857	–	0.0038	0.0040

- Daily data 31 July 2002 – 18 September 2026, $T = 6\,073$
  - target $\hat{\bar{Q}}_{12} = -0.248$
- Persistence $a + b = 0.986$
  - the heuristic measure $\ln 0.5 / \ln 0.986 \approx 48$ days, not the exact half-life of Q_t
  - $\mathbb{E}_t[\varepsilon_{t+1}\varepsilon_{t+1}^\top] = R_{t+1} \neq Q_{t+1}$, so the deviation of Q does not decay exactly as $(a+b)^h$
  - the actual response of the correlation comes from a conditional simulation
- LR (Likelihood Ratio, test) test against CCC: $LR = 2(\hat{\ell}^{DCC} - \hat{\ell}^{CCC}) = 2(460.26 - 193.54) = 533.4$
  - bootstrap of both steps under CCC ($B = 199$): the largest simulated value under H_0 is 15.7
  - so $p = 1/(B+1) = 0.005$, the smallest attainable value: CCC is rejected

DCC Estimates for SPY and TLT (2/2)

- In this bootstrap experiment, re-estimating GARCH raises the dispersion of $\hat{a}$ by 2% and of $\hat{b}$ by 1%
  - with $B = 199$, the Monte Carlo error of a bootstrap standard deviation is about 5%
  - here the step-1 uncertainty is therefore small, but the result is not general: the step-1 contribution can stay first-order as T grows
- The Hessian SEs exceed the bootstrap standard deviations, which does not prove that their intervals cover the true value more often than the nominal level
- The SPY–TLT correlation reacts little to a single day's shock (small a) but keeps its level for a long time (large b)
- TLT volatility is more persistent than SPY's ($\alpha + \beta$: 0.993 versus 0.975)
  - but reacts less to news (smaller α)
- These standard DCC estimates illustrate the method: with correlation targeting they are inconsistent in general (Two-Step Estimation (2/2))

DCC Correlation of SPY and TLT, 2002–2026

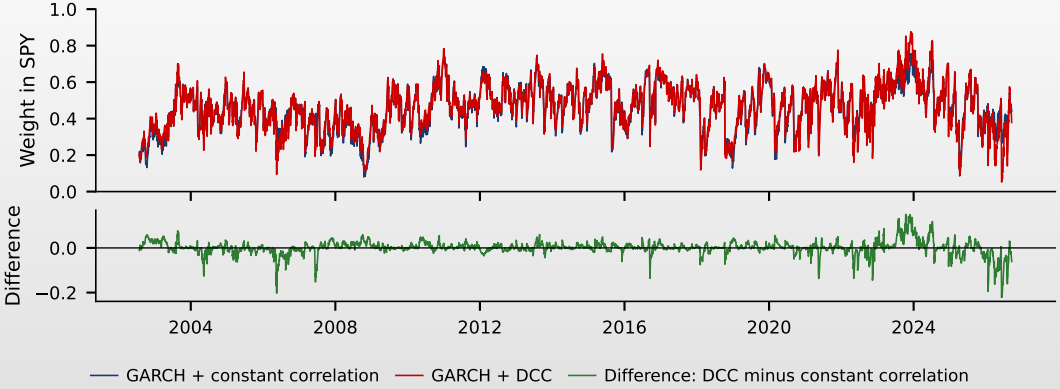


- DCC(1,1) conditional correlation, the 252-day rolling correlation and the target $\hat{\bar{Q}}_{12} = -0.248$
- The standard DCC path illustrates the method, subject to the targeting inconsistency (Two-Step Estimation (2/2))
- Range -0.79 (2 November 2011) to 0.49 (21 May 2026)
- Last value 0.36
- Average DCC correlation
  - -0.54 from 15 September to 31 December 2008
  - -0.56 in 2012, 0.00 in 2022
- Since January 2022 the path is above the target on 92% of days (36% before)
  - a single long-run correlation is doubtful
- Type of evidence: description

MFM_ch6_dcc_spy_tlt ▶ Interactive

▸ Further Topics: dating the break

Application: Minimum-Variance Weights with DCC (1/2)



MFM_ch6_dcc_minvar ▶ Interactive

Application: Minimum-Variance Weights with DCC (2/2)

- Weight in SPY of the daily minimum-variance SPY–TLT portfolio
  - $w_t = (\sigma^2_{2t} - \sigma_{12,t})/(\sigma^2_{1t} + \sigma^2_{2t} - 2\sigma_{12,t})$, truncated to $[0, 1]$ (no short sales)
  - $\sigma_{12,t} = \rho_t \sigma_{1t} \sigma_{2t}$: the conditional covariance
- Same GARCH volatilities, correlation from DCC or constant
  - the constant one equals the full-sample 2002–2026 return correlation (-0.30)
- Annualised volatility
  - 60/40 fixed: 10.99%
  - GARCH with constant correlation: 9.10%
  - GARCH with DCC: 9.06%
- Most of the gain comes from GARCH volatilities, DCC adds little
  - the DCC and constant-correlation weights differ by 0.019 on average, by at most 0.22 (8 June 2026)
- In-sample illustration: the parameters are estimated on the same data used for the evaluation
  - a fair test re-estimates on expanding windows and adds transaction costs (Chapter 4)
  - out-of-sample evaluation: Chapter 8
- Type of evidence: comparison of procedures

MFM_ch6_dcc_minvar

Contagion or Interdependence?



Lehman Brothers headquarters, New York, 15 September 2008, the day of the bankruptcy filing

Photo: Robert Scoble (2008); CC BY 2.0; Wikimedia Commons

- Common finding: correlations measured during crises are higher than in calm periods
- A high crisis correlation is not yet contagion
  - it must exceed the calm one
- Two very different explanations
  - **Contagion**: the transmission mechanism itself changes
  - **Volatility bias** (heteroskedasticity bias): the mechanism is unchanged, but a more volatile source market mechanically raises the correlation
- Forbes & Rigobon (2002): after correcting the bias, most "contagion" became interdependence (links present in all states)
  - in the 1987, 1994 and 1997 crises
- Longin & Solnik (2001), Ang & Chen (2002): correlations rise in bear markets, not in bull markets
  - an asymmetry the Normal distribution cannot produce

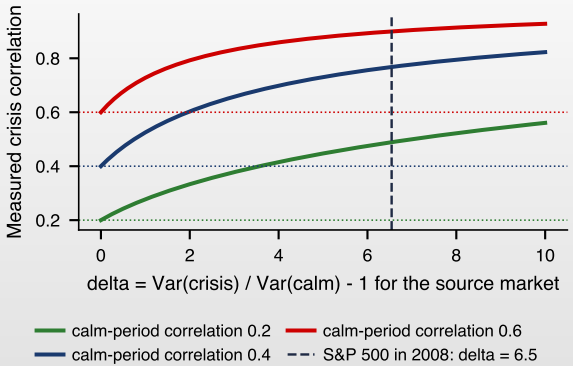
The Heteroskedasticity Bias (1/2): Correlation in a Calm Period

- Model: $y_t = \alpha + \beta x_t + \varepsilon_t$, $\mathrm{Cov}(x_t, \varepsilon_t) = 0$, β and σ_ε^2 constant
  - x_t: return of the source market
  - y_t: return of the other market
  - β: transmission coefficient
  - ε_t: own shock of y, variance σ_ε^2
  - σ_x^2: variance of x
- Two moments of the model
  - $\mathrm{Cov}(x_t, y_t) = \beta\sigma_x^2$ and $\mathrm{Var}(y_t) = \beta^2\sigma_x^2 + \sigma_\varepsilon^2$, because ε_t is uncorrelated with x_t
- Correlation: $\rho^2 = \dfrac{\beta^2\sigma_x^2}{\beta^2\sigma_x^2 + \sigma_\varepsilon^2} = \left(1 + \dfrac{\sigma_\varepsilon^2}{\beta^2\sigma_x^2}\right)^{-1}$
- The same in one number: the signal-to-noise ratio $k = \beta^2\sigma_x^2/\sigma_\varepsilon^2$
  - $\rho^2 = \dfrac{k}{1+k}$, equivalently $k = \dfrac{\rho^2}{1-\rho^2}$: the correlation depends on σ_x^2 only through k
  - a more volatile source market (larger σ_x^2) raises k, hence $|\rho|$, with β unchanged

The Heteroskedasticity Bias (2/2): From Calm to Crisis

- Crisis (subscript c): $\sigma^2_{x,c} = (1+\delta)\,\sigma^2_x$
  - $\delta > 0$: the relative rise of the source variance
  - nothing else changes
- Only σ^2_x grows, so $k_c = (1+\delta)\,k$ and
  - $\rho^2_c = \dfrac{k_c}{1+k_c} = \dfrac{(1+\delta)\rho^2/(1-\rho^2)}{1+(1+\delta)\rho^2/(1-\rho^2)} = \dfrac{(1+\delta)\rho^2}{1+\delta\rho^2}$
- Then $\rho_c = \rho\sqrt{\dfrac{1+\delta}{1+\delta\rho^2}}$, so $|\rho_c| > |\rho|$ for $0 < |\rho| < 1$
  - the measured correlation rises in absolute value although β is the same
- Inverting: $\rho^* = \dfrac{\rho_c}{\sqrt{1+\delta(1-\rho_c^2)}}$ is the crisis correlation adjusted for the volatility increase
- Quick check: the source variance doubles ($\delta = 1$) and $\rho = 0.5$
  - $k = 0.25/0.75 = 1/3$, $k_c = 2/3$, $\rho^2_c = (2/3)/(5/3) = 0.4$, so $\rho_c = 0.632$: a rise of 0.13 with no change in β
  - Adjusting: $\rho^* = 0.632/\sqrt{1+1\cdot(1-0.4)} = 0.632/1.265 = 0.5$, the calm value is recovered

How Much Can Volatility Alone Raise a Correlation?



- Measured crisis correlation ρ_c as a function of δ, for three calm-period correlations
  - β is constant: only the volatility of the source market rises
- Vertical line: the 2008 crisis window
  - the variance of 2-day S&P 500 returns was 7.5 times its calm-year level ($\delta = 6.54$)
- At this δ a calm-period correlation of 0.4 is measured as 0.77 in the crisis
  - and 0.2 as 0.49
- For an investor: a crisis rise in correlation of this size does not by itself prove contagion

MFM_ch6_fr_bias ▶ Interactive

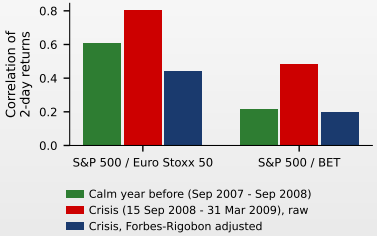
The Forbes–Rigobon Test and Its Inference

- Procedure
  - Choose the source market x, a calm window and a crisis window
  - Use 2-day returns (closing times differ)
  - Compute $\delta = \hat{\sigma}^2_{x,c}/\hat{\sigma}^2_{x,calm} - 1$ and ρ^*
  - Test $H_0 : \rho^* \leq \rho_{calm}$ (no contagion) with Fisher z statistics, the convention of Forbes & Rigobon (2002)
- Inference
  - With δ fixed, $\mathrm{Var}(\operatorname{artanh}\rho^*) \approx 1/[n_c(1+\delta(1-\rho_c^2))] < 1/n_c$ (n_c: crisis observations): the Fisher test is conservative (it rejects H_0 too rarely)
  - But δ is estimated and 2-day returns overlap: use a stationary bootstrap of the calm and crisis windows, resampled jointly, and test $D = \rho^* - \rho_{calm}$
  - Stationary bootstrap: resamples blocks of consecutive days (random starts, geometric lengths, mean 10 days; $B = 1\,999$ draws), keeping the serial dependence inside each block

Assumptions to remember

- No omitted common shock
  - a global factor hitting both markets raises σ_ε as well as σ_x
  - and the adjustment over-corrects (Corsetti, Pericoli & Sbracia (2005); ▸ Appendix)
- No feedback from y to x (no endogeneity)
- The choice of windows matters: results can flip with a few weeks
- The residual variance σ^2_ε does not change between periods

2008: Raw versus Adjusted Crisis Correlations



- Correlation of 2-day returns with the S&P 500
  - calm year, crisis (raw) and crisis after the Forbes–Rigobon adjustment
- Both raw rises (Euro Stoxx 50 and BET) are significant at 5%, neither survives the adjustment
- Stationary bootstrap: 90% interval for δ [3.9, 9.6]
- Euro Stoxx 50: 90% interval for ρ^* [0.38, 0.52], bootstrap $p = 0.98$
- BET: 90% interval for ρ^* [0.16, 0.24], bootstrap $p = 0.62$
  - bootstrap SE of ρ^* 0.025, against 0.047 from the Normal formula
- Verdict: no evidence of contagion once volatility is adjusted for, under the maintained assumptions
- Caveats
  - in 2008 a common global shock hit every market, so the adjustment may over-correct (Corsetti, Pericoli & Sbracia (2005))
  - the test also has low power (the 2020 application: Seminar B3)
- Type of evidence: association (the variables move together; this does not prove causation)

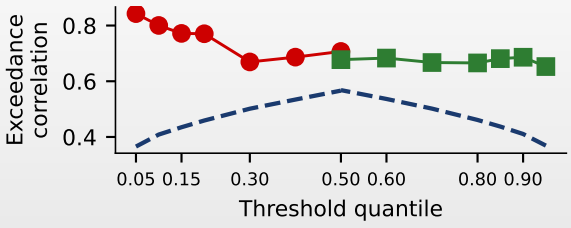
MFM_ch6_forbes_rigobon

▶ Interactive

Exceedance Correlations

- The question of this section: how likely are joint extreme losses?
  - and which dependence model reproduces them when correlation cannot?
- **Exceedance correlation** (Longin & Solnik (2001)): correlation conditional on both returns being beyond a threshold
- Lower tail at quantile q: $\rho^{-}(q) = \mathrm{Corr}\big(x, y \mid x \le F_x^{-1}(q),\, y \le F_y^{-1}(q)\big)$
  - $F_x^{-1}(q)$: the q-quantile of x
  - e.g. $q = 0.10$: the worst 10% of weeks
- Upper tail: the same with $x \ge F_x^{-1}(1-q)$, $y \ge F_y^{-1}(1-q)$
- Benchmark: for a bivariate Normal distribution, $\rho^{\pm}(q) \to 0$ as $q \to 0$
  - Conditioning on extremes lowers the correlation even when nothing changes: compare with the Normal curve, not with ρ

Exceedance Correlations: S&P 500 and Euro Stoxx 50



- Weekly returns 2000–2026 ($n = 1\,393$)
  - on the common trading days of the S&P 500 and the Euro Stoxx 50
  - overall $\hat{\rho} = 0.78$
  - dashed: the bivariate Normal distribution with the same correlation
- At $q = 0.10$
  - joint losses: $\rho^- = 0.80$ (86 weeks)
  - joint gains: $\rho^+ = 0.69$ (70 weeks)
  - Normal benchmark: 0.41
- Both tails lie far above the Normal benchmark
  - joint losses are more strongly correlated than joint gains
- A single coefficient cannot describe this
  - we need a model of the whole dependence structure: a copula (defined below)
- Type of evidence: description

MFM_ch6_exceedance ▶ Interactive

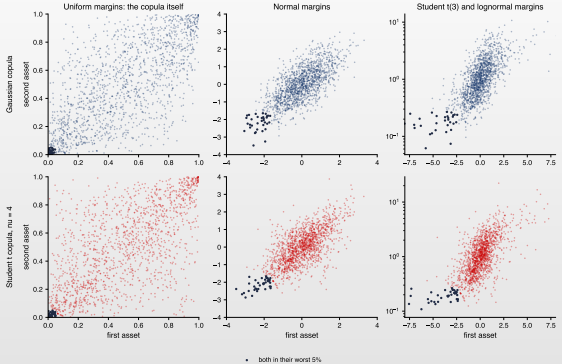
Why Correlation Is Not Enough

- Embrechts, McNeil & Straumann (2002): three fallacies about linear correlation
- Correlation measures only *linear* dependence
  - it is not invariant to monotone transformations: $\text{Corr}(x, y) \neq \text{Corr}(e^x, e^y)$
- Margins and correlation do not determine the joint distribution
  - Two portfolios with identical margins and ρ can have very different Value-at-Risk (Chapter 7)
- Attainable correlations: for given margins, $\rho \in [\rho_{\min}, \rho_{\max}]$
  - reached by countermonotone (one is a decreasing function of the other) and comonotone (increasing function) pairs
  - Lognormal $X = e^{\sigma_1 Z}$, $Y = e^{\pm\sigma_2 Z}$, Z standard Normal ($+$: comonotone, $-$: countermonotone): $\rho_{\max} = \dfrac{e^{\sigma_1\sigma_2} - 1}{\sqrt{(e^{\sigma_1^2} - 1)(e^{\sigma_2^2} - 1)}}$, $\rho_{\min}$ with $e^{-\sigma_1\sigma_2}$ in the numerator
  - $\sigma_1 = 1$, $\sigma_2 = 3$: $\rho_{\max} = 0.162$ and $\rho_{\min} = -0.008$: perfectly dependent, yet almost uncorrelated
- The remedy: separate the margins from the dependence structure (next slide)

What Is a Copula, Intuitively? (1/2)

- Step 1: replace each return by its percentile
  - $u = 0.05$ is a day among the worst 5% of that asset, $u = 0.5$ an ordinary day
- Step 2: the cloud of points (u, v) in the square $[0, 1]^2$ is the **copula**
  - it shows only how often the good and the bad days of the two assets coincide
  - the volatility, the fat tails and the units of each asset (the margins) are gone; only the link remains
- The margins say how each asset moves; the copula says how they move together
  - the same copula can be combined with any margins (Sklar's theorem, next slide)
- In the data: when JPM has one of its worst 5% days, BAC has one too in about 61% of cases
  - this is a property of the copula, not of the margins

What Is a Copula, Intuitively? (2/2)



- 1 500 simulated pairs with the same Kendall τ (0.5); each row keeps one copula (Gaussian at the top, Student t with $\nu = 4$ at the bottom), each column changes the margins
- Dark points: both assets in their worst 5% days, 34 pairs for the Gaussian, 40 for the t copula

Sklar's Theorem: Margins plus Copula

- A **copula** is a joint distribution function on $[0,1]^d$ with uniform margins: it describes only how the variables move together, once each is put on the probability scale $[0,1]$
- Sklar (1959): for any joint distribution F with margins $F_1, \ldots, F_d$ there is a copula C
  - with $F(x_1, \ldots, x_d) = C\big(F_1(x_1), \ldots, F_d(x_d)\big)$
  - d: number of variables
  - F_i: marginal distribution function of x_i
  - $u_i = F_i(x_i) \in [0,1]$: its probability level
  - C is unique if the margins are continuous
  - Conversely, any copula and any margins give a valid joint distribution
  - With densities: $f(x) = c\big(F_1(x_1), \ldots, F_d(x_d)\big) \prod_i f_i(x_i)$, c the copula density, f_i the marginal densities (the comonotone copula has none)
- Margins carry volatility, skewness and fat tails of each asset (Chapters 1 and 5)
- The copula carries all the dependence
  - it is invariant to increasing transformations of each variable, unlike ρ
- Three reference copulas and the Fréchet–Hoeffding bounds $W \le C \le M$ (Nelsen (2006))
  - Independence $\Pi(u,v) = uv$
  - Comonotonicity $M(u,v) = \min(u,v)$
  - Countermonotonicity $W(u,v) = \max(u+v-1, 0)$

Rank Correlations: Kendall and Spearman

- **Kendall's tau** (Kendall (1938)): concordant minus discordant pairs
  - $\tau = P\big[(X_1 - X_2)(Y_1 - Y_2) > 0\big] - P\big[(X_1 - X_2)(Y_1 - Y_2) < 0\big]$
  - (X_1, Y_1), (X_2, Y_2): two independent draws of the pair
  - Concordant: both coordinates move in the same direction
  - Copula form: $\tau = 4 \int\int C \, dC - 1 = 4\,\mathbb{E}[C(U, V)] - 1$
- **Spearman's rho** (Spearman (1904)): Pearson correlation of the ranks
  - Copula form: $\rho_S = 12 \int\int (C(u, v) - uv) \, du \, dv$
- Both lie in $[-1, 1]$ and depend only on the copula: 1 for M, 0 for Π, -1 for W
  - so, unlike ρ, they do not change under increasing transformations of the variables
- θ: the copula parameter

Links to copula parameters

Copula	τ
Gaussian, t	$\frac{2}{\pi} \arcsin \rho$
Clayton	$\theta/(\theta + 2)$
Gumbel	$1 - 1/\theta$
Frank	$1 - \frac{4}{\theta}\left[1 - D_1(\theta)\right]$

- $D_1(\theta) = \frac{1}{\theta} \int_0^\theta \frac{t}{e^t - 1} dt$ (Debye function)
- Gaussian: $\rho_S = \frac{6}{\pi} \arcsin(\rho/2)$
- Example, Gaussian copula with $\rho = 0.5$: $\tau = \frac{2}{\pi} \cdot \frac{\pi}{6} = 1/3$ and $\rho_S = \frac{6}{\pi} \arcsin 0.25 = 0.483$

Five Copula Families (1/2)

Family	$C(u, v)$	Parameter
Gaussian	$\Phi_{\mathbf{2},\rho}\left(\Phi^{-\mathbf{1}}(u), \Phi^{-\mathbf{1}}(v)\right)$	$-1 < \rho < 1$
Student t (Demarta & McNeil (2005))	$t_{\mathbf{2},\rho,\nu}\left(t_\nu^{-\mathbf{1}}(u), t_\nu^{-\mathbf{1}}(v)\right)$	$-1 < \rho < 1,\ \nu > 0$
Clayton (Clayton (1978))	$\left(u^{-\theta} + v^{-\theta} - 1\right)^{-\mathbf{1}/\theta}$	$\theta > 0$
Gumbel (Gumbel (1960))	$\exp\left(-\left[(-\ln u)^\theta + (-\ln v)^\theta\right]^{\mathbf{1}/\theta}\right)$	$\theta \geq 1$
Frank (Frank (1979))	$-\frac{\mathbf{1}}{\theta}\ln\left(1 + \frac{(e^{-\theta u}-\mathbf{1})(e^{-\theta v}-\mathbf{1})}{e^{-\theta}-\mathbf{1}}\right)$	$\theta \neq 0$

- Notation in the table
  - $\Phi, \Phi^{-\mathbf{1}}, \Phi_{\mathbf{2},\rho}$: the Normal distribution function, quantile and bivariate distribution function (correlation ρ)
  - t_ν, $t_{\mathbf{2},\rho,\nu}$: the Student t analogues with ν degrees of freedom; θ: dependence parameter
- Why these five: they cover the four patterns of tail dependence (the tendency of extremes to occur together, defined below)
  - none (Gaussian, Frank), both symmetric (t), lower tail only (Clayton), upper tail only (Gumbel)
- Elliptical (Gaussian, t): built from the multivariate Normal and Student t distributions
  - extend to d dimensions with a correlation matrix, the standard choice in risk systems
  - $\nu \to \infty$ turns the t copula into the Gaussian one

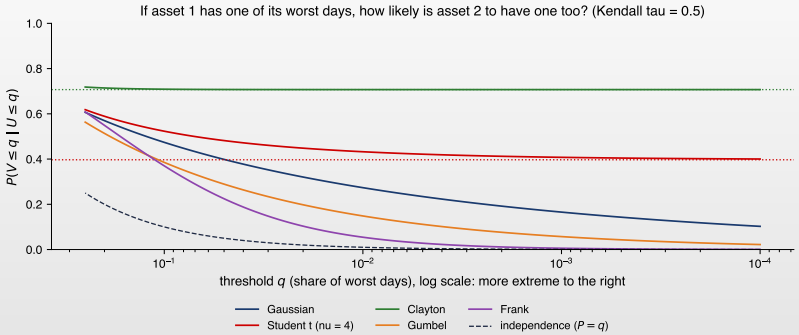
Five Copula Families (2/2)

- Where the names come from: the Gaussian copula is the link inside a pair with the bivariate Normal distribution (correlation ρ), after each variable is mapped to $[0,1]$ by Φ
  - the name says where the link comes from, not that the returns are Normal: the margins can be anything (the copula, intuitively)
  - the Student t copula: the same, from the bivariate t distribution; Clayton, Gumbel and Frank are named after the authors who proposed them
- Archimedean (Clayton, Gumbel, Frank): $C(u,v) = \varphi^{-1}\big(\varphi(u) + \varphi(v)\big)$, φ: generator (decreasing, $\varphi(1) = 0$)
  - One parameter, closed form, easy simulation, but exchangeable (symmetric in the variables, $C(u,v) = C(v,u)$) and poor in high dimension
  - for example, Clayton: $\varphi(t) = (t^{-\theta} - 1)/\theta$; Gumbel: $\varphi(t) = (-\ln t)^{\theta}$
  - the name: the generator turns combining into adding; applied n times to the same $u < 1$, the copula gives $\varphi^{-1}(n\varphi(u)) \to 0$, as in Archimedes' axiom ($n\varepsilon$ exceeds any bound for n large enough)
- Rotation by 180° moves the tail dependence to the opposite corner
  - the survival copula of the pair $(1-U, 1-V)$, where U and V are the uniform variables of the copula
  - e.g. a rotated Gumbel has dependence only in the lower tail (joint crashes)
  - densities and generators in the ▸ Appendix

Tail Dependence, Intuitively (1/2)

- The question: on the very bad days of one asset, how often does the other also have a very bad day?
- We count: take the worst 1% days of the first asset and see on how many of them the second is also in its own worst 1% days
  - independent assets: on only 1% of them
  - Gaussian copula with Kendall τ 0.5: 27.3% at the 1% threshold and 16.5% at 0.1%; the share falls towards 0
  - Student t copula ($\nu = 4$) with the same τ: 43.2% at 1% and 40.8% at 0.1%; the share stays around 40%
- **Tail dependence** λ is the value this share tends to as the threshold becomes ever more extreme
  - $\lambda = 0$: joint crashes become ever rarer than single ones; $\lambda > 0$: a fixed part of the crashes stays joint
- For risk: at the same τ, the portfolio loses on both assets at once far more often when $\lambda > 0$
  - in the data, at the 5% threshold: the very bad days of JPM coincide with those of BAC in about 61% of cases

Tail Dependence, Intuitively (2/2)



- Horizontal axis: the threshold q, more extreme to the right; vertical axis: the share of joint bad days, $P(V \leq q \mid U \leq q)$
- Dotted lines: the limits λ_L, 0.707 (Clayton) and 0.397 (Student t); the Gaussian, Gumbel and Frank tend to 0, although they have the same τ

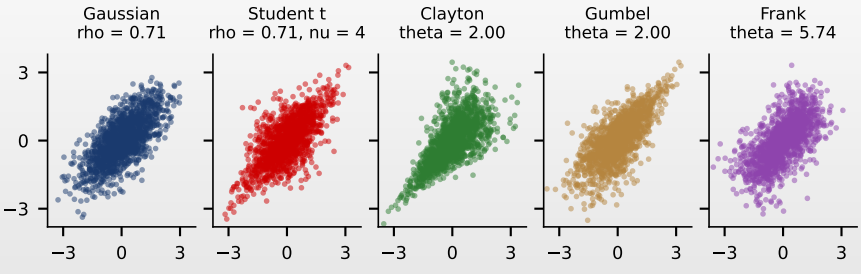
Tail Dependence

- The question: if one asset has an extreme day, how likely is the other to have one too?
- Lower tail:
  $\lambda_L = \lim_{q \downarrow \mathbf{0}} P(V \le q \mid U \le q) = \lim_{q \downarrow \mathbf{0}} C(q,q)/q$
  - $U = F_{\mathbf{1}}(X_{\mathbf{1}})$, $V = F_{\mathbf{2}}(X_{\mathbf{2}})$; q: the tail probability
- Upper tail: $\lambda_U = \lim_{q \downarrow \mathbf{0}} P(V > 1 - q \mid U > 1 - q)$
- $\lambda > 0$: the probability of a joint extreme stays positive however extreme the threshold
- $\lambda = 0$: asymptotic independence, the conditional probability of a joint extreme goes to 0 as the threshold becomes more extreme
- Tail dependence is a copula property, invisible in ρ (next slide)
- At a finite q, $C(q,q)/q$ is a conditional probability, not yet the limit: data only give such values
- $t_{\nu+\mathbf{1}}$ in the table: Student t distribution function with $\nu + 1$ degrees of freedom

Closed forms

	λ_L	λ_U
Gaussian	0	0
Student t	$2\, t_{\nu+\mathbf{1}}\Big(-\sqrt{\frac{(\nu+\mathbf{1})(\mathbf{1}-\rho)}{\mathbf{1}+\rho}}\Big)$	same
Clayton	$2^{-\mathbf{1}/\theta}$	0
Gumbel	0	$2 - 2^{\mathbf{1}/\theta}$
Frank	0	0

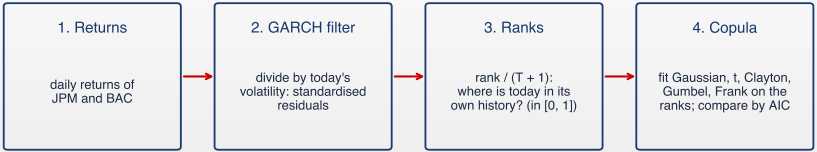
Five Copulas with the Same Kendall's Tau



- 1 500 simulated pairs from each copula with $\tau = 0.5$ on Normal-score margins $(\Phi^{-1}(u), \Phi^{-1}(v))$
  - Gaussian and t $\rho = 0.707$ (t with $\nu = 4$), Clayton and Gumbel $\theta = 2.0$, Frank $\theta = 5.74$
- Corners: Gaussian and Frank thin out, t stays dense in both
  - Clayton is sharp lower-left (joint crashes), Gumbel sharp upper-right
- Equal Kendall's tau does not determine joint-tail risk: the copula family matters

MFM_ch6_copula_zoo ▶ Interactive

Estimating Copulas, Intuitively

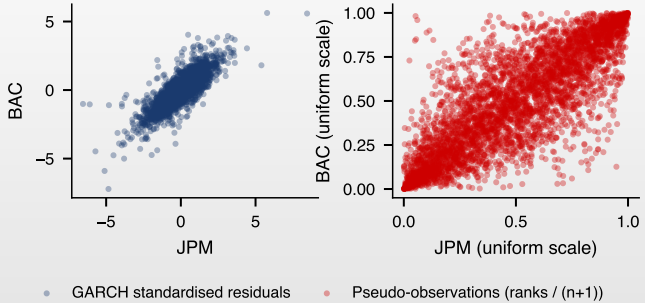


- Step 2, the GARCH filter: -3% means something else on a calm day than in a crisis
  - we divide each return by the volatility of the day, to compare days from different regimes
- Step 3, the ranks: each day becomes its percentile in the asset's own history
  - we assume no shape for the margins; a wrong margin cannot spoil the copula estimate
- Step 4, the likelihood: for each family we choose the parameter that makes the observed pairs of percentiles most probable
  - the copula density is high where the points are dense, for example in the corner of the joint bad days
- Then we compare the families: the AIC (Akaike Information Criterion) criterion penalises extra parameters; for JPM and BAC the t copula wins (next slides)

Estimating Copulas

- **Full maximum likelihood**: margins and copula estimated jointly
  - Efficient, but computationally demanding and sensitive to misspecification
- **IFM** (inference functions for margins, Joe (2005)): two steps
  - Fit parametric margins (e.g. GARCH with Student t)
  - Set $\hat{u}_{it} = \hat{F}_i(x_{it})$ ($\hat{F}_i$: fitted marginal distribution function)
  - Maximise $\sum_t \ln c(\hat{\boldsymbol{u}}_{\mathbf{1}t}, \hat{\boldsymbol{u}}_{\mathbf{2}t}; \theta)$, c the copula density
- **CML** (canonical maximum likelihood, Genest, Ghoudi & Rivest (1995)): no parametric margins
  - Pseudo-observations (rescaled ranks): $\hat{u}_{it} = \text{rank}(x_{it})/(T+1)$
  - Then maximise the copula likelihood
  - Robust to misspecified margins
  - The $T+1$ keeps $\hat{u}$ away from 0 and 1
- For returns: filter each series with GARCH first
  - without filtering, volatility clustering would mix dependence with changing volatility
  - then use the ranks of the standardised residuals (CML, the method used from here on)
  - the static CML inference shown here assumes i.i.d. innovations; GARCH filtering alone does not guarantee this
- Model choice: most of the families compared are not nested (none is a special case of another)
  - exception: the Gaussian is the t copula with $\nu = \infty$; for it we use an LR test calibrated by bootstrap (Seminar 6, B2)
  - for the others we use AIC $= 2k - 2\ell$ (Akaike information criterion), BIC $= k \ln T - 2\ell$ (Bayesian information criterion)
  - k: number of parameters, ℓ: maximised log-likelihood

From Residuals to Pseudo-Observations: JPM and BAC



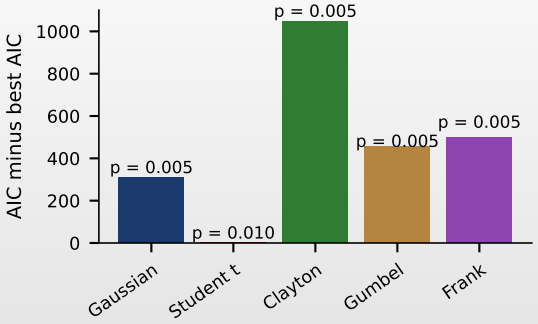
- JPMorgan Chase (JPM) and Bank of America (BAC), daily 2010–2026 ($T = 4\,202$), GARCH(1,1)-filtered
- Kendall's $\hat{\tau} = 0.641$, Spearman's $\hat{\rho}_S = 0.822$
- Both corners of the uniform scatter are dense: joint crashes and joint rallies
  - next question: which copula family reproduces this pattern?

MFM_ch6_copula_banks ▶ Interactive

Goodness-of-Fit Tests, Intuitively

- The question: do the data look as if they were generated by the estimated copula?
- The idea of the Rosenblatt test: if the model is right, each pair (u, v) turns into two numbers with no link between them
  - the first stays u; the second says where v falls in the distribution the model predicts for v, given u
  - with a good model, the transformed cloud fills the square $[0,1]^2$ evenly, with no denser areas
- The statistic S_n measures how far the transformed cloud is from an even cloud: 0 for a perfect fit, large values for a poor fit
  - JPM and BAC: t 0.082, Frank 0.365, Gaussian 0.968, Clayton 1.24, Gumbel 1.42
- The bootstrap p-value: simulate data B times from the estimated copula, redo the whole computation and count how often the simulated S_n is at least the observed one
  - here no simulation reaches the observed value, so p is the smallest possible, $1/(B+1)$: every family is rejected
- With 4 202 days the test sees even small departures: a rejection says the model is not exact, not how wrong it is; to choose we also use AIC

Copula Estimates and Goodness of Fit: JPM and BAC (1/2)



Numbers above bars: goodness-of-fit p-value
(Rosenblatt Cramer-von Mises, parametric bootstrap)

MFM_ch6_copula_banks ▶ Interactive

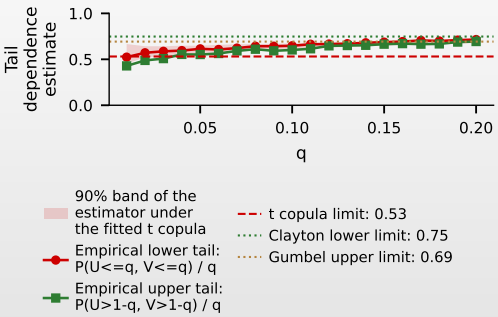
Copula Estimates and Goodness of Fit: JPM and BAC (2/2)

	$\hat{\theta}$ or $\hat{\rho}$	$\hat{\nu}$	λ_L	λ_U
Gaussian	0.833	–	0.000	0.000
Student t	0.842	4.24	0.531	0.531
Clayton	2.384	–	0.748	0.000
Gumbel	2.587	–	0.000	0.693
Frank	9.152	–	0.000	0.000

- λ_L, λ_U in the table: tail dependence implied by the estimated parameters
- The bars (previous slide) show the AIC distance from the best model
  - the numbers above them: the bootstrap goodness-of-fit p-values (full table in the ▸ Appendix)
- Student t has the lowest AIC, by a wide margin
- The Rosenblatt test (Further Topics) checks whether the data, transformed with the fitted copula, become independent uniforms: it rejects every family
  - each p-value is the smallest attainable, $1/(B+1)$
- Test statistics (smaller means better fit): t 0.082, Frank 0.365, Gaussian 0.968, Clayton 1.24, Gumbel 1.42
  - Frank is second by fit but fourth by AIC
- For a risk manager: the t copula is the best of the five, but not an exact model
- Type of evidence: comparison of procedures

MFM_ch6_copula_banks

Empirical Tail Dependence: JPM and BAC



- $\hat{\lambda}_L(q) = \hat{C}(q,q)/q$ and $\hat{\lambda}_U(q)$ for q from 1% to 20%
  - band: pointwise 90% range of the same estimator on data simulated from the fitted t copula
  - it conditions on the estimated parameters and is not simultaneous across thresholds
- $q = 5\%$: $\hat{\lambda}_L = 0.61$, $\hat{\lambda}_U = 0.55$
  - $q = 1\%$: 0.52 and 0.43
- When JPM has one of its 5% worst days
  - BAC is in its own worst 5% about 61% of the time
- Lower tail at 5%: inside the t-copula band (0.57–0.65)
- Upper tail (0.55 at 5%): below the t-copula band
  - joint rallies are weaker than the model predicts
- At $q = 1\%$ only 42 days remain in the tail, so the estimate is imprecise
- Type of evidence: description
  - the gap $\hat{\lambda}_L - \hat{\lambda}_U$ is not significant (Further Topics)

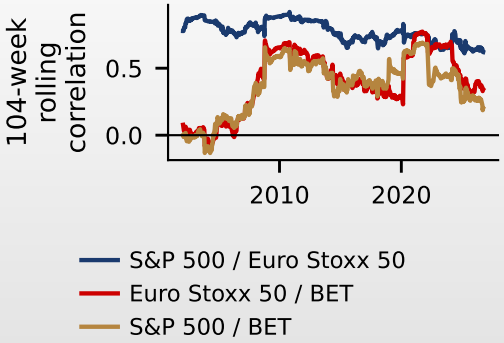
MFM_ch6_copula_banks ▶ Interactive

▸ Further Topics: asymmetry, dynamic t copula

The Gaussian Copula in the 2008 Crisis

- Li (2000) modelled joint default times with a Gaussian copula
  - One correlation parameter, calibrated to market spreads
  - It priced whole credit portfolios and became the standard for tranches of collateralised debt obligations (CDO)
- Zero asymptotic tail dependence ($\lambda_L = \lambda_U = 0$)
  - for the most extreme joint defaults it gives lower probabilities than tail-dependent models calibrated to the same data
- A static correlation calibrated in calm years
  - exactly the Forbes–Rigobon and DCC lessons of this chapter
- The lesson: joint extremes need a copula with tail dependence, re-estimated when the regime changes

US, Euro-Area and Romanian Equities, 2000–2026 (1/2)



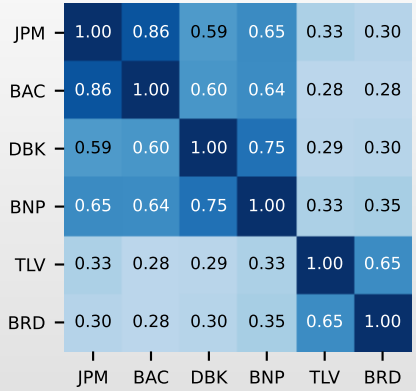
MFM_ch6_integration ▶ Interactive

US, Euro-Area and Romanian Equities, 2000–2026 (2/2)

- The question of this section: is the Romanian market increasingly integrated with the large ones?
  - and can one still diversify across banking systems in crises?
- Weekly returns, 104-week window, common trading days, Friday closes
- Euro Stoxx 50–BET, correlation over each whole subperiod: 0.33 in 2000–2009, 0.45 in 2010–2019, 0.61 since 2020
  - S&P 500–Euro Stoxx 50 in the same periods: 0.84, 0.78, 0.70
- The rolling Euro Stoxx 50–BET correlation jumps in crises and partly reverts afterwards
  - 2008: from 0.31 (end-2007) to 0.65 (end-2009)
  - 2020: from 0.29 (end-2019) to 0.76 (end-2021), then 0.34 in September 2026
- Integration or the Forbes–Rigobon volatility bias? (Seminar C2)
  - in 2008, the rise of the BET–S&P 500 correlation is not significant after the adjustment (Forbes–Rigobon section)
- For a euro-area investor, BET diversifies less on average than in 2000–2009, and least in crises
- Type of evidence: description

MFM_ch6_integration

Banks: Correlations of Weekly Returns, 2010–2026 (1/2)



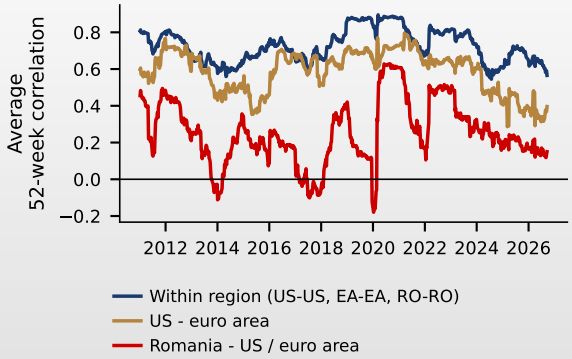
MFM_ch6_banks ▶ Interactive

Banks: Correlations of Weekly Returns, 2010–2026 (2/2)

- The question: do banks move more with their own country than with banks of other regions?
- The banks
  - US: JPM and BAC
  - Euro area: DBK (Deutsche Bank) and BNP (BNP Paribas)
  - Romania: TLV (Banca Transilvania) and BRD (BRD-Groupe Société Générale)
- 871 weeks from 15 January 2010
  - the choice: two of the largest listed banks of each system, with continuous data since 2010
- Correlation of the two banks of each region: US 0.86, euro area 0.75, Romania 0.65
- Cross-region links are weakest for Romania
  - BNP–TLV 0.33, JPM–BRD 0.30
- Bank returns share country-level factors
  - sovereign risk, domestic rates, regulation
  - does this diversification survive crises? (next slide)

MFM_ch6_banks

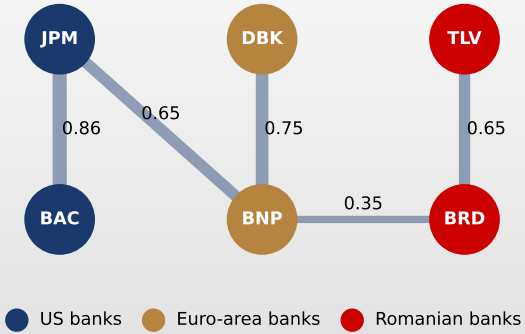
Bank Correlations over Time



- Average pairwise 52-week correlation of weekly returns
  - within regions (legend: US, EA: euro area, RO: Romania)
  - between the US and the euro area, between Romania and the others
- Full period
  - within regions 0.73, US–euro area 0.59, Romania–others 0.25
- March–June 2020: 0.87, 0.69 and 0.56
  - the Romanian link more than doubled in a global shock
  - latest 0.57, 0.40, 0.15
- Diversification across banking systems works in normal times
  - and weakens exactly in systemic events
- Type of evidence: description

MFM_ch6_banks ▶ Interactive

A Correlation Network of Banks



- Minimum spanning tree (Mantegna (1999)) on $d_{ij} = \sqrt{2(1-\rho_{ij})}$
  - the 5 links that join all 6 banks without cycles, with the smallest total distance
- Edges: JPM–BAC (0.86); JPM–BNP (0.65); DBK–BNP (0.75); BNP–BRD (0.35); TLV–BRD (0.65)
  - Romania is attached through BNP–BRD
- The distance $\sqrt{(1-\rho)/2}$ of Chapter 4 gives the same tree
  - both distances increase with $1-\rho$
- The tree is undirected and contemporaneous: it does not identify transmission paths
  - only hypotheses for the network measures of Chapter 18
  - literature in the ▸ Appendix

MFM_ch6_banks ▸ Interactive

Deeper study

Case Study (1/2): Which Events Shake All Markets at Once?

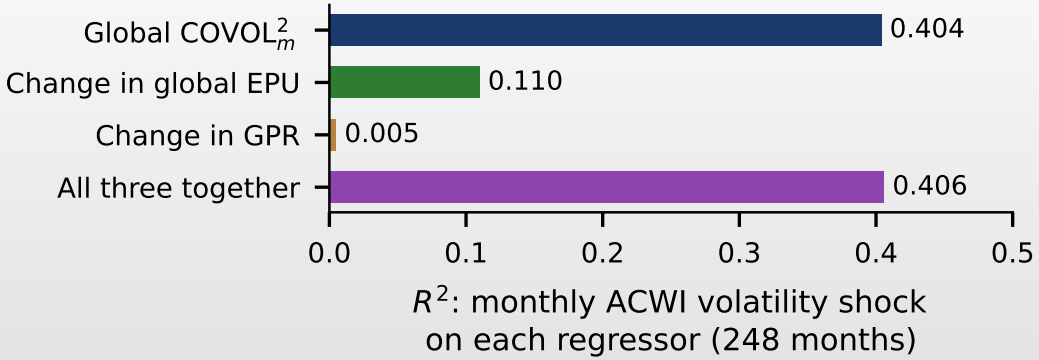
- Questions of Engle & Campos-Martins (2023)
  - which events move the volatilities of almost all countries and asset classes on the same day?
  - how large is this common shock?
- The daily size of the common shock is global COVOL (common volatility)
- The distinct idea: a correct DCC gives residuals e_t with $E_{t-\mathbf{1}}(e_t e_t^\top) = I$
  - but uncorrelated is not independent: the squares can comove
  - $\Psi_t = E_{t-\mathbf{1}}\left[(e_t^{\mathbf{2}} - \iota)(e_t^{\mathbf{2}} - \iota)^\top\right]$
  - e_t: vector of the N DCC residuals of day t
  - $e_t^{\mathbf{2}}$: its element-by-element square
  - ι: vector of ones
  - $E_{t-\mathbf{1}}$: expectation given the information at $t-1$
  - Off-diagonal elements $\Psi_{ij,t} \neq 0$: the squared shocks of assets i and j are large on the same days
  - The cross-sectional analogue of ARCH (autoregressive conditional heteroskedasticity, Chapter 5): correlated squares across assets on one day, not over time
  - Herskovic, Kelly, Lustig & Van Nieuwerburgh (2016) document such a factor in the idiosyncratic volatility of US stocks

Deeper study

Case Study (2/2): One Common Factor Scales Every Variance

- Model: one common factor x_t scales the variance of every residual
  - $e_{it} = \sqrt{s_i x_t + 1 - s_i}\,\varepsilon_{it}$, $\quad \varepsilon_{it}$ i.i.d. $\mathcal{N}(0,1)$, $\quad x_t > 0$, $\quad E_{t-\mathbf{1}}(x_t) = 1$
  - x_t: the squared global COVOL of day t ($x_t > 1$: a common shock day)
  - $s_i \in [0,1]$: loading of asset i, the weight of x_t in its conditional variance $g_{it} = s_i x_t + 1 - s_i = 1 + s_i(x_t - 1)$
- Estimation: only the products $s_i(x_t - 1)$ are identified
  - so the scale must be fixed by a normalisation
  - Equal loadings $s_i = 1$: $\hat{x}_t = \frac{\mathbf{1}}{N}\sum_i e_{it}^{\mathbf{2}}$, the cross-sectional mean of the squared residuals, on the variance scale of the residuals
  - Unequal loadings: alternating maximum likelihood over x_t (one cross-section per day) and s_i (one time series per asset), with $s^\top \underline{s} = 1$ and mean $x_t = 1$
  - Under $s^\top s = 1$ the loadings are relative sizes, not variance shares (equal loadings would be $s_i = 1/\sqrt{N}$)
- Test ξ: the standardised mean over days and pairs of the products $(e_{it}^{\mathbf{2}} - 1)(e_{jt}^{\mathbf{2}} - 1)$
  - under $H_{\mathbf{0}}$: $\Psi_{ij} = 0$, $i \neq j$, it is standard Normal
  - rejecting $H_{\mathbf{0}}$: the squared residuals comove, so there is a common volatility shock
- Data: 47 country ETFs listed on the NYSE, 2000–2021, and 23 US asset-class series
  - ETF: exchange-traded fund; NYSE: New York Stock Exchange
  - details in the Appendix

Realised Shocks against News Indices (1/2)



MFM_ch6_covol

Realised Shocks against News Indices (2/2)

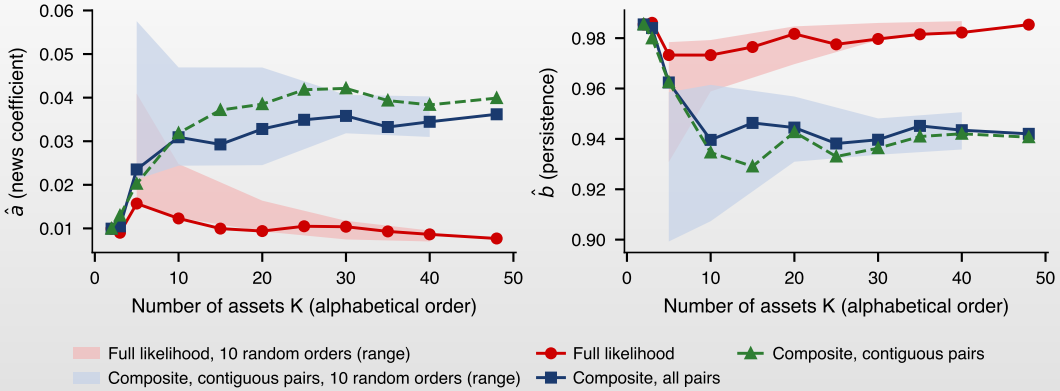
- The question: do news indices explain the global volatility shock as well as COVOL? (248 months, Table 8 of Engle & Campos-Martins (2023))
- The monthly ACWI (MSCI All Country World Index) volatility shock regressed on each index
  - On COVOL_m^2, the monthly mean of $\hat{x}_t$ minus 1: $R^2 = 0.404$ (R^2: share of variance explained)
  - On the change in EPU (economic policy uncertainty, Baker, Bloom & Davis (2016)): 0.110
  - On GPR (geopolitical risk, Caldara & Iacoviello (2022)): 0.005
  - Both news indices added: only 0.406
  - Reading: news indices measure what people fear, COVOL the shocks that hit prices
- Top days in the paper: 9 March 2020 ($\hat{x}_t = 51.4$) and 24 June 2016 (Brexit)
- 28 April 2020 is also a top day, but with a positive mean return (+2.33%)
  - COVOL measures the size of the common shock, not its sign (event table in the Appendix)
- For the risk manager: correlations between the country ETFs rise with COVOL (same Appendix), so diversification weakens on common-shock days
- Type of evidence: association
  - the authors find the squares still correlated after one factor, possibly a second factor

MFM_ch6_covol

DCC in Vast Dimensions: Estimation with Many Assets

- The question of Pakel, Shephard, Sheppard & Engle (2021): can the dynamic DCC parameters (a, b) be estimated on hundreds of assets?
  - the VaR of a large portfolio and hedging need the conditional correlations of all assets
- The claim of the paper, for scalar BEKK (A, B multiples of the identity) and the corrected DCC (cDCC, DCC section): with correlation targeting, the full likelihood gives strongly biased $(\hat{a}, \hat{b})$ as K grows
  - K: the number of assets; the target $\bar{Q}$ has $K(K-1)/2$ correlations, estimated before (a, b)
  - these incidental parameters (needed, but of no direct interest) grow in number with K, and their errors accumulate in the likelihood
  - every evaluation of the likelihood inverts T matrices of size $K \times K$: the cost grows with K^3
- The remedy: the composite likelihood (CL), the sum of the bivariate likelihoods of pairs of assets
  - $\ell_{CL}(a, b) = \sum_{(i,j)\in\mathcal{P}} \ell_{ij}(a, b)$; each pair uses its own 2×2 block of $\hat{\bar{Q}}$
  - $\mathcal{P}$: all $K(K-1)/2$ pairs or only the contiguous pairs $(1, 2), (2, 3), \ldots, (K-1, K)$
  - no large matrix is inverted, and the theory of the paper lets K grow with T
- Data of the paper: the stocks of the S&P 100 and of the S&P 500, daily, 1997–2006

DCC Estimates as the Number of Assets Grows



MFM_ch6_vast_dcc ▶ Interactive

Method and What the Course Adaptation Changes

- Data: all 48 US-listed stocks and ETFs of the course data with a complete history from 18 July 2013
  - 16 stocks and 32 equity, bond and commodity ETFs
  - 3,245 common days, up to 18 September 2026
- Step 1: GARCH(1,1) with the Normal distribution on each series; step 2: DCC with correlation targeting, as in the DCC section
- The cross-section grows in the alphabetical order of the tickers, as in the paper
  - $K = 2, 3, 5, 10, 15, \ldots, 40, 48$
  - for each K: the full likelihood, CL over all pairs and CL over contiguous pairs
  - at $K = 2$ the three estimators coincide: there is a single pair
- Robustness: 10 random orders of the assets, for $K \in \{5, 10, 20, 30, 40\}$ (the bands in the chart)
  - the assets included and the contiguous pairs depend on the order, and the result should not depend on the alphabetical one
- What differs from the paper
  - 48 assets of different types, not only stocks, and the period 2013–2026
  - the standard DCC of the chapter, with the target set to the covariance of the standardised residuals; the paper analyses scalar BEKK and the cDCC of Aielli (2013)
  - so a comparative adaptation on the standard DCC, not a replication of the theoretical result
  - no standard errors for CL, which would need a sandwich matrix

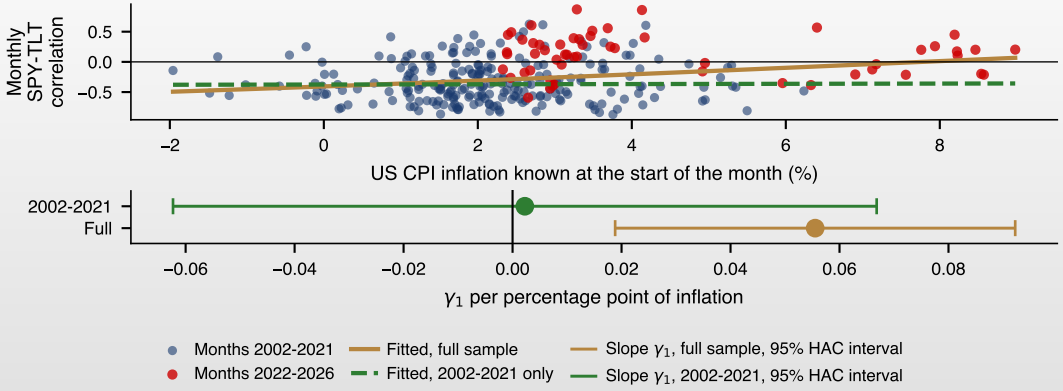
The Result: with the Full Likelihood, Correlations React Much More Slowly

- At $K = 48$, the full likelihood gives $\hat{a} = 0.0077$ (SE 0.0002) and $\hat{b} = 0.9854$
  - CL over all pairs: $\hat{a} = 0.0362$, $\hat{b} = 0.9420$; CL over contiguous pairs: $\hat{a} = 0.0399$, $\hat{b} = 0.9407$
  - SE: standard error from the step-2 Hessian
- For $K \geq 10$, the full-likelihood $\hat{a}$ stays between 0.008 and 0.012, the all-pairs CL $\hat{a}$ between 0.029 and 0.036
  - at $K = 40$, over the 10 random orders: full-likelihood $\hat{a}$ between 0.007 and 0.010, contiguous-pairs CL $\hat{a}$ between 0.031 and 0.040
- The consequence: the mean conditional correlation of the 48 assets (mean over the period 0.34)
  - its standard deviation over time is 0.034 with the full-likelihood estimates and 0.058 with CL
  - its range: 0.28–0.47 against 0.23–0.57
  - the heuristic measure $\ln 0.5 / \ln(a + b)$, roughly the days after which the effect of a shock on the correlation halves: 99 days and 31 days, respectively
  - for a large portfolio, the estimator decides how fast the correlations, hence VaR and hedges, adapt after a shock
- Type of evidence: comparison of procedures
  - the data alone do not tell which (a, b) is the true one
  - the simulations of the paper support the advantage of CL in the models and settings they study; they do not show the absence of bias in our application
  - a single (a, b) for 48 different assets may itself be a misspecification
  - the economic test, out-of-sample hedging and VaR (Chapter 8), remains to be done

An Open Question: Will the Stock–Bond Correlation Stay Positive?

- The question: is the positive stock–bond correlation since 2022 a durable regime?
  - is this regime governed by inflation?
  - can the correlation be forecast?
- Campbell, Pflueger & Viceira (2020): the sign of the comovement follows the inflation–output relation
  - and the monetary-policy regime
  - Molenaar, Sénéchal, Swinkels & Wang (2024) extend the evidence across countries and a long history
- The facts of the chapter: the SPY–TLT correlation is positive since 2022
  - the break is dated August 2020 (Further Topics)
- Hypothesis: $z_m = \operatorname{artanh} \rho_m = \gamma_0 + \gamma_1 \pi_{m-2} + u_m$
  - ρ_m: the monthly realised correlation, Fisher-transformed; u_m: the error
  - π: the annual US CPI (consumer price index) inflation
- H_0: $\gamma_1 = 0$ against H_1: $\gamma_1 > 0$, with a HAC t-statistic
- Refutation: insignificance is not enough; it needs a precisely estimated γ_1 inside a margin fixed in advance (an equivalence test), in periods where inflation varied but the regime did not
- The loop: an AI assistant drafts, the researcher checks
  - AI: the literature map, competing hypotheses, the code, a robustness plan and a referee critique
  - the researcher resolves every DOI (Digital Object Identifier), replicates the chapter's numbers first and fixes the plan before the first run
  - prompts in the Appendix

Mini-Case: Does Inflation Drive the Stock–Bond Correlation? (1/2)



MFM_ch6_ai_inflation ▶ Interactive

Mini-Case: Does Inflation Drive the Stock–Bond Correlation? (2/2)

- Monthly SPY–TLT correlation from daily returns against the inflation known at the start of month m
  - CPI for month $m-1$ is released only in mid-month m, so the latest known value is π_{m-2}
  - current CPI vintage, so a retrospective test, not a real-time one
- Full sample, 288 months: $\hat{\gamma}_1 = 0.056$ per percentage point
  - HAC SE 0.019, $t = 2.97$, $R^2 = 0.047$
  - mean ρ -0.12 when inflation exceeds 3% (89 months), -0.30 otherwise
- 2002–2021 only (233 months, 53 with inflation above 3%): $\hat{\gamma}_1 = 0.002$, $t = 0.07$
  - the 95% HAC interval for γ_1 on 2002–2021 is $[-0.062;\ 0.067]$, which also contains the full-sample slope
- We find no evidence of a positive association before 2022, but the estimate is imprecise
  - so we do not refute the hypothesis either: the interval is too wide for an equivalence test
  - a difference between the slopes would need a model with a regime indicator and an interaction
  - a single regime change cannot separate the inflation level from the policy regime (Campbell, Pflueger & Viceira (2020))
- Type of evidence: association

MFM_ch6_ai_inflation

What the Researcher Checks, and a Project Seed

Checks before any claim

- Look-ahead bias and calendars
  - release calendar and data vintages (the lag handles publication delay, not revisions)
  - align prices on common days before computing returns
- Persistent regressor, few regimes: large t-statistics can arise by chance
  - subsamples, long HAC lags, a placebo replacing inflation by simulated series of the same persistence
- Forbes–Rigobon bias: control for realised volatility
- Multiple testing: fix measures and windows in advance
- Invented references
  - an unresolvable DOI is dropped and logged in `AI_ERRORS.md`

Your turn: a team-project seed

- Does a DCC with a moving target forecast the SPY–TLT correlation better out of sample?
  - the benchmark: the plain DCC
  - the target depends on inflation and the policy rate
  - replicate first: $a + b = 0.986$, $\hat{\bar{Q}}_{12} = -0.248$ and the slope $\hat{\gamma}_1 = 0.056$
- Deliverables
  - a one-page pre-registration: hypothesis, data and tests, fixed before the first run
  - code on the course market data plus CPI
  - a Diebold–Mariano test, which compares the forecast errors of the two models, and a referee-style self-critique
  - the AI-use statement, in `AI_USE.md`

Decision Checklist: Question, Estimator, Uncertainty, Failure

Question	Estimator	Uncertainty	Fails when
Did a correlation change?	Pearson on aligned prices	HAC delta method	Fisher under clustering; daily data across time zones
Covariance for risk today?	EWMA, DCC	bootstrap of both steps	the target $\bar{Q}$ moves with the regime
When did it break?	Wied–Krämer–Dehling on GARCH residuals	Brownian-bridge limit	estimated GARCH residuals (approximate p-value); unmodelled variances
Contagion?	Forbes–Rigobon ρ^*	stationary bootstrap, δ estimated	a common shock or feedback
Joint extremes?	copula on GARCH pseudo-observations	rank sandwich, parametric bootstrap	Gaussian copula ($\lambda = 0$), static copula when dependence moves

- Next: Chapter 7 turns DCC covariances and simulated copulas into portfolio VaR and ES

Check Your Understanding

- Why is the Fisher interval for the SPY–TLT correlation too narrow on daily data?
  - which two features does the HAC interval take into account?
- Why is the same-day correlation of SPY and Euro Stoxx 50 biased towards zero?
  - how do 2-day or weekly returns reduce the bias?
- Which source of uncertainty do step-2 DCC standard errors ignore?
  - how does the bootstrap of both steps restore it?
- A crisis correlation is higher than the calm one
  - what must you compute before calling the rise contagion?
  - which assumption can make the adjustment excessive?
- Five copulas share Kendall's $\tau = 0.5$
  - why can their probabilities of joint crashes still differ?
- Why can the full DCC likelihood fail on dozens of assets?
  - what does the composite likelihood change?

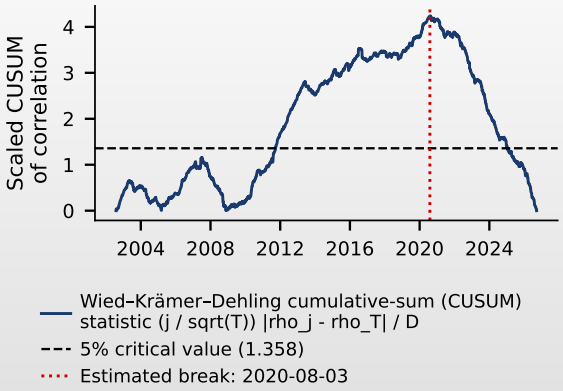
The Practical Decision

- **What I choose**: GARCH(1,1) volatilities with a DCC(1,1) correlation for the daily covariance
  - SPY–TLT: $a + b = 0.986$; CCC is rejected (LR 533.4, bootstrap $p = 0.005$)
  - For joint extremes: a Student t copula on GARCH pseudo-observations (lowest AIC for JPM–BAC)
- **Under which conditions**
  - Prices aligned on common trading days; 2-day or weekly returns across time zones
  - One regime per target $\bar{Q}$: the SPY–TLT correlation breaks in August 2020
  - A second break, in July 2007, is only an exploratory candidate (Further Topics)
  - HAC delta-method intervals for correlations, valid only with finite fourth moments
- **What result would make me drop it**
  - A DCC gain that stays insignificant out of sample, as for the BET hedge (26.2% against 24.7% of the variance removed, one-sided $p = 0.098$): then a constant correlation per regime suffices
  - A target $\bar{Q}$ that moves with the regime, as since 2022 (DCC path above the target on 92% of days): then estimate the target separately per regime

Testing Constant Correlation at an Unknown Date

- The question: has the correlation changed and, if so, when, without choosing the date in advance?
- Tse (2000): a Lagrange multiplier test of H_0: constant correlation
  - against $\rho_{ij,t} = \rho_{ij} + \kappa_{ij}\,\varepsilon_{i,t-1}\varepsilon_{j,t-1}$ (H_0: $\kappa_{ij} = 0$)
  - computed from the fitted CCC model alone, without fitting the alternative
- Picking a break date from a chart and then testing it with a fixed-date Fisher z test
  - is oversized: it rejects a true H_0 more often than 5% of the time
  - the date is itself a function of the data
- Wied, Krämer & Dehling (2012): $Q_T = \max_{2 \le j \le T} \dfrac{j}{\sqrt{T}\,\hat{D}} \left|\hat{\rho}_j - \hat{\rho}_T\right|$
  - Compares the correlation of the first j days, $\hat{\rho}_j$, with the full-sample $\hat{\rho}_T$
  - estimated break date $\hat{j} = \arg\max$: the day with the largest standardised gap
  - $\hat{D}^2$: the delta-method long-run variance of $\hat{\rho}$, Bartlett kernel with bandwidth $\lfloor T^{1/4} \rfloor$ ($\lfloor\cdot\rfloor$: integer part)
  - Under H_0 (constant correlation, means and variances): $Q_T \to \sup_{0 \le s \le 1} |B(s)|$, B a Brownian bridge (a Brownian motion tied to 0 at $s = 0$ and $s = 1$), $s = j/T$
  - 5% critical value 1.358
  - Constant variances are required, so we apply it to GARCH-standardised residuals

Dating the SPY–TLT Correlation Break (1/2)



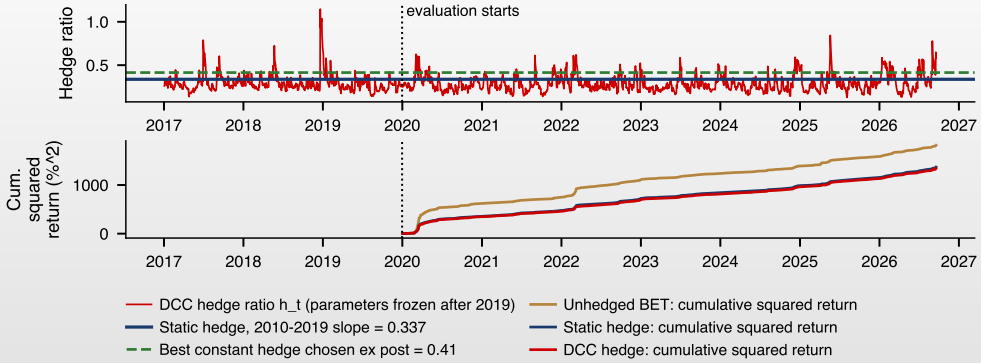
MFM_ch6_corr_inference

Dating the SPY–TLT Correlation Break (2/2)

- Wied–Krämer–Dehling process on the GARCH(1,1)-standardised residuals
  - $T = 6\,073$, bandwidth 8
  - dashed: 5% critical value
- $Q_T = 4.25$ against 1.358: constant correlation rejected ($p < 0.001$)
- Estimated break 3 August 2020, not January 2022
  - residual correlation -0.36 before, 0.07 after
- Binary segmentation, exploratory: the test repeated on each sub-period
  - before the main break: $Q = 3.01$ and a candidate break on 9 July 2007
  - after the main break (August 2020 onwards): $Q = 1.56 > 1.358$, nominal $p = 0.015$
- The limit law of Q_T is proved for directly observed series, not for estimated residuals
  - so the p-values here are approximate
  - tests repeated on segments chosen after the first break do not keep the overall 5% level
- Consequence for DCC: one target $\bar{Q}$ over 2002–2026 may mix different regimes
  - confirmation needs a bootstrap of the whole procedure: GARCH filtering, the tests and the selection of segments

MFM_ch6_corr_inference

Dynamic Hedging Out of Sample: BET with Euro Stoxx 50 (1/2)



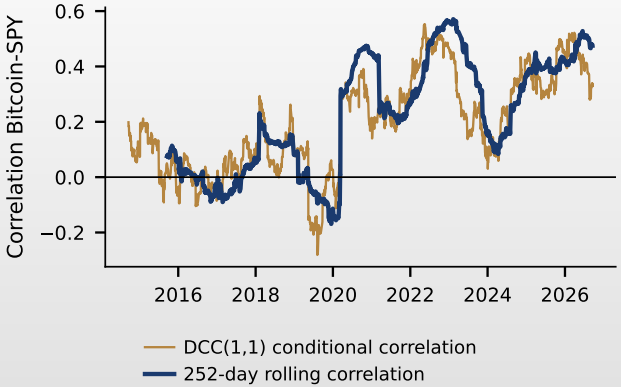
MFM_ch6_corr_inference

Dynamic Hedging Out of Sample: BET with Euro Stoxx 50 (2/2)

- The BET investor hedges by selling h units of Euro Stoxx 50 per unit of BET
  - there are no liquid BET futures
  - the two markets close almost synchronously
- Daily estimation from 5 January 2010 to the end of 2019 ($n = 2\,491$)
- Out-of-sample evaluation in 2020–2026 ($n = 1\,663$) with frozen parameters
  - the top panel starts in 2017
- Static hedge: the slope of BET on Euro Stoxx 50 in the estimation window, $h = 0.337$
- DCC hedge: $h_t = \rho_t \sigma_{S,t} / \sigma_{F,t}$
  - S: BET, F: Euro Stoxx 50; ρ_t and the volatilities from DCC, with $\hat{a} = 0.009$, $\hat{b} = 0.964$
  - mean 0.29, range 0.13 to 0.84 in 2020–2026
- Variance removed in 2020–2026
  - 24.7% with the static hedge, 26.2% with the DCC hedge
  - 25.6% with the best constant hedge chosen ex post ($h = 0.41$)
- The Diebold & Mariano (1995) test: $t = 1.29$, one-sided $p = 0.098$
  - d_t: squared static-hedged return minus squared DCC-hedged return; the t-test on the mean of d_t uses a HAC SE
  - so the DCC gain is small and not significant at 5%

MFM_ch6_corr_inference

Bitcoin and US Equities: DCC Correlation (1/2)



MFM_ch6_dcc_btc_spy

Bitcoin and US Equities: DCC Correlation (2/2)

- Bitcoin and SPY on common trading days since 18 September 2014 ($T = 3\,018$)
  - weekend Bitcoin moves enter the Monday return
- Average DCC correlation 0.04 before 2020, 0.31 since January 2020
  - peak 0.55 on 19 May 2022
- Very persistent: $a = 0.015$, $b = 0.982$, $a + b = 0.9968$
- LR 105.6
  - bootstrap under CCC ($B = 199$): $p = 1/(B + 1) = 0.005$, the smallest attainable value
  - CCC rejected
- From diversifier to risk asset
  - the full-sample correlation (0.24) hides two regimes
- Liu & Tsyvinski (2021), Pele et al. (2023)
  - until 2018 cryptos had small exposures to stock-market factors
  - after 2020 DCC shows a much closer link with equities

MFM_ch6_dcc_btc_spy

Worked Example: The 2008 Crisis Step by Step (1/2): Adjusting

- Step 1, source market: S&P 500
  - calm year: 14 September 2007 – 12 September 2008 ($n = 242$ 2-day returns)
  - crisis: 15 September 2008 – 31 March 2009 ($n = 130$)
- Step 2: S&P 500 crisis variance over calm-year variance, $1 + \delta = 7.54$, so $\delta = 6.54$
- Step 3: the adjustment $\rho^* = \rho_c / \sqrt{1 + \delta(1 - \rho_c^2)}$ for each market (ρ_{calm}, ρ_c: correlations with the S&P 500 in the calm year and in the crisis)
- Euro Stoxx 50: $\rho_{calm} = 0.610$, $\rho_c = 0.804$
  - $\rho^* = 0.804 / \sqrt{1 + 6.54\,(1 - 0.804^2)} = 0.442$
- BET: $\rho_{calm} = 0.216$, $\rho_c = 0.485$
  - $\rho^* = 0.485 / \sqrt{1 + 6.54\,(1 - 0.485^2)} = 0.198$
- Both raw crisis correlations exceed the calm ones by far
  - after the adjustment both fall below them; the next slide tests whether the differences are significant

Worked Example: The 2008 Crisis Step by Step (2/2): Testing

- One-sided Fisher test of $H_0 : \rho_1 \leq \rho_{calm}$, with $\rho_1 = \rho_c$ (raw) or $\rho_1 = \rho^*$ (adjusted)
  - $z = \dfrac{\operatorname{artanh} \rho_1 - \operatorname{artanh} \rho_{calm}}{\sqrt{1/(n_c - 3) + 1/(n_{calm} - 3)}}, \quad p = 1 - \Phi(z)$
- 2-day returns overlap, so the sample sizes must be adjusted
  - We use effective sizes $n_{calm} = 121$ and $n_c = 65$, half of 242 and 130
  - For i.i.d. Normal daily pairs, the Bartlett approximation gives an effective size near $2n/3$, so halving is a cautious convention in that benchmark, not a general bound
- Euro Stoxx 50, raw: $z = \dfrac{\operatorname{artanh} 0.804 - \operatorname{artanh} 0.610}{\sqrt{1/62 + 1/118}} = 2.56$
  - One-sided p-value 0.005; with ρ^*, $z = -1.49$ and $p = 0.932$
- BET
  - $z = 1.97$ raw ($p = 0.024$), $z = -0.12$ adjusted ($p = 0.548$)
- Check: with all 130 and 242 overlapping returns
  - the raw Euro Stoxx 50 statistic would be $z = 3.66$, $p \approx 0.0001$
  - treating overlapping returns as independent overstates the evidence
- The Fisher test treats δ as known; the stationary bootstrap re-estimates it and gives $p = 0.98$ (Euro Stoxx 50) and 0.62 (BET)

Inference for Copula Parameters (1/2): Inverting Kendall's Tau

- **Inversion of $\hat{\tau}$**: $\hat{\tau}$ is a U-statistic, i.e. an average over all pairs of observations
- Its standard error comes from the Hoeffding projection, which approximates the U-statistic by an average of independent terms (Appendix, "Rank Estimators and Their Precision")
- For $\hat{\theta} = g(\hat{\tau})$, the delta method (Genest & Rivest (1993)) multiplies the standard error of $\hat{\tau}$ by the slope $|g'(\hat{\tau})|$
  - g: the inverse of the link $\tau(\theta)$ on the slide "Rank Correlations"; estimates: estimates table
  - Elliptical copulas: $\tau = \frac{2}{\pi} \arcsin \rho$ for *every* ν, so τ identifies ρ, never ν
  - JPM–BAC: $\hat{\tau} = 0.641$ (SE 0.0065), $\hat{\rho}_\tau = \sin(\pi\hat{\tau}/2) = 0.845$ (SE 0.0055)
  - CML $\hat{\rho} = 0.842$ (SE 0.0053): practically the same estimate and precision

Inference for Copula Parameters (2/2): Ranks, Fit and GARCH Filtering

- **CML: the ranks are estimated**, and the inverse-Hessian SE ignores this uncertainty; under misspecification it can be too small or too large
- The Genest, Ghoudi & Rivest (1995) sandwich adds the rank terms
  - With the sandwich, the SE rises by 60% for the Gaussian ρ (S&P 500–Euro Stoxx 50 weekly, from 0.0086 to 0.0138)
  - For JPM–BAC, the SE rises by 17% for the t-copula ρ (from 0.0046 to 0.0053) and by 1% for ν
- Not all of the +60% is the price of ranking
  - If the Gaussian copula were correct, estimating the margins by ranks would multiply the SE by $\sqrt{1+\rho^2}$ (Appendix)
  - At the weekly correlation 0.78 the factor would be 1.27; Monte Carlo at $\rho = 0.8$ gives 1.28
  - The observed factor, 1.60, points to a misspecified Gaussian copula
  - The sandwich stays valid in this case too (fit: Seminar B2)
- **After GARCH filtering** (Chen & Fan (2006))
  - CML on the ranks of estimated residuals has the same limit as on the true innovations
  - Conditions: i.i.d. innovations with copula C_θ and a correctly specified GARCH
  - Then the sandwich stays valid and, unlike DCC, no step-1 correction is needed (Monte Carlo in the Appendix)

Goodness of Fit: The Rosenblatt Test

- If $(U, V) \sim C_\theta$, then $E_1 = U$ and $E_2 = h(V \mid U) = \partial C_\theta(U, V)/\partial u$ are independent uniforms
  - h: the conditional distribution function of V given U; so we test whether (E_1, E_2) are independent uniforms
- Statistic (Genest, Rémillard & Beaudoin (2009)): $S_n = n \int_{[0,1]^2} \left(D_n(e) - e_1 e_2\right)^2 de$
  - $e = (e_1, e_2) \in [0, 1]^2$
  - $D_n(e)$: share of the n pairs with $E_1 \leq e_1$ and $E_2 \leq e_2$
  - $e_1 e_2$: the same share under independence
  - $S_n \geq 0$, large values mean a poor fit
- p-value by parametric bootstrap, because θ is estimated
  - Simulate n pairs from $C_{\hat{\theta}}$, re-rank them into pseudo-observations, re-estimate, recompute $S_n^{(b)}$
  - $p = (1 + \#\{S_n^{(b)} \geq S_n\})/(B + 1)$
  - For JPM–BAC $B = 200$ (100 for Student t)
  - The smallest attainable p-value is therefore $1/201 \approx 0.005$ ($1/101 \approx 0.010$ for t)
- On GARCH residuals, Rémillard (2017) shows that rank-based tests keep their limit
  - and the parametric bootstrap applies
- With $T = 4\,202$ the test has high power against small departures
  - a rejection does not say which feature is missing

Is the Tail Asymmetry Real?

- Schmidt & Stadtmüller (2006): $\hat{\lambda}_L = \frac{n}{k}\,\hat{C}_n\left(\frac{k}{n}, \frac{k}{n}\right)$, the chart estimator with $q = k/n$
  - n: sample size
  - k: number of observations in the tail
  - $\hat{C}_n$: empirical copula
  - A larger k lowers the variance but raises the bias (conditions in the Appendix)
- $H_0 : \lambda_L(q) = \lambda_U(q)$ at a fixed q (threshold probabilities, not the limits)
- Bootstrap of residual pairs, re-ranked, $B = 999$

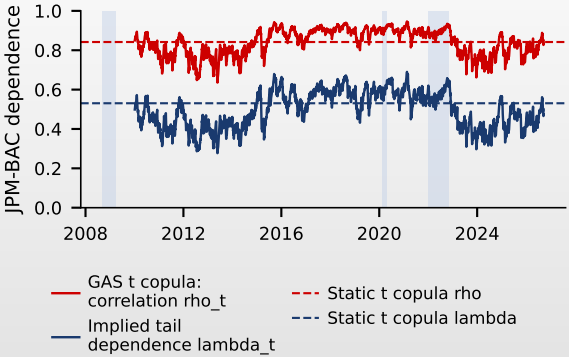
q	k	$\hat{\lambda}_L - \hat{\lambda}_U$	95% interval, i.i.d.	95% interval, blocks
5%	210	0.062	[−0.014, 0.148]	[−0.014, 0.138]
1%	42	0.095	[−0.095, 0.286]	[−0.119, 0.286]

- The block version allows serial dependence in the tail indicators
  - stationary bootstrap, mean block 1.6 and 1.0 days, Politis & White (2004) rule
- All intervals contain zero: the asymmetry is not detected statistically
  - the test only checks that the two threshold probabilities are equal, not their level nor the whole copula
  - so it does not validate the t specification, which the global goodness-of-fit test rejects

The Dynamic Copula, Intuitively

- The static copula assumes the same link between JPM and BAC over the whole period, although the banks move more tightly together in crises
- The idea: let the copula correlation, ρ_t, change from one day to the next, as GARCH lets the volatility change
- The updating rule (GAS): each day we ask whether today's pair $(\hat{u}_t, \hat{v}_t)$ would have been more likely with a higher or a lower correlation, and move the correlation a little in that direction
  - a day on which both banks have extreme percentiles in the same direction pushes the correlation up; a day with opposite moves pushes it down
  - A says how much we react to today (0.025, little), B how much we keep from the past (0.9926, much): the correlation moves slowly
  - tanh keeps ρ_t between -1 and 1
- The result: ρ_t between 0.64 (2013) and 0.95 (2021); the tail dependence λ_t between 0.28 and 0.69, against 0.53 in the static copula
- Why it matters: part of the joint extremes of the static copula came from mixing calm periods with crises ($\hat{\nu}$ rises from 4.24 to 5.35)

A Dynamic t Copula for JPM and BAC (1/2)



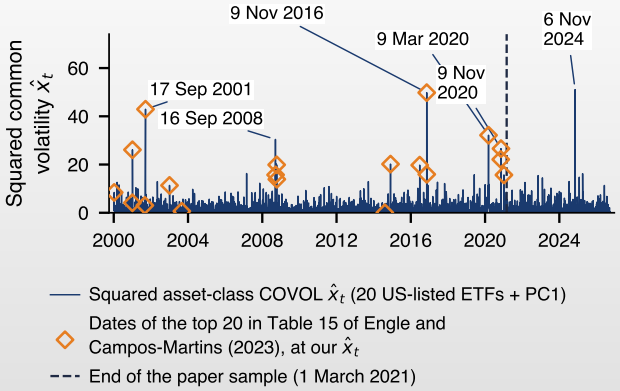
MFM_ch6_copula_inference

A Dynamic t Copula for JPM and BAC (2/2)

- GAS (generalised autoregressive score) t copula, Creal, Koopman & Lucas (2013), same pseudo-observations
  - $\rho_t = \tanh f_t$, $f_{t+1} = \omega + A s_t + B f_t$
  - f_t: the transformed correlation, on the whole real line; ω: constant
  - $s_t = \partial \ln c(\hat{u}_t, \hat{v}_t; \rho_t, \nu)/\partial f_t$: the score, unscaled; f_t moves in the direction that raises the likelihood of today's observation
  - A: reaction to the score; B: persistence
- Estimates: $A = 0.025$, $B = 0.9926$, $\nu = 5.35$ (static copula $\nu = 4.24$)
- ρ_t from 0.64 (23 May 2013) to 0.95 (18 February 2021)
  - λ_t (the tail dependence implied by ρ_t and ν) from 0.28 to 0.69, against a static 0.53
  - shaded: autumn 2008 – March 2009, February–April 2020, January–October 2022
- Once ρ moves, $\hat{\nu}$ rises (thinner tails)
  - part of the static tail dependence was time variation in dependence
- LR against the static t copula: 245.5
  - under $A = 0$, B is not identified, so the χ^2 limit does not apply
  - the p-value comes from a bootstrap (Seminar B6)
- Time variation is one explanation consistent with the rejection of every static family

MFM_ch6_copula_inference

Asset-Class COVOL on the Course ETFs, 2000–2026 (1/2)



MFM_ch6_covol

Asset-Class COVOL on the Course ETFs, 2000–2026 (2/2)

- 20 US-listed ETFs and their PC1 (first principal component)
  - 4 January 2000 – 18 September 2026, 6 717 days
  - AR(1) mean (first-order autoregressive) with PC1, GARCH(1,1), then the alternating estimation
- Diamonds: the 20 dates of Table 15 of Engle & Campos-Martins (2023)
  - placed at our replication's value on that day
  - days with the market closed: at the next trading day
- Top days up to 1 March 2021, the end of the paper sample
  - 9 November 2016 ($\hat{x}_t = 49.8$): the 2016 US election
  - 17 September 2001 (42.9): the reopening after 9/11
  - 9 March 2020 (32.2): the COVID-19 crash
- 13 of the 20 largest days of our replication before March 2021
  - are among the 20 dates of Table 15: the replication recovers most of the paper's events
- With $s^\top s = 1$ and mean $\hat{x}_t = 1$ the scale depends on the panel
  - replication values are not comparable with the paper's: compare rankings, not levels
- Largest in the full sample: 6 November 2024 (51.1)
  - the day after the 2024 US election
  - mean ETF return 0.95%: a large common shock on an up day

MFM_ch6_covol

Further Study — Courses, Code and Data

Advanced courses on quantinar.com

- Measuring Statistical Risk
- Multivariate Statistical Analysis
- Minimum Spanning Tree
- FRM: Financial Risk Meter

Code, data and textbooks

- All charts and tables: github.com/danpele/MFM/tree/main/Quantlets/Ch_06
- Textbooks: Franke, Härdle & Hafner (2019); Härdle & Simar (2019); Nelsen (2006)

Back

Inference on a Correlation, the Delta Method

- $\hat{\rho} = g(\hat{m})$, with $\hat{m}$ the sample mean of $m_t = (x_t,\ y_t,\ x_t^2,\ y_t^2,\ x_t y_t)$
  - g: the function that turns these five moments into the correlation; ∇g: its gradient
- Delta method: $\sqrt{T}(\hat{\rho} - \rho) \to N\big(0,\ \nabla g^\top \Omega\, \nabla g\big)$
  - Ω: the long-run covariance of m_t, estimated by HAC with a Bartlett kernel and L lags (Newey & West (1987))
- The same form with the daily contributions ψ_t (influence function): $\hat{\rho} - \rho \approx \frac{1}{T}\sum_t \psi_t$ of the main text
- Conditions: stationarity, finite moments of order higher than 4 of (x_t, y_t) and weak dependence (mixing)
  - with infinite fourth moments the Normal limit fails
  - nor does the ordinary bootstrap automatically fix missing moments: this applies to the Forbes–Rigobon bootstrap and to tests on squared losses too
- The HAC interval, built on the Fisher scale: $\tanh\big(\operatorname{artanh}\hat{\rho} \pm 1.96\,\widehat{\mathrm{SE}}_{\mathrm{HAC}}/(1-\hat{\rho}^2)\big)$
  - with $\widehat{\mathrm{SE}}_{\mathrm{HAC}} = \sqrt{\nabla\hat{g}^\top \hat{\Omega}\, \nabla\hat{g}/T}$; $1/(1-\hat{\rho}^2)$ is the derivative of artanh (delta method)
- The classic Fisher interval replaces $\widehat{\mathrm{SE}}_{\mathrm{HAC}}/(1-\hat{\rho}^2)$ by $1/\sqrt{T-3}$
- Example: the last 252-day rolling window of the main text, ending on 17 September 2026
  - $\hat{\rho} = 0.245$, Fisher interval [0.13, 0.36], HAC interval [0.07, 0.41]
  - largest HAC-to-Fisher width ratio over all windows: 2.02

Back

The VEC and BEKK Models

- **VEC** (Bollerslev, Engle & Wooldridge (1988)): in the equation of the main text, c is a constant vector
- VEC: $N(N+1)/2 \cdot [1 + N(N+1)]$ parameters: the vector c plus the matrices A and B, each with $[N(N+1)/2]^2$ elements
- **BEKK** (Engle & Kroner (1995)): positive definite by construction if $CC^\top$ is positive definite
- BEKK: $N(N+1)/2 + 2N^2$ parameters (C triangular, A and B of size $N \times N$), hard to interpret
  - and the likelihood is flat in high dimension
- DCC: $3N + 2$ parameters, univariate GARCH for each variance plus a and b; the target $\bar{Q}$ is estimated separately

Back

Two-Step Estimators, the Correct Covariance

- The question: how do the standard errors of (a, b) include the estimation error of step 1?
- Stacked estimating equations
  - $\sum_t s_{1t}(\theta_1) = 0$ (step 1: GARCH and, with correlation targeting, vech $\bar{Q}$)
  - $\sum_t s_{2t}(\theta_1, \theta_2) = 0$ (step 2, $\theta_2 = (a, b)$)
- s_{jt}: score of step j on day t; $A_{jk} = \mathbb{E}\left[\partial s_{jt} / \partial \theta_k^\top\right]$
  - A_{21} measures how the correlation score reacts to the GARCH parameters
- Newey & McFadden (1994), Section 6: $\sqrt{T}(\hat{\theta}_2 - \theta_2) \to N(0, V_2)$
  - $V_2 = A_{22}^{-1}\, \Omega\, A_{22}^{-1\top}$, with Ω the long-run variance of the adjusted score $u_t = s_{2t} - A_{21} A_{11}^{-1} s_{1t}$
- The usual step-2 standard errors use only s_{2t}
  - they drop $A_{21} A_{11}^{-1} s_{1t}$, the GARCH estimation error carried into step 2
- The target moment is not a martingale difference: $\mathbb{E}_{t-1}[\varepsilon_t \varepsilon_t^\top - \bar{Q}] = R_t - \bar{Q}$
  - so Ω is estimated by HAC
  - the contemporaneous variance of u_t suffices only if the whole system, the target moment included, is a martingale difference
- Engle & Sheppard (2001) derive V_2 for DCC
- A simple alternative bootstraps the whole two-step procedure (next slide)

Back

Bootstrap of Both DCC Steps

- Goal: standard errors for (a, b) that include the uncertainty of both steps; parametric residual bootstrap, $B = 199$
  - decorrelate the residuals with the Cholesky factor of R_t and resample them
  - simulate DCC innovations with $(\hat{a}, \hat{b}, \bar{Q})$
- Step 2 only: re-estimate (a, b) on the simulated innovations
- Both steps: rebuild returns with the fitted GARCH
  - re-estimate GARCH, $\bar{Q}$ and (a, b) on the same innovations
- The CCC-null bootstrap of the LR test uses the same machinery with $a = b = 0$
  - constant-correlation innovations, simulated GARCH returns, both steps re-estimated

Back

Robert Engle and the DCC Model



Robert F. Engle at the Western Economic Association International conference, Santiago, 2017

Photo: Econterms (2017); CC BY-SA 4.0; Wikimedia Commons

- Nobel Memorial Prize in Economic Sciences 2003 for ARCH models (Chapter 5)
- Engle (2002): "Dynamic Conditional Correlation"
  - Keeps the flexibility of univariate GARCH for each variance
  - Adds a simple, estimable law of motion for correlations
- Refinement: the corrected DCC of Aielli (2013), which makes $\bar{Q}$ a consistent target

Back

Contagion Tests beyond Forbes–Rigobon

- Corsetti, Pericoli & Sbracia (2005): with a common factor the Forbes–Rigobon adjustment over-corrects
  - biasing the test towards "no contagion"
  - they give a factor-model version
- Rigobon (2003): identification through heteroskedasticity
  - the variance shift between regimes identifies the structural coefficient β, an instrumental-variables idea
- The common-factor critique applies to the 2008 crisis as well (Seminar A1, B3)

Back

Copula Densities and Constructions

- Gaussian density: $c = \frac{1}{\sqrt{1-\rho^2}} \exp\left(- \frac{\rho^2(x^2+y^2)-2\rho xy}{2(1-\rho^2)} \right)$
  - $x = \Phi^{-1}(u)$, $y = \Phi^{-1}(v)$
- t copula: the tail dependence is large for small ν
- Archimedean generators φ in $C(u,v) = \varphi^{-1}(\varphi(u) + \varphi(v))$
  - Clayton $\varphi(t) = (t^{-\theta} - 1)/\theta$
  - Gumbel $\varphi(t) = (-\ln t)^{\theta}$
  - Frank $\varphi(t) = -\ln \frac{e^{-\theta t}-1}{e^{-\theta}-1}$
- Asymmetric tails
  - Clayton for joint crashes of returns; Gumbel, applied to losses (returns with the sign reversed), describes the same crashes
- The 180° rotation (the survival copula) $\bar{C}(u,v) = u + v - 1 + C(1-u, 1-v)$ moves the tail to the opposite corner

Back

Copula Estimates for JPM and BAC, Full Table

Copula	$\hat{\theta}$ or $\hat{\rho}$	$\hat{\nu}$	ℓ	ΔAIC	λ_L	λ_U	fit p
Gaussian	0.833	–	2 486.0	309.3	0.000	0.000	0.005
Student t	0.842	4.24	2 641.7	0.0	0.531	0.531	0.010
Clayton	2.384	–	2 116.7	1 048.0	0.748	0.000	0.005
Gumbel	2.587	–	2 412.2	456.9	0.000	0.693	0.005
Frank	9.152	–	2 390.8	499.8	0.000	0.000	0.005

- CML on pseudo-observations of GARCH residuals
- ℓ: maximised log-likelihood; λ_L, λ_U: tail dependence implied by the parameters
- ΔAIC relative to the best model
- Goodness-of-fit p: Rosenblatt test with parametric bootstrap

Back

Rank Estimators and Their Precision

- $\hat{\tau}$ is a U-statistic
  - its Hoeffding projection gives $\sqrt{n}(\hat{\tau} - \tau) \to N\big(0,\ 16\mathrm{Var}[2C(U,V) - U - V]\big)$
- Plug-in: $\widehat{\mathrm{SE}}(\hat{\tau}) = 4\,\mathrm{sd}\big(2\hat{C}_n(\hat{U}_i, \hat{V}_i) - \hat{U}_i - \hat{V}_i\big)/\sqrt{n}$, with $\hat{C}_n$ the empirical copula
- Inverting $\hat{\tau}$ gives a robust starting value for the likelihood
- CML sandwich (Genest, Ghoudi & Rivest (1995)): $\sqrt{n}(\hat{\theta} - \theta) \to N(0,\ A^{-1}\Sigma A^{-1})$
  - A: the Fisher information of the copula
- $\Sigma = \mathrm{Var}\big[\dot{\ell}_\theta(U,V) + W_1(U) + W_2(V)\big]$, $\dot{\ell}_\theta$ the score of $\ln c$
- $W_1(u) = \int\!\!\int 1\{u \le u'\}\,\partial_u \dot{\ell}_\theta(u', v')\,dC(u', v')$
  - the correction for estimating the margins by ranks; W_2 likewise in v
- Gaussian copula: variance $(1-\rho^2)^2$ for CML, against $(1-\rho^2)^2/(1+\rho^2)$ with known margins
- The ratio of the standard errors is therefore $\sqrt{1+\rho^2}$, the factor used in Further Topics

Back

Chen and Fan (2006), GARCH Filtering and the Limit

- Model: $r_{it} = \mu_i + \sigma_{it}\varepsilon_{it}$ with GARCH σ_{it}
- Innovations $(\varepsilon_{1t}, \varepsilon_{2t})$ i.i.d. with copula C_θ
- Chen & Fan (2006): CML on the ranks of the *estimated* residuals $\hat{\varepsilon}_{it}$
  - has the same limit distribution as CML on the ranks of the true ε_{it}
- Monte Carlo: 500 samples, $T = 2\,000$
  - Gaussian copula $\rho = 0.8$, Student t(5) innovations, GARCH as fitted to JPM and BAC
  - Standard deviation of $\hat{\rho}$: 0.0081 (true innovations) and 0.0082 (GARCH residuals)
  - Mean sandwich SE 0.0081, mean inverse-Hessian SE 0.0063
- Conclusion: GARCH filtering barely changes the precision; the sandwich estimates it correctly, the inverse Hessian understates it

Back

Nonparametric Tail Dependence, Conditions

- Schmidt & Stadtmüller (2006): the estimator $\hat{\lambda}_L$ of the main text is consistent and asymptotically Normal if
  - $k \to \infty$, $k/n \to 0$ and the tail copula exists
  - a second-order bias condition holds: $\sqrt{k}\, A(n/k) \to 0$
  - $A(\cdot)$: the rate at which $C(q,q)/q$ approaches λ_L
- The i.i.d. bootstrap of residual pairs assumes a static innovation copula
- The stationary block bootstrap keeps the serial dependence of the difference
  - $1\{U \le q,\, V \le q\} - 1\{U > 1-q,\, V > 1-q\}$
  - its mean divided by q estimates $\lambda_L - \lambda_U$

Back

Dependence in High Dimension and over Time

- DCC for large N: the full likelihood is slow and biased
- Remedies
  - composite likelihood over pairs of assets (Pakel, Shephard, Sheppard & Engle (2021))
  - DCC-NL (Dynamic Conditional Correlation with NonLinear shrinkage), nonlinear shrinkage of the target $\bar{Q}$ (Engle, Ledoit & Wolf (2019); the shrinkage of Chapter 4)
- Vine (pair-copula) constructions (Aas, Czado, Frigessi & Bakken (2009)): $d(d-1)/2$ bivariate copulas, some conditional
  - in nested trees, each with its own family and tails
- Hierarchical Archimedean copulas (Okhrin, Okhrin & Schmid (2013)): groups nested by the strength of dependence
  - e.g. banks within a country and countries within a region
- Dynamic copulas: score-driven parameters (Creal, Koopman & Lucas (2013); the GAS model of Further Topics)
  - extend to many assets through factor copulas
- Oh & Patton (2018) measure time-varying systemic risk in credit default swap spreads
- Conditional copulas: Patton (2006)

Back

From Correlations to Systemic Risk

- Correlation networks are the first step towards the systemic-risk measures of Chapter 18
  - Connectedness from Granger causality: Billio, Getmansky, Lo & Pelizzon (2012)
  - Spillover indices from variance decompositions: Diebold & Yilmaz (2009)
  - Connectedness from quantile regressions: CoVaR (Conditional Value-at-Risk of the system given an institution in distress), Adrian & Brunnermeier (2016)
  - Tail-event networks: the Financial Risk Meter applied to EU (European Union) energy firms, Toma, Dima, Grecu, Pele & Lessmann (2026)
- Oil price shocks and the US stock market: Kilian & Park (2009)
- Energy prices and EU stock markets, Grecu, Cramer, Pele & Lessmann (2025)
  - dependence across regions and industrial groups
- Correlation networks inherit all the problems of correlations
  - non-linearity, time variation, tails

Back

COVOL: Data and Estimation

- **Model** (Eq. 8–11 of Engle & Campos-Martins (2023)): for $i \neq j$, $\Psi_{ij,t} = s_i s_j v_t$, $v_t = \mathrm{Var}_{t-\mathbf{1}}(x_t)$
  - e_{it}: the standardised residual of asset i; $\Psi_{ij,t} = \mathrm{Cov}_{t-\mathbf{1}}(e_{it}^{\mathbf{2}}, e_{jt}^{\mathbf{2}})$
  - so Ψ has one factor and its first principal component gives starting values (Eq. 12–13)
  - Equal loadings (before the normalisation $s^\top s = 1$): $g_{it} = x_t$, and the cross-sectional mean $\hat{x}_t = \frac{\mathbf{1}}{N} \sum_i e_{it}^{\mathbf{2}}$ of the main text is the maximum-likelihood estimate (Eq. 17)
- **Estimation** (Section 5.3): maximise $-\frac{\mathbf{1}}{\mathbf{2}} \sum_{i,t} \left[\ln g_{it} + e_{it}^{\mathbf{2}} / g_{it} \right]$
  - $g_{it} = s_i x_t + 1 - s_i$, alternating x_t and s_i until convergence
  - After each step: $s^\top s = 1$, mean of x_t equal to 1 (Remark 2)
  - Test ξ (Eq. 27): standardised mean of $(e_{it}^{\mathbf{2}} - 1)(e_{jt}^{\mathbf{2}} - 1)$ over the $N(N-1)/2$ pairs, $\mathcal{N}(0,1)$ under $H_{\mathbf{0}}$
- **Data** (Section 7.1, Table B.1): 47 country ETFs listed on the NYSE
  - daily USD (US dollar) log returns, 2 June 2000 – 1 March 2021
  - same closing time, so no asynchrony; unbalanced panel
  - Mean: AR(1) plus two factors, the ACWI and PC1 of the ETF returns
  - GARCH(1,1) variances
  - COVOL fitted to the 47 residual series and the two factors
  - Also: 23 US asset-class series (Section 9) and Monte Carlo studies of the estimator and of the test (Sections 5.4, 6.1)

Back

COVOL: Key Results

Day	$\hat{x}_t$	$\bar{r}_t$ (%)	Event (label of the paper)
9 March 2020	51.416	−9.650	COVID-19: Italy locked down; oil price war
24 June 2016	49.190	−7.270	Brexit
5 August 2019	39.457	−2.574	Trade war: China devalues its currency
28 April 2020	38.883	2.334	COVID-19: tech share sell-off
17 September 2001	30.077	−8.630	NYSE reopens after 9/11

$\bar{r}_t$: mean return across the country ETFs; top five of the twenty largest values in Table 6 of Engle & Campos-Martins (2023)

- Loadings (Table 7): PC1 0.261 and ACWI 0.250 lead
  - France 0.217, Spain 0.200 against New Zealand 0.065, Poland 0.080
  - factors load more than idiosyncrasies, so correlations rise with COVOL (Eq. 41)
- Regressions on COVOL_m^2, the monthly mean of $\hat{x}_t$ minus 1
  - ACWI squared residual: 0.326 (SE 0.025), $R^2 = 0.404$ (Table 8)
  - mean pairwise ETF correlation: +0.053 per unit (SE 0.009, Table 10, column 4)
  - change in VIX (Cboe index of the implied volatility of the S&P 500): 4.034 (SE 0.249), $R^2 = 0.517$ (Table 12)

Back

COVOL: Replication Protocol

- **Replicate** the asset-class version (Section 9)
  - with the US-listed ETFs from the course market data, which all close at the same time
- The ETFs
  - 11 sector ETFs, Russell 1000 growth and value, Russell 2000
  - developed and emerging markets outside the US
  - gold, long Treasuries, corporate bonds, crude oil
- One factor, PC1, on an unbalanced panel
- $\hat{x}_t$ by Eq. 17, then the iteration of Section 5.3
- The ξ test before and after dividing e_{it}^2 by $\hat{g}_{it}$ (Section 7.2)
- Table 15 of the paper ranks among the top four days
  - 9 November 2016 (43.04) and 9 March 2020 (40.47)
- The dollar and the Russell 2000 growth/value split are not in the course data, so the replication panel differs from the paper's

Back

COVOL: Limitations and the Data after 2021

- Limitations stated by the authors
  - after dividing by $\hat{g}_{it}$ the squares are still correlated, possibly a second factor (Section 7.2)
  - multifactor COVOL is open (Section 9)
- With $s^\top s = 1$ the loadings shrink as N grows
  - R_s^2: the fit between estimated and true loadings
  - R_s^2 is 0.640 for $N = 50$, $T = 1000$, against 0.924 for $N = 10$ (Table 2)
- The test loses power with unequal loadings, so the authors advise $T \geq 1000$
- After the sample of the paper
  - 3 April 2025 ($\hat{x}_t = 16.2$, mean return -3.67%), after the US tariff announcement
  - 10 November 2022 ($\hat{x}_t = 15.4$, $+4.56\%$), a lower-than-expected US inflation print
- Open question (Seminar C3): do elections, tariffs and inflation news share one factor?
- Open question: does COVOL_m^2 explain the monthly DCC correlations beyond the change in VIX?

The Discovery Loop, Reusable Prompts (1/2)

- Tools are interchangeable
  - an AI assistant (Claude, ChatGPT, Gemini, GitHub Copilot)
  - literature engines (Semantic Scholar, Elicit, Google Scholar)
- **Literature**
  - *"List studies since 2015 on the drivers of the stock–bond correlation, with DOI, sample, measure of correlation, identification strategy and main estimate."*
  - Resolve every DOI on Crossref and check the estimate in the paper, not in the summary
- **Hypotheses**
  - *"Give competing explanations for the 2022 sign change (inflation level, monetary-policy regime, term-premium shock, a one-off tightening) and what each predicts for ρ before 2022 and out of sample."*
- **Data and code**
  - *"Write Python that aligns SPY and TLT adjusted prices from the course market data on common trading days, computes monthly realised correlations and merges US CPI inflation published before each month starts."*
  - Replicate first: -0.40 (2002–2021) and 0.11 (2022–2026)

Back

The Discovery Loop, Reusable Prompts (2/2)

- **Robustness**
  - *"Write a pre-registered plan: the pre-2022 subsample, core CPI instead of headline, weekly instead of daily returns, DCC instead of realised correlation, control for realised volatility, HAC lag 12."*
  - Placebo: replace π by simulated random-walk series of the same persistence
  - A persistent regressor can produce large t-statistics by chance
- **Critique**
  - *"Act as a referee: is γ_1 identified by a single episode, and would the Forbes–Rigobon volatility bias explain it? Give the test that would expose each."*
- **Report**: every inflation measure, window and frequency tried
  - subsample estimates next to the full sample
  - any specification chosen after seeing results is exploratory and reported as such
- Code of the mini-case: Quantlets/Ch_06/ai_discovery_case.py
- Inflation: US consumer price index for all urban consumers
  - source: U.S. Bureau of Labor Statistics

Back

Acronyms used in this material (1/4)

- **ACWI**: All Country World Index (MSCI)
- **AI**: Artificial Intelligence
- **AIC**: Akaike Information Criterion
- **AR**: AutoRegressive (model)
- **ARCH**: AutoRegressive Conditional Heteroskedasticity
- **BAC**: Bank of America Corporation (NYSE ticker)
- **BEKK**: Baba–Engle–Kraft–Kroner (multivariate GARCH model)
- **BET**: Bucharest Exchange Trading (index)
- **BIC**: Bayesian Information Criterion
- **BNP**: BNP Paribas SA (Euronext Paris ticker; Banque Nationale de Paris)
- **BRD**: BRD–Groupe Société Générale (BVB ticker)
- **CCC**: Constant Conditional Correlation (model)

Acronyms used in this material (2/4)

- **CDO**: Collateralised Debt Obligation
- **CL**: Composite Likelihood
- **CML**: Canonical Maximum Likelihood
- **CoVaR**: Conditional Value-at-Risk (of the system, given an institution)
- **COVID-19**: COronaVIrus Disease 2019
- **COVOL**: COmmon VOLatility (global COVOL: the size of the volatility shock common to all assets)
- **CPI**: Consumer Price Index
- **DBK**: Deutsche Bank AG (Xetra ticker)
- **DCC**: Dynamic Conditional Correlation
- **DCC-NL**: Dynamic Conditional Correlation with NonLinear shrinkage
- **DOI**: Digital Object Identifier
- **EA**: Euro Area
- **EPU**: Economic Policy Uncertainty (index)
- **ES**: Expected Shortfall
- **ETF**: Exchange-Traded Fund

Acronyms used in this material (3/4)

- **EU**: European Union
- **EWMA**: Exponentially Weighted Moving Average
- **FRM**: Financial Risk Meter
- **GARCH**: Generalised AutoRegressive Conditional Heteroskedasticity
- **GAS**: Generalised Autoregressive Score (model)
- **GPR**: GeoPolitical Risk (index)
- **HAC**: Heteroskedasticity and Autocorrelation Consistent (standard errors)
- **i.i.d.**: independent and identically distributed
- **IFM**: Inference Functions for Margins
- **JPM**: JPMorgan Chase & Co. (NYSE ticker)
- **LR**: Likelihood Ratio (test)
- **MSCI**: Morgan Stanley Capital International (index provider)

Acronyms used in this material (4/4)

- **NYSE**: New York Stock Exchange
- **PC1**: first Principal Component
- **SE**: Standard Error
- **SPY**: SPDR S&P 500 ETF Trust
- **TLT**: iShares 20+ Year Treasury Bond ETF
- **TLV**: Banca Transilvania (BVB ticker)
- **US**: United States
- **USD**: United States Dollar
- **VaR**: Value-at-Risk
- **VEC**: VECtorised multivariate GARCH model (vech of the covariance matrix)
- **VIX**: Cboe Volatility Index (implied volatility of the S&P 500)

References (1/5)

- Aas, K., Czado, C., Frigessi, A., Bakken, H. (2009). Pair-copula constructions of multiple dependence. *Insurance: Mathematics and Economics*, 44(2), 182–198.
- Adrian, T., Brunnermeier, M.K. (2016). CoVaR. *American Economic Review*, 106(7), 1705–1741.
- Aielli, G.P. (2013). Dynamic conditional correlation: on properties and estimation. *Journal of Business & Economic Statistics*, 31(3), 282–299.
- Ang, A., Chen, J. (2002). Asymmetric correlations of equity portfolios. *Journal of Financial Economics*, 63(3), 443–494.
- Baker, S.R., Bloom, N., Davis, S.J. (2016). Measuring economic policy uncertainty. *Quarterly Journal of Economics*, 131(4), 1593–1636.
- Billio, M., Getmansky, M., Lo, A.W., Pelizzon, L. (2012). Econometric measures of connectedness and systemic risk in the finance and insurance sectors. *Journal of Financial Economics*, 104(3), 535–559.
- Bollerslev, T. (1990). Modelling the coherence in short-run nominal exchange rates: a multivariate generalized ARCH model. *Review of Economics and Statistics*, 72(3), 498–505.
- Bollerslev, T., Engle, R.F., Wooldridge, J.M. (1988). A capital asset pricing model with time-varying covariances. *Journal of Political Economy*, 96(1), 116–131.
- Caldara, D., Iacoviello, M. (2022). Measuring geopolitical risk. *American Economic Review*, 112(4), 1194–1225.
- Campbell, J.Y., Pflueger, C., Viceira, L.M. (2020). Macroeconomic drivers of bond and equity risks. *Journal of Political Economy*, 128(8), 3148–3185.
- Chen, X., Fan, Y. (2006). Estimation and model selection of semiparametric copula-based multivariate dynamic models under copula misspecification. *Journal of Econometrics*, 135(1–2), 125–154.
- Clayton, D.G. (1978). A model for association in bivariate life tables and its application in epidemiological studies of familial tendency in chronic disease incidence. *Biometrika*, 65(1), 141–151.
- Corsetti, G., Pericoli, M., Sbracia, M. (2005). 'Some contagion, some interdependence': more pitfalls in tests of financial contagion. *Journal of International Money and Finance*, 24(8), 1177–1199.

References (2/5)

- 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.
- Diebold, F.X., Mariano, R.S. (1995). Comparing predictive accuracy. *Journal of Business & Economic Statistics*, 13(3), 253–263.
- Diebold, F.X., Yilmaz, K. (2009). Measuring financial asset return and volatility spillovers, with application to global equity markets. *Economic Journal*, 119(534), 158–171.
- Embrechts, P., McNeil, A.J., Straumann, D. (2002). Correlation and dependence in risk management: properties and pitfalls. In M.A.H. Dempster (ed.), *Risk Management: Value at Risk and Beyond*, 176–223. Cambridge University Press.
- 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.
- Engle, R.F., Kroner, K.F. (1995). Multivariate simultaneous generalized ARCH. *Econometric Theory*, 11(1), 122–150.
- Engle, R.F., Ledoit, O., Wolf, M. (2019). Large dynamic covariance matrices. *Journal of Business & Economic Statistics*, 37(2), 363–375.
- Engle, R.F., Sheppard, K. (2001). Theoretical and empirical properties of dynamic conditional correlation multivariate GARCH. NBER Working Paper 8554.
- Forbes, K.J., Rigobon, R. (2002). No contagion, only interdependence: measuring stock market comovements. *Journal of Finance*, 57(5), 2223–2261.
- Frank, M.J. (1979). On the simultaneous associativity of F(x, y) and x + y − F(x, y). *Aequationes Mathematicae*, 19, 194–226.

References (3/5)

- Franke, J., Härdle, W.K., Hafner, C.M. (2019). Statistics of Financial Markets: An Introduction, 5th ed. Springer.
- 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., 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.
- Genest, C., Rivest, L.-P. (1993). Statistical inference procedures for bivariate Archimedean copulas. *Journal of the American Statistical Association*, 88(423), 1034–1043.
- Grecu, R.A., Cramer, A.A., Pele, D.T., Lessmann, S. (2025). The link between energy prices and stock markets in European Union countries. *The North American Journal of Economics and Finance*, 78, 102420.
- Gumbel, E.J. (1960). Bivariate exponential distributions. *Journal of the American Statistical Association*, 55(292), 698–707.
- Härdle, W.K., Simar, L. (2019). Applied Multivariate Statistical Analysis, 5th ed. Springer.
- Herskovic, B., Kelly, B., Lustig, H., Van Nieuwerburgh, S. (2016). The common factor in idiosyncratic volatility: quantitative asset pricing implications. *Journal of Financial Economics*, 119(2), 249–283.
- Joe, H. (2005). Asymptotic efficiency of the two-stage estimation method for copula-based models. *Journal of Multivariate Analysis*, 94(2), 401–419.
- J.P. Morgan/Reuters (1996). RiskMetrics – Technical Document, 4th ed. New York.
- Kendall, M.G. (1938). A new measure of rank correlation. *Biometrika*, 30(1–2), 81–93.
- Kilian, L., Park, C. (2009). The impact of oil price shocks on the U.S. stock market. *International Economic Review*, 50(4), 1267–1287.
- Li, D.X. (2000). On default correlation: a copula function approach. *Journal of Fixed Income*, 9(4), 43–54.
- Liu, Y., Tsyvinski, A. (2021). Risks and returns of cryptocurrency. *Review of Financial Studies*, 34(6), 2689–2727.

References (4/5)

- Longin, F., Solnik, B. (2001). Extreme correlation of international equity markets. *Journal of Finance*, 56(2), 649–676.
- Mantegna, R.N. (1999). Hierarchical structure in financial markets. *European Physical Journal B*, 11(1), 193–197.
- Markowitz, H. (1952). Portfolio selection. *Journal of Finance*, 7(1), 77–91.
- Molenaar, R., Sénéchal, E., Swinkels, L., Wang, Z. (2024). Empirical evidence on the stock–bond correlation. *Financial Analysts Journal*, 80(3), 17–36.
- 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.
- Newey, W.K., West, K.D. (1987). A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix. *Econometrica*, 55(3), 703–708.
- Oh, D.H., Patton, A.J. (2018). Time-varying systemic risk: evidence from a dynamic copula model of CDS spreads. *Journal of Business & Economic Statistics*, 36(2), 181–195.
- Okhrin, O., Okhrin, Y., Schmid, W. (2013). On the structure and estimation of hierarchical Archimedean copulas. *Journal of Econometrics*, 173(2), 189–204.
- Pakel, C., Shephard, N., Sheppard, K., Engle, R.F. (2021). Fitting vast dimensional time-varying covariance models. *Journal of Business & Economic Statistics*, 39(3), 652–668.
- Patton, A.J. (2006). Modelling asymmetric exchange rate dependence. *International Economic Review*, 47(2), 527–556.
- 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.

References (5/5)

- Politis, D.N., White, H. (2004). Automatic block-length selection for the dependent bootstrap. *Econometric Reviews*, 23(1), 53–70.
- Rémillard, B. (2017). Goodness-of-fit tests for copulas of multivariate time series. *Econometrics*, 5(1), 13.
- Rigobon, R. (2003). Identification through heteroskedasticity. *Review of Economics and Statistics*, 85(4), 777–792.
- Schmidt, R., Stadtmüller, U. (2006). Non-parametric estimation of tail dependence. *Scandinavian Journal of Statistics*, 33(2), 307–335.
- Sklar, A. (1959). Fonctions de répartition à n dimensions et leurs marges. *Publications de l'Institut de Statistique de l'Université de Paris*, 8, 229–231.
- Spearman, C. (1904). The proof and measurement of association between two things. *American Journal of Psychology*, 15(1), 72–101.
- Toma, F.-M., Dima, D.M., Grecu, R., Pele, D.T., Lessmann, S. (2026). Energy Financial Risk Meter: systemic risk in the EU energy sector. SSRN Working Paper.
- Tse, Y.K. (2000). A test for constant correlations in a multivariate GARCH model. *Journal of Econometrics*, 98(1), 107–127.
- 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.