








Time Series Analysis

Chapter 14: Multivariate GARCH models



Daniel Traian PELE

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

Acronyms used in this material

- **AI**: Artificial Intelligence
- **BEKK**: Baba–Engle–Kraft–Kroner (multivariate GARCH model)
- **BET**: Bucharest Exchange Trading (index)
- **BNR**: Banca Națională a României (the National Bank of Romania)
- **CCC**: Constant Conditional Correlation (model)
- **CEE**: Central and Eastern Europe
- **CET**: Central European Time
- **COVID-19**: COronaVIrus Disease 2019
- **DCC**: Dynamic Conditional Correlation
- **EODHD**: EOD Historical Data (market-data provider)
- **EU**: European Union
- **FHS**: Filtered Historical Simulation
- **GARCH**: Generalised AutoRegressive Conditional Heteroskedasticity
- **LR**: Likelihood Ratio (test)
- **MGARCH**: Multivariate GARCH
- **OLS**: Ordinary Least Squares
- **QML**: Quasi-Maximum Likelihood
- **SE**: Standard Error
- **US**: United States
- **UTC**: Coordinated Universal Time (Temps Universel Coordonné)
- **VAR**: Vector AutoRegression
- **VEC**: Vector (half-vectorised) GARCH model
- **VECM**: Vector Error Correction Model

Why this chapter

- **Question**: how do the risks of several assets move **together**? Chapter 5 modelled the volatility of one series at a time
  - the risk of a portfolio, a hedge or the spread of a crisis depends on **covariances and correlations**, and these change in time
- **Route** of the chapter
  - the conditional covariance matrix $\mathbf{H}_t$ and its conditions
  - the VEC and BEKK models; the number of parameters
  - CCC and DCC: volatilities and correlations estimated in two steps
  - applications: correlations in crises, portfolio VaR 1%, hedge ratios
- **Self-study chapter**
  - read the slides in order, redo the worked examples, run the notebook, then answer the self-assessment and the quiz
  - prerequisites: Chapter 5 (GARCH), Chapter 6 (VAR models), Chapter 7 (cointegration)

Learning outcomes

- Explain why portfolio risk, hedging and contagion need a model for time-varying covariances
- Write the VEC, BEKK, CCC and DCC models and count their parameters
- Compute one step of the DCC recursion by hand and interpret a, b and $a + b$
- Estimate a DCC model in two steps and test it against constant correlations
- Compute and backtest a portfolio VaR 1% from $\mathbf{H}_t$
- Compute dynamic minimum-variance hedge ratios and their hedging effectiveness

Reading and tools

- Textbook background: Huang and Petukhina (2022) (GARCH models); surveys: Bauwens, Laurent and Rombouts (2006), Silvennoinen and Teräsvirta (2009)
  - the original papers: Bollerslev, Engle and Wooldridge (1988) (VEC), Engle and Kroner (1995) (BEKK), Bollerslev (1990) (CCC), Engle (2002) (DCC)
- Python Quantlets of this chapter: Quantlets/Ch_14
  - step 1 (univariate GARCH) uses the `arch` package; step 2 (DCC) is written out in about 40 lines of `numpy`, because `arch` has no DCC
- Lecture notebook: open in Google Colab
- Video course: Applied Time Series Analysis with Python

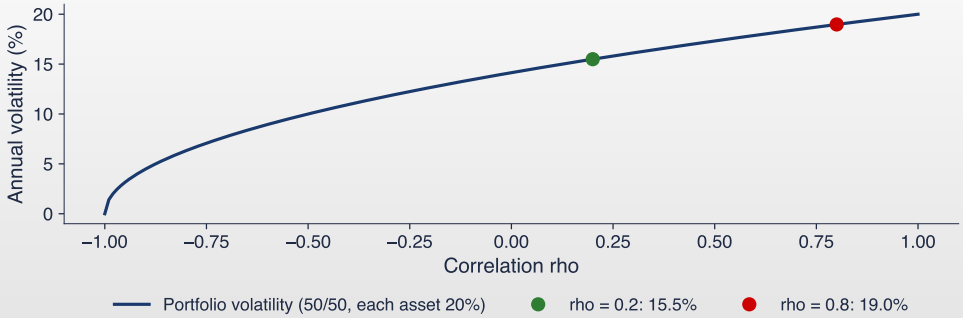
Three questions that need covariances

- **Portfolio risk**
  - the variance of a portfolio with weights $\mathbf{w}$ is $\sigma_p^2 = \mathbf{w}^\top \mathbf{H}\mathbf{w}$ Markowitz (1952)
  - $\mathbf{w}$: the vector of portfolio weights; $\mathbf{H}$: the covariance matrix of the returns; σ_p^2: the portfolio variance
  - two assets: $\sigma_p^2 = w_1^2\sigma_1^2 + w_2^2\sigma_2^2 + 2w_1w_2\rho\,\sigma_1\sigma_2$, with volatilities σ_1, σ_2 and correlation ρ
  - the correlation ρ decides how much diversification helps
- **Hedging**
  - how many units of an asset to sell to offset the risk of another
  - the answer is a ratio of a covariance to a variance (Section 7)
- **Contagion**
  - do markets move more closely together in a crisis?
  - if correlations rise exactly when volatility rises, diversification fails when it is needed most Longin and Solnik (2001)
- A univariate GARCH (Chapter 5) gives $\sigma_{1,t}$ and $\sigma_{2,t}$; this chapter adds the time-varying ρ_t

Worked example: correlation and portfolio risk

- Two assets with annual volatility 20% each, weights 50% and 50%
- $\rho = 0.2$: $\sigma_p^2 = 0.25 \cdot 400 + 0.25 \cdot 400 + 2 \cdot 0.25 \cdot 0.2 \cdot 400 = 240$, so $\sigma_p = 15.5\%$
- $\rho = 0.8$: $\sigma_p^2 = 200 + 2 \cdot 0.25 \cdot 0.8 \cdot 400 = 360$, so $\sigma_p = 19.0\%$
- The same assets and weights: the risk of the portfolio rises by more than a fifth only because the correlation rose
  - a risk model that keeps ρ fixed at its calm-period value underestimates the risk in a crisis

Portfolio volatility against the correlation



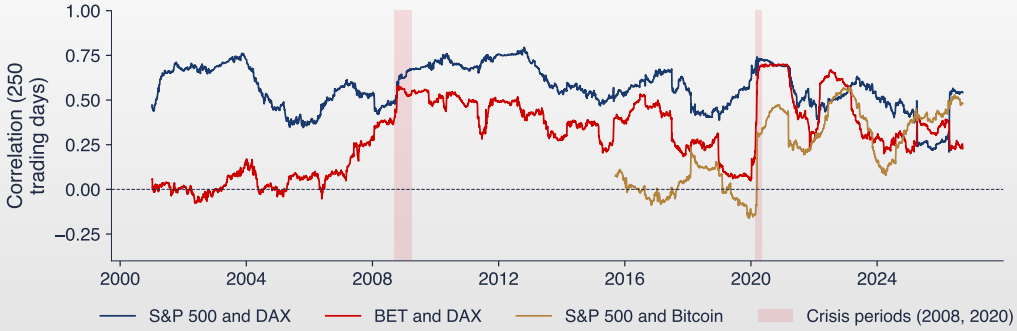
- A 50/50 portfolio of two assets with 20% volatility each; the points are the two cases of the worked example

TSA_ch14_correlations

Interpreting the diversification curve

- $\rho = 1$: no diversification, $\sigma_p = 20\%$; $\rho = 0$: $\sigma_p = 20/\sqrt{2} \approx 14.1\%$; $\rho = -1$: the risk disappears
- The curve is concave: the same rise of ρ adds more volatility at low correlations than at high ones
- The question of the chapter: does ρ stay put in real markets, or does it move with the state of the market?

Correlations change over time



- Correlation of daily log returns over rolling windows of 250 trading days (about one year), dated at the window end
  - returns on the days on which both markets trade; shaded: September 2008–March 2009 and February–April 2020

TSA_ch14_correlations

Interpreting the rolling correlations

- S&P 500 and DAX: between 0.21 (April 2026) and 0.80 (October 2012); full-sample value 0.60
  - the euro-area debt crisis of 2010–2012 tied the two markets even more closely than 2008
- BET and DAX: close to zero before 2007 (minimum -0.08), up to 0.70 in August 2020: the Romanian market became integrated with Europe after EU accession
- S&P 500 and Bitcoin: from -0.16 (December 2019) to 0.57 (February 2023); Bitcoin became a risk asset correlated with equities after 2020
- Correlations wander far from any constant value: a model must let them move

Asynchronous trading and the bias of daily correlations

- Markets close at different hours: Bucharest at 16:45 CET (17:45 local time), Frankfurt at 17:30 CET, New York at 22:00 CET; Bitcoin trades all day, every day
  - news that arrives after the European close shows up in the European prices only the next day
- Example: 9 April 2025, the pause of the US tariffs announced in the afternoon in New York
  - 9 April: S&P 500 $+9.1\%$, DAX -3.1%; 10 April: S&P 500 -3.5%, DAX $+4.4\%$
  - the same news, on two different days: the daily correlation is biased towards zero
- Correlations of daily and weekly returns since 2015: S&P 500 and BET 0.32 against 0.46; S&P 500 and DAX 0.56 against 0.68; BET and DAX 0.42 against 0.47
  - the gap is small for BET and DAX, which close almost together; remedies: weekly returns or returns between synchronous times

Recap: Why multivariate volatility

- Portfolio variance $\mathbf{w}^\top \mathbf{H}\mathbf{w}$, hedge ratios and contagion all depend on covariances
- Rolling correlations of real markets move a lot: integration, crises, new asset classes
- Check the trading hours: asynchronous closes bias daily correlations towards zero

The model for N return series

- $\mathbf{r}_t = \boldsymbol{\mu} + \boldsymbol{\varepsilon}_t, \quad \boldsymbol{\varepsilon}_t = \mathbf{H}_t^{1/2}\boldsymbol{\eta}_t, \quad \boldsymbol{\eta}_t$ i.i.d. with mean $\mathbf{0}$ and covariance $\mathbf{I}_N$
  - $\mathbf{r}_t$: the vector of N returns on day t; $\boldsymbol{\mu}$: the means (or a VAR model for the mean, Chapter 6)
  - $\mathbf{H}_t^{1/2}$: a square root of $\mathbf{H}_t$, e.g. the Cholesky factor $\mathbf{L}_t$ with $\mathbf{L}_t\mathbf{L}_t^\top = \mathbf{H}_t$
- $\mathbf{H}_t = \mathrm{Var}(\mathbf{r}_t \mid \mathcal{F}_{t-1})$: the **conditional covariance matrix**, given the information $\mathcal{F}_{t-1}$ up to day $t-1$
  - diagonal: the conditional variances $h_{ii,t} = \sigma_{i,t}^2$ (Chapter 5)
  - off-diagonal: the conditional covariances $h_{ij,t}$; the conditional correlation is $\rho_{ij,t} = h_{ij,t}/\sqrt{h_{ii,t}h_{jj,t}}$
- A **multivariate GARCH** (MGARCH) model is a rule that computes $\mathbf{H}_t$ from past shocks $\boldsymbol{\varepsilon}_{t-1}$ and past $\mathbf{H}_{t-1}$

Two conditions on $\mathbf{H}_t$

- **Symmetry**
  - $h_{ij,t} = h_{ji,t}$, so only $N(N+1)/2$ elements are distinct
  - the operator vech stacks the lower triangle: $\text{vech}\begin{pmatrix} h_{11} & h_{12} \\ h_{12} & h_{22} \end{pmatrix} = (h_{11}, h_{12}, h_{22})^\top$
  - $N = 2$: 3 elements; $N = 10$: 55; $N = 50$: 1 275
- **Positive definiteness**
  - $\mathbf{w}^\top \mathbf{H}_t \mathbf{w} > 0$ for every $\mathbf{w} \neq \mathbf{0}$, i.e. every portfolio has a positive variance
  - $N = 2$: $h_{11,t} > 0$ and $h_{11,t} h_{22,t} - h_{12,t}^2 > 0$, i.e. $|\rho_{12,t}| < 1$
  - example: $h_{11} = 4$, $h_{22} = 1$, $h_{12} = 2.5$ gives $4 - 6.25 < 0$: an impossible "correlation" of 1.25
- A good MGARCH model guarantees both conditions for **every** t, with few parameters

Estimation by maximum likelihood

- If $\boldsymbol{\eta}_t$ follows the multivariate Normal distribution, the log-likelihood is
  - $\ell(\theta) = -\frac{1}{2}\sum_{t=1}^{T}\left(N\log 2\pi + \log|\mathbf{H}_t| + \boldsymbol{\varepsilon}_t^{\top}\mathbf{H}_t^{-1}\boldsymbol{\varepsilon}_t\right)$
  - $|\mathbf{H}_t|$ is the determinant; for $N = 1$ this is the GARCH likelihood of Chapter 5
- Returns have fat tails, so the Normal likelihood is used as **quasi-maximum likelihood** (QML)
  - the estimates stay consistent; the standard errors need a robust (sandwich) correction
- Every evaluation of $\ell(\theta)$ needs $\mathbf{H}_t^{-1}$ and $|\mathbf{H}_t|$ for all t: the cost grows fast with N and with the number of parameters

Recap: The multivariate GARCH framework

- $\mathbf{H}_t$: conditional variances on the diagonal, covariances off the diagonal; $\rho_{ij,t} = h_{ij,t}/\sqrt{h_{ii,t}h_{jj,t}}$
- Conditions: symmetry ($N(N+1)/2$ distinct elements) and positive definiteness
- Estimation by (quasi-)maximum likelihood with the multivariate Normal density

The VEC model

- Bollerslev, Engle and Wooldridge (1988): every element of $\mathbf{H}_t$ depends on every past squared shock, cross-product and element of $\mathbf{H}_{t-\mathbf{1}}$
  - VEC(1,1): $\operatorname{vech}(\mathbf{H}_t) = \mathbf{c} + \mathbf{A}\operatorname{vech}(\boldsymbol{\varepsilon}_{t-\mathbf{1}}\boldsymbol{\varepsilon}_{t-\mathbf{1}}^{\top}) + \mathbf{B}\operatorname{vech}(\mathbf{H}_{t-\mathbf{1}})$
  - $\boldsymbol{\varepsilon}_{t-\mathbf{1}}\boldsymbol{\varepsilon}_{t-\mathbf{1}}^{\top}$: the matrix of yesterday's squared shocks and cross-products; $\mathbf{c}$: a vector of constants; $\mathbf{A}$: the reaction to shocks; $\mathbf{B}$: the persistence
  - with $k = N(N+1)/2$: $\mathbf{c}$ has k elements, $\mathbf{A}$ and $\mathbf{B}$ are $k \times k$: in total $k + 2k^{\mathbf{2}}$ parameters
- $N = 2$: $k = 3$ and 21 parameters; $N = 10$: 6 105
  - and positive definiteness of $\mathbf{H}_t$ is **not** guaranteed: it needs complicated restrictions
- **Diagonal VEC**
  - $h_{ij,t} = c_{ij} + a_{ij}\varepsilon_{i,t-\mathbf{1}}\varepsilon_{j,t-\mathbf{1}} + b_{ij}h_{ij,t-\mathbf{1}}$
  - each element is a GARCH(1,1) of its own; $3k$ parameters (9 for $N = 2$); still no guarantee of positive definiteness

The BEKK model

- Engle and Kroner (1995): $\mathbf{H}_t = \mathbf{C}\mathbf{C}^\top + \mathbf{A}^\top \boldsymbol{\varepsilon}_{t-1}\boldsymbol{\varepsilon}_{t-1}^\top \mathbf{A} + \mathbf{B}^\top \mathbf{H}_{t-1}\mathbf{B}$
  - BEKK: Baba, Engle, Kraft and Kroner; $\mathbf{C}$ lower triangular, $\mathbf{A}$ and $\mathbf{B}$ are $N \times N$
  - $\mathbf{A}$: the reaction to yesterday's shocks; $\mathbf{B}$: the persistence of $\mathbf{H}_{t-1}$; c_{ij}^*: the elements of $\mathbf{C}\mathbf{C}^\top$
- **Positive definite by construction**
  - each term is a "square" $\mathbf{X}^\top\mathbf{M}\mathbf{X}$ of a positive (semi)definite matrix $\mathbf{M}$
  - $\mathbf{M}$: $\mathbf{I}$, $\boldsymbol{\varepsilon}_{t-1}\boldsymbol{\varepsilon}_{t-1}^\top$ or $\mathbf{H}_{t-1}$; $\mathbf{X}$: $\mathbf{C}^\top$, $\mathbf{A}$ or $\mathbf{B}$
  - parameters: $N(N+1)/2 + 2N^2$, i.e. 11 for $N = 2$ and 255 for $N = 10$
- **Diagonal BEKK** ($\mathbf{A}$, $\mathbf{B}$ diagonal)
  - $h_{12,t} = c_{12}^* + a_{11}a_{22}\,\varepsilon_{1,t-1}\varepsilon_{2,t-1} + b_{11}b_{22}\,h_{12,t-1}$
  - **scalar BEKK**: $\mathbf{A} = a\mathbf{I}$, $\mathbf{B} = b\mathbf{I}$, the same dynamics for all elements
- The off-diagonal elements of a full $\mathbf{A}$ measure **volatility spillovers**: a shock to asset 1 raises the variance of asset 2

Worked example: one step of a diagonal BEKK

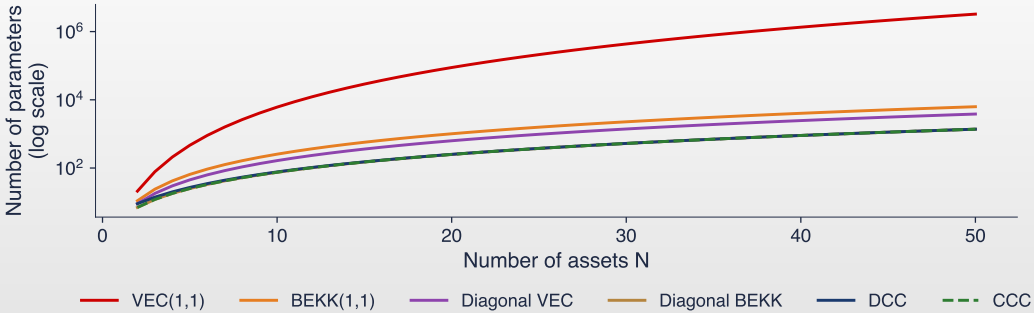
- Given: $c_{12}^{*} = 0.02$, $a_{11} = 0.30$, $a_{22} = 0.25$, $b_{11} = 0.94$, $b_{22} = 0.95$, $h_{12,t-1} = 0.5$; yesterday's shocks $\varepsilon_{1,t-1} = -2$, $\varepsilon_{2,t-1} = -1.5$
- Step 1: the coefficients $a_{11}a_{22} = 0.075$ and $b_{11}b_{22} = 0.893$
- Step 2: the cross-product of shocks $\varepsilon_{1,t-1}\varepsilon_{2,t-1} = 3.000$ (both negative: a joint fall)
- Step 3: $h_{12,t} = 0.020 + 0.075 \cdot 3.000 + 0.893 \cdot 0.5 = 0.020 + 0.225 + 0.4465 = 0.6915$
- A joint fall raises the covariance from 0.5 to 0.6915; a shock of opposite signs would have lowered it
  - note: a diagonal BEKK has no spillovers, because $a_{12} = a_{21} = 0$

The curse of dimensionality

Model	$N = 2$	$N = 5$	$N = 10$	$N = 50$
VEC(1,1)	21	465	6 105	3 252 525
Diagonal VEC	9	45	165	3 825
BEKK(1,1)	11	65	255	6 275
Diagonal BEKK	7	25	75	1 375
Scalar BEKK	5	17	57	1 277
CCC	7	25	75	1 375
DCC	9	27	77	1 377

- Number of parameters of the variance equation (the means μ are not counted); CCC and DCC: 3 GARCH parameters per asset plus the $N(N-1)/2$ correlations, and 2 more for DCC
- A full VEC for 50 assets has more parameters than any sample has observations

Parameters against the number of assets



Number of parameters of the variance equation for $N = 2, \ldots, 50$ assets, log scale

TSA_ch14_parameters

Interpreting the parameter counts

- VEC grows like $N^4/2$, full BEKK like $2N^2$: both become impossible to estimate beyond a handful of assets
- CCC and DCC grow like $N^2/2$, but the correlations are **not** found by numerical optimisation: they are sample correlations (Section 4)
  - the DCC optimiser works on only 2 parameters, a and b, whatever N

Recap: VEC and BEKK

- VEC: the most general linear model; $k + 2k^2$ parameters, no positive-definiteness guarantee
- BEKK: positive definite by construction; off-diagonal a_{ij} measure volatility spillovers
- Diagonal and scalar versions cut the parameters but remove the spillovers

Splitting volatilities and correlations

- Decomposition: $\mathbf{H}_t = \mathbf{D}_t \mathbf{R}_t \mathbf{D}_t, \quad \mathbf{D}_t = \mathrm{diag}(\sigma_{1,t}, \ldots, \sigma_{N,t})$
  - $\mathbf{R}_t$: the conditional correlation matrix (ones on the diagonal); element by element: $h_{ij,t} = \rho_{ij,t}\, \sigma_{i,t}\sigma_{j,t}$
  - each $\sigma_{i,t}$ is a univariate GARCH(1,1) Bollerslev (1986): $\sigma^2_{i,t} = \omega_i + \alpha_i \varepsilon^2_{i,t-1} + \beta_i \sigma^2_{i,t-1}$
- **Standardised residuals**
  - $z_{i,t} = \varepsilon_{i,t}/\sigma_{i,t}$
  - their conditional covariance matrix is exactly $\mathbf{R}_t$
- **CCC** (constant conditional correlation, Bollerslev (1990))
  - $\mathbf{R}_t = \mathbf{R}$ for all t
  - $\mathbf{H}_t$ is positive definite whenever $\mathbf{R}$ is; $\mathbf{R}$ is estimated by the sample correlation of the $z_{i,t}$
  - covariances still move, but only through the volatilities; the rolling correlations of Section 1 reject this

Robert Engle and the DCC model



Robert F. Engle (b. 1942), Nobel Prize 2003

Photo: Jon Demske (2022); CC BY-SA 4.0; Wikimedia Commons

- Engle (2002): let the correlations follow a GARCH-like recursion
  - $\mathbf{Q}_t = (1 - a - b)\,\bar{\mathbf{Q}} + a\,\mathbf{z}_{t-1}\mathbf{z}_{t-1}^\top + b\,\mathbf{Q}_{t-1}$
  - $\mathbf{Q}_t$: an auxiliary matrix with elements $q_{ij,t}$; $\mathbf{z}_{t-1}$: yesterday's vector of standardised residuals
  - $\rho_{ij,t} = q_{ij,t} / \sqrt{q_{ii,t}\, q_{jj,t}}$: the rescaling puts ones on the diagonal
- $\bar{\mathbf{Q}}$: the unconditional correlation of $\mathbf{z}_t$, set to its sample value (**correlation targeting**)
  - $a \geq 0$: reaction to yesterday's joint shock; $b \geq 0$: persistence; $a + b < 1$: mean reversion to $\bar{\mathbf{Q}}$
- $a = b = 0$ gives back CCC

Two-step estimation

- The log-likelihood splits into a volatility part and a correlation part: $\ell = \ell_V(\theta_{\mathbf{1}}) + \ell_C(\theta_{\mathbf{1}}, a, b)$
  - $\theta_{\mathbf{1}}$: the GARCH parameters of all series; ℓ_V depends only on them, ℓ_C also on a and b
- **Step 1**
  - fit a GARCH(1,1) to each series separately (Chapter 5)
  - keep $\hat{\sigma}_{i,t}$ and $\hat{z}_{i,t} = \hat{\varepsilon}_{i,t}/\hat{\sigma}_{i,t}$
- **Step 2**
  - set $\bar{\mathbf{Q}}$ to the sample correlation of $\hat{\mathbf{z}}_t$ and maximise over (a, b) only
  - $\ell_C(a, b) = -\frac{\mathbf{1}}{\mathbf{2}} \sum_t \left(\log|\mathbf{R}_t| + \hat{\mathbf{z}}_t^\top \mathbf{R}_t^{-\mathbf{1}} \hat{\mathbf{z}}_t - \hat{\mathbf{z}}_t^\top \hat{\mathbf{z}}_t \right)$
- Consistent, fast, works for large N; the step-2 standard errors ignore the estimation error of step 1 (they are too small)
- **Test against CCC**
  - likelihood ratio $LR = 2(\ell_C^{DCC} - \ell_C^{CCC})$, compared with $\chi^{\mathbf{2}}(2)$ (5% critical value 5.99)
  - also Tse (2000)
  - ℓ_C^{DCC}, ℓ_C^{CCC}: the maximised correlation log-likelihoods of the two models; a large LR rejects constant correlation
  - $a = b = 0$ lies on the boundary of the parameter space, so the $\chi^{\mathbf{2}}(2)$ critical value is conservative

Worked example: one step of the DCC recursion

- Given: $a = 0.05$, $b = 0.93$, $\bar{q}_{\mathbf{12}} = 0.5$ (and $\bar{q}_{\mathbf{11}} = \bar{q}_{\mathbf{22}} = 1$), $\mathbf{Q}_{t-\mathbf{1}} = \begin{pmatrix} 1 & 0.5 \\ 0.5 & 1 \end{pmatrix}$, yesterday $z_{\mathbf{1},t-\mathbf{1}} = -2$, $z_{\mathbf{2},t-\mathbf{1}} = -2.5$
- Step 1: $1 - a - b = 0.020$
- Step 2: $q_{\mathbf{12},t} = 0.020 \cdot 0.5 + 0.05 \cdot (-2)(-2.5) + 0.93 \cdot 0.5 = 0.010 + 0.250 + 0.465 = 0.725$
- Step 3: $q_{\mathbf{11},t} = 0.020 + 0.05 \cdot 4 + 0.93 = 1.150$; $\quad q_{\mathbf{22},t} = 0.020 + 0.05 \cdot 6.25 + 0.93 = 1.2625$
- Step 4: $\rho_{\mathbf{12},t} = 0.725/\sqrt{1.150 \cdot 1.2625} = 0.602$
  - a large joint fall raises tomorrow's correlation from 0.50 to 0.602; without step 3 the "correlation" $q_{\mathbf{12},t}$ would be 0.725

Reading a and b

- a is the **news** coefficient
  - how much one day's joint shock moves the correlation
  - typical daily values 0.01–0.05
- $a + b$ is the **persistence**
  - a deviation of $q_{ij,t}$ from $\bar{q}_{ij}$ shrinks by the factor $a + b$ each day
  - half-life: $\ln 0.5 / \ln(a + b)$ days; $a + b = 0.98$ gives 34 days, 0.99 gives 69 days
- The same reading as α and $\alpha + \beta$ in a GARCH(1,1) (Chapter 5)
- Variants: asymmetric DCC (joint falls move correlations more than joint rises) Cappiello, Engle and Sheppard (2006); corrected DCC Aielli (2013)

Recap: CCC and DCC

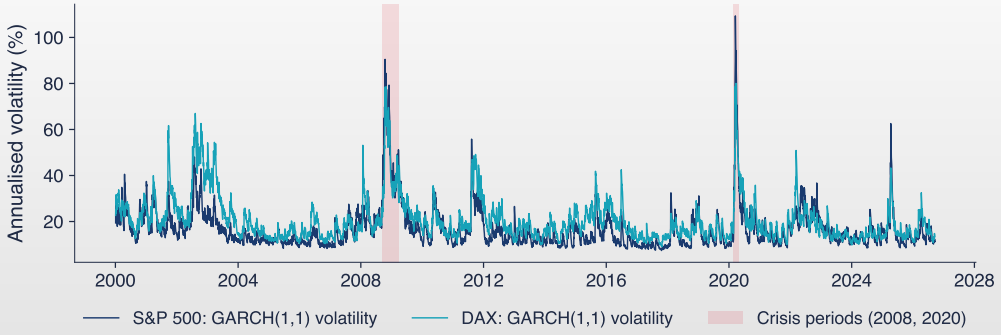
- $\mathbf{H}_t = \mathbf{D}_t \mathbf{R}_t \mathbf{D}_t$: GARCH volatilities times a correlation matrix
- CCC: $\mathbf{R}$ constant; DCC: $\mathbf{Q}_t$ follows a GARCH-like recursion with parameters a, b
- Two steps: univariate GARCH, then (a, b) with correlation targeting; LR test against CCC

Step 1: two univariate GARCH models

Index	μ	ω	α	β	$\alpha+\beta$	kurt. r_t	kurt. z_t
S&P 500	0.063	0.027	0.122	0.857	0.980	10.5	1.8
DAX	0.067	0.030	0.098	0.886	0.984	6.2	1.4

- Daily log returns in %, 6 605 common trading days, 4 January 2000–18 September 2026 (EODHD); GARCH(1,1) with a constant mean, Gaussian QML; kurt.: excess kurtosis
- Standardising by $\hat{\sigma}_{i,t}$ removes most of the excess kurtosis, but not all: the tails of z_t are still fatter than Normal
- Correlation of the returns 0.599; correlation of the residuals $\bar{q}_{\mathbf{12}} = 0.575$, the CCC estimate

Step 1: conditional volatilities



⊡ One-step-ahead GARCH(1,1) volatility $\hat{\sigma}_{i,t}\sqrt{252}$, annualised, in %; shaded: the crises of 2008 and 2020

TSA_ch14_dcc_estimation

Interpreting the volatilities

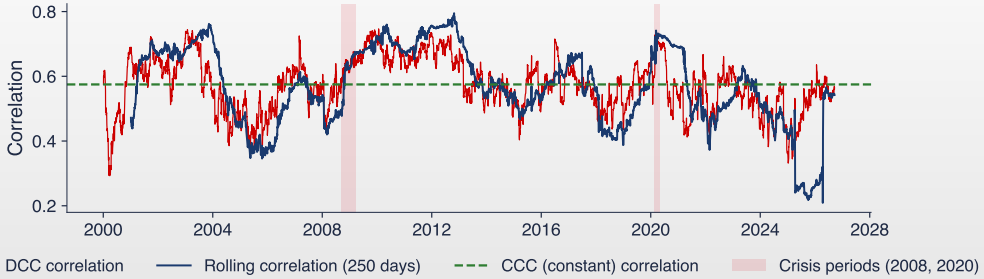
- Both volatilities peak in the same weeks: S&P 500 109% on 17 March 2020, DAX 80% on 25 March 2020
  - volatility clustering happens **jointly**: the covariance rises even with a constant correlation
- Persistence 0.980 and 0.984: volatility shocks fade slowly, as in Chapter 5
- Step 2 asks a different question: once the volatilities are removed, is the remaining correlation constant?

Step 2: the estimated DCC

Quantity	Estimate	SE
a	0.0152	0.0035
b	0.9766	0.0065
$a + b$	0.9918	
Half-life of a correlation shock (days)	85	
LR against CCC	131.8	
$\chi^2(2)$ critical value, 5%	5.99	

- a is small but well determined; $a + b$ is close to 1: correlation shocks last for months
- LR $= 131.8$ is far above 5.99: constant correlation is rejected
- SE: standard errors from the step-2 Hessian (too small, they ignore step 1)

DCC correlation of the S&P 500 and the DAX



One-step-ahead DCC correlation $\hat{\rho}_{\mathbf{12},t}$, the 250-day rolling correlation and the CCC value $\bar{q}_{\mathbf{12}}$; daily, 4 January 2000–18 September 2026

TSA_ch14_dcc_estimation

Interpreting the DCC correlation

- The DCC correlation ranges from 0.29 (March 2000) to 0.75 (January 2010), around a mean of 0.57
- It follows the rolling correlation, but reacts within days instead of a year and has no window to choose
  - the rolling estimate is more erratic (standard deviation 0.122 against 0.083) and jumps when an extreme day enters or leaves the window
- The rolling estimate falls sharply in 2025–2026: the asynchronous days of April 2025 (Section 1) pull it down for exactly 250 days

Correlations in crises



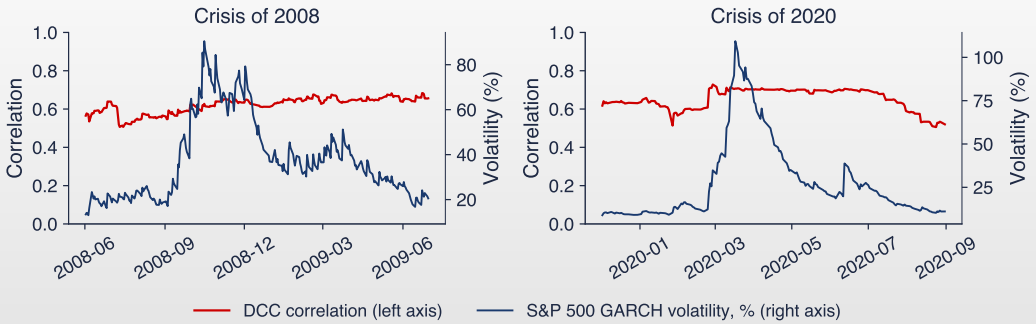
Lehman Brothers, New York, 15 September 2008

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

Period	$\bar{\rho}^{DCC}$	$\bar{\sigma}^{S\&P}$ (%)
Calm, 2017–2019	0.56	13
Sept 2008–March 2009	0.63	51
Feb–April 2020	0.70	58

- Average DCC correlation and average annualised S&P 500 volatility in each period
- Correlations are higher in crises, but volatility rises far more
- Caution Forbes and Rigobon (2002): the sample correlation rises mechanically when volatility rises, even if the link is unchanged; DCC reduces this effect because it works with z_t, not with r_t

DCC correlation in the crises of 2008 and 2020



DCC correlation of the S&P 500 and the DAX (left axis) and the GARCH volatility of the S&P 500 (right axis)

TSA_ch14_dcc_estimation

Interpreting the two crises

- 2008: volatility peaks at 90% on 16 October 2008, while the correlation drifts up from 0.57 to 0.68 only by 23 June 2009
- 2020: the correlation jumps to 0.73 on 28 February 2020, before the volatility peak (109% on 17 March 2020)
  - a global shock hit both markets on the same days: joint falls drive the DCC recursion
- Correlations and volatilities need not peak together: modelling them separately, as DCC does, shows the difference

Recap: S&P 500 and DAX

- Step 1: persistent GARCH volatilities that peak together
- Step 2: $a + b = 0.9918$; LR 131.8: the correlation is not constant
- Higher correlations in 2008 and 2020, but the largest change is in volatility

A five-asset DCC

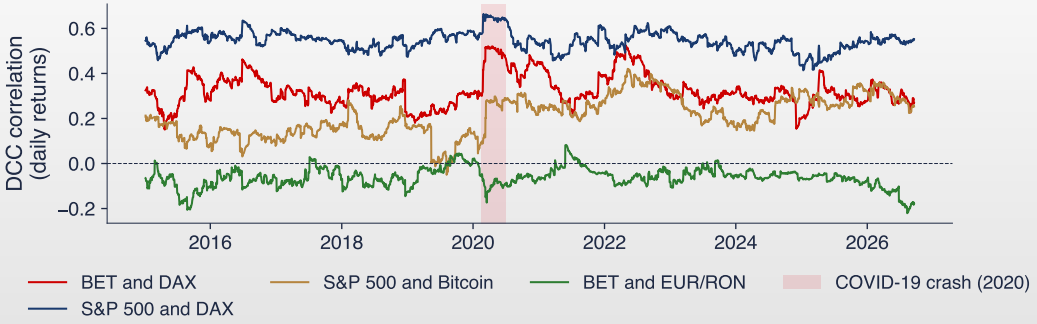


The Bucharest Stock Exchange

Photo: Corina Chitu (2024); CC BY-SA 4.0; Wikimedia Commons

- Daily log returns on the 2 815 days on which all five markets trade, from 6 January 2015
  - EUR/RON: the BNR reference rate (fixed at 13:00 Bucharest time); BET: Bucharest Stock Exchange; Bitcoin: close of the day (UTC)
- Step 1: GARCH persistence S&P 500 0.961, DAX 0.965, BET 0.941, EUR/RON 0.985, Bitcoin 0.958
- Step 2: $\hat{a} = 0.0088$ (SE 0.0019), $\hat{b} = 0.9782$ (SE 0.0062); half-life 53 days; LR against CCC 88.5
  - 10 correlation pairs, one pair of dynamics parameters (a, b) for all of them

Dynamic correlations of four pairs



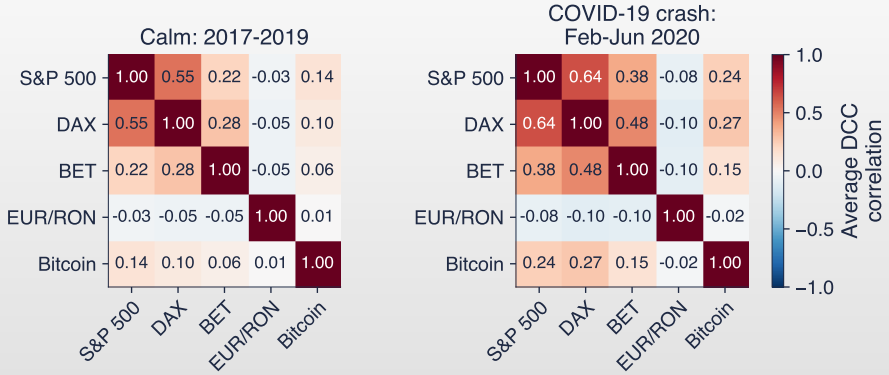
One-step-ahead DCC correlations from the five-asset model; shaded: the COVID-19 crash, mid-February–June 2020

TSA_ch14_dcc_panel

Interpreting the four pairs

- S&P 500 and DAX: the strongest link, mean 0.54; BET and DAX: mean 0.32, between 0.15 and 0.53 (April 2022)
- S&P 500 and Bitcoin: from -0.05 to 0.42 (May 2022); mean 0.21
  - the diversification benefit of Bitcoin shrank after 2020
- BET and EUR/RON: close to zero (mean -0.06); a weaker leu tends to come with a weaker BET, a small negative link
- One (a, b) for all pairs: pairs with different speeds are forced to the same rhythm, the price of DCC's parsimony

Average correlations: calm against crisis



- Average one-step-ahead DCC correlation matrix: 2017–2019 and mid-February–June 2020 (89 common trading days)

TSA_ch14_dcc_panel

Interpreting the two matrices

- The average pairwise correlation rises from 0.12 to 0.19
  - BET and DAX: $0.28 \to 0.48$; S&P 500 and BET: $0.22 \to 0.38$; DAX and Bitcoin: $0.10 \to 0.27$
- The largest rises are for the BET and for Bitcoin, which are weakly linked to the large markets in calm times: exactly the assets that were supposed to diversify
- EUR/RON stays near zero or slightly negative with all equity markets: the only diversifier in this set

Recap: Five markets

- One DCC for 5 assets: 2 dynamics parameters for 10 correlations
- Correlations of the Romanian market with Europe and of Bitcoin with equities move a lot
- In the 2020 crash the weak links strengthened most: diversification weakens in crises

From $\mathbf{H}_t$ to the VaR of a portfolio

- Portfolio return $r_{p,t} = \mathbf{w}^\top \mathbf{r}_t$; conditional mean $\mu_p = \mathbf{w}^\top \boldsymbol{\mu}$ and volatility $\sigma_{p,t} = \sqrt{\mathbf{w}^\top \mathbf{H}_t \mathbf{w}}$
- **VaR 1%**
  - the loss exceeded with probability 1%, $\text{VaR}_{0.01,t} = -q_{0.01}(r_{p,t}) = -(\mu_p + q\,\sigma_{p,t})$
  - Normal: $q = z_{0.01} = -2.326$
  - **filtered historical simulation** (FHS): q = the empirical 1% quantile of the past standardised portfolio returns $(r_{p,t} - \mu_p)/\sigma_{p,t}$
- Three choices of $\mathbf{H}_t$: DCC; CCC (GARCH volatilities with a constant correlation); **static** (one sample covariance matrix for all days)

Worked example: VaR 1% of two assets

- Today's forecasts: $\sigma_1 = 1.5\%$, $\sigma_2 = 1.2\%$ (daily), weights 50% and 50%, mean 0
- Calm, $\rho_t = 0.3$: $\sigma_p^2 = 0.25 \cdot 2.25 + 0.25 \cdot 1.44 + 2 \cdot 0.25 \cdot 0.3 \cdot 1.8$, so $\sigma_p = 1.092\%$ and VaR 1% $= 2.326 \cdot 1.092 = 2.54\%$
- Crisis, $\rho_t = 0.9$ (the same volatilities): $\sigma_p = 1.316\%$ and VaR 1% $= 3.06\%$
  - the jump of the correlation alone adds a fifth to the VaR; in a real crisis the volatilities rise too
- For a portfolio of 1 million lei: VaR 1% of about 25.4 thousand lei in calm times and 30.6 thousand lei in the crisis case

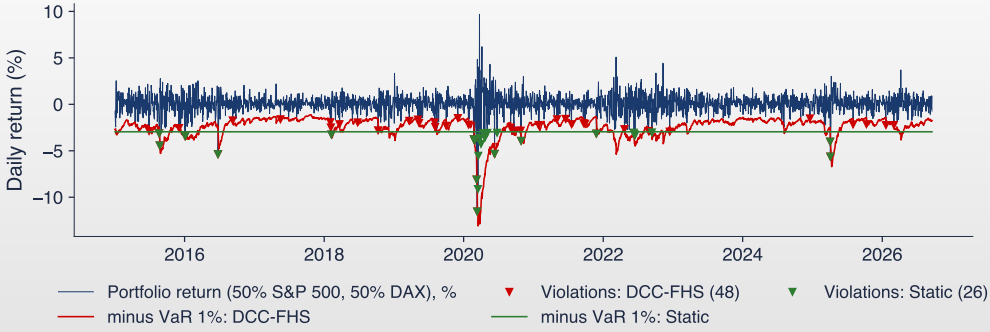
Backtesting a VaR (1/2): frequency

- A **violation** on day t
  - $r_{p,t} < -\text{VaR}_{\mathbf{0.01},t}$
  - a correct VaR 1% has violations on 1% of the days, at random times
- **Kupiec test** Kupiec (1995) (frequency)
  - x violations in n days, $\hat{p} = x/n$
  - $LR_{uc} = -2 \log \dfrac{(1-0.01)^{n-x}\, 0.01^{x}}{(1-\hat{p})^{n-x}\, \hat{p}^{x}} \sim \chi^{\mathbf{2}}(1)$; reject at 5% if $LR_{uc} > 3.84$
  - uc: unconditional coverage; the ratio compares the likelihood of the violations under the rate 1% with that under the observed rate $\hat{p}$

Backtesting a VaR (2/2): independence

- **Christoffersen test** Christoffersen (1998) (independence)
  - is a violation more likely the day after a violation?
  - compares $\pi_{11} = P(\text{violation} \mid \text{violation yesterday})$ with $\pi_{01} = P(\text{violation} \mid \text{no violation yesterday})$
- The statistic
  - $LR_{ind} = -2 \log \dfrac{(1-\hat{\pi})^{n_{00}+n_{10}}\,\hat{\pi}^{\,n_{01}+n_{11}}}{(1-\hat{\pi}_{01})^{n_{00}}\,\hat{\pi}_{01}^{n_{01}}\,(1-\hat{\pi}_{11})^{n_{10}}\,\hat{\pi}_{11}^{n_{11}}} \sim \chi^2(1)$
  - n_{ij}: the number of days in state j after a day in state i (1: violation, 0: none); $\hat{\pi}_{01} = n_{01}/(n_{00}+n_{01})$, $\hat{\pi}_{11} = n_{11}/(n_{10}+n_{11})$
  - $\hat{\pi}$: the overall violation rate; the numerator assumes $\pi_{01} = \pi_{11}$ (no clustering); a small p-value rejects independence
- **Out-of-sample design**
  - estimate all parameters on 2000–2014 (3 712 days), then run the filters forward on 2015–2026 (2 893 days) without re-estimating

VaR 1% of an S&P 500 and DAX portfolio, 2015–2026



⊡ Daily return of a 50/50 portfolio and minus VaR 1% from DCC with filtered historical simulation and from a static covariance matrix; triangles: violations

TSA_ch14_portfolio_var

Interpreting the backtest

Model	Violations	Rate	LR_{uc}	p	p indep.	Feb–Apr 2020
DCC, Normal	65	2.25%	33.6	< 0.001	0.003	4
DCC, FHS	48	1.66%	10.6	0.001	0.008	4
CCC, Normal	57	1.97%	21.4	< 0.001	0.030	4
Static, Normal	26	0.90%	0.3	0.577	0.001	10

- Expected: about 29 violations in 2 893 days
- Dynamic models with Normal quantiles have too many violations: the residuals have fat tails (Chapter 5); FHS ($q = -2.59$) brings the rate down to 1.66%
- The static VaR passes the frequency test, but its violations cluster: 10 in February–April 2020 alone, and after a violation the next day is a violation with probability 11.5%
- The static VaR is too high in calm years (2.95% every day) and far too low in a crash; a dynamic VaR adapts in both directions

Recap: Portfolio VaR

- $\sigma_{p,t} = \sqrt{\mathbf{w}^\top \mathbf{H}_t \mathbf{w}}$; VaR 1% $= -(\mu_p + q\,\sigma_{p,t})$
- Backtest out of sample: Kupiec (frequency) and Christoffersen (independence)
- Dynamic covariances fix the clustering of violations; fat tails need a better quantile than the Normal

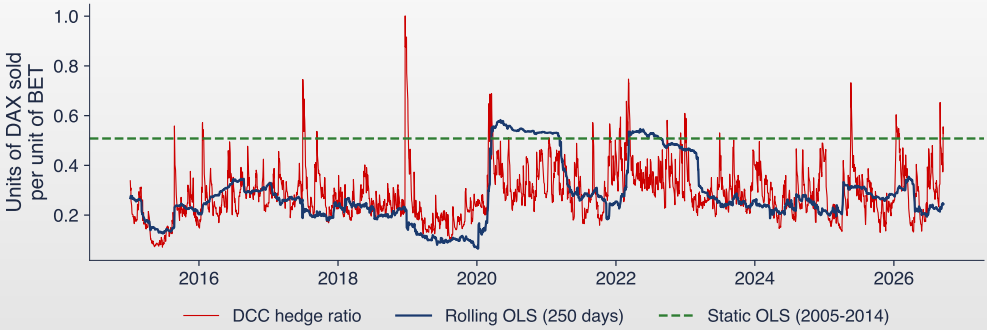
The minimum-variance hedge ratio

- Hold one unit of asset s (e.g. Romanian shares, the BET) and sell h units of a hedging asset f (e.g. a DAX futures contract)
  - hedged return $r_{s,t} - h\, r_{f,t}$, with variance $\sigma_s^2 - 2h\,\sigma_{sf} + h^2\sigma_f^2$
  - σ_s^2, σ_f^2: the variances of the two returns; σ_{sf}: their covariance
- Setting the derivative in h to zero: $h^* = \dfrac{\sigma_{sf}}{\sigma_f^2} = \rho\,\dfrac{\sigma_s}{\sigma_f}$ Ederington (1979)
  - static h^*: the OLS slope of r_s on r_f
  - **dynamic**: $h_t^* = h_{sf,t}/h_{ff,t}$ from an MGARCH model Kroner and Sultan (1993); $h_{sf,t}$, $h_{ff,t}$: elements of $\mathbf{H}_t$
- **Hedging effectiveness**
  - $HE = 1 - \mathrm{Var}(r_s - h\, r_f)/\mathrm{Var}(r_s)$, the share of the variance removed
  - with $h = h^*$ constant: $HE = \rho^2$; a hedge is only as good as the correlation

Worked example: a hedge ratio

- Given: $\sigma_s = 1.2\%$, $\sigma_f = 1.0\%$ (daily), $\rho = 0.8$
- Step 1: $\sigma_{sf} = \rho\,\sigma_s\sigma_f = 0.8 \cdot 1.2 \cdot 1.0 = 0.96$
- Step 2: $h^* = \sigma_{sf}/\sigma_f^2 = 0.96/1 = 0.96$: sell 0.96 units of f per unit of s
- Step 3: $HE = \rho^2 = 0.64$; the hedged volatility is $\sigma_s\sqrt{1-\rho^2} = 0.72\%$ instead of 1.2%
- If tomorrow DCC forecasts $\rho_t = 0.5$ and $\sigma_{f,t} = 2\%$, the hedge ratio becomes $0.5 \cdot 1.2/2 = 0.3$: the position must be adjusted

Hedging the BET with the DAX, 2015–2026



- Daily returns on common trading days since 2005; all parameters estimated on 2005–2014; DCC: $h_t = \hat{\rho}_t \hat{\sigma}_{BET,t} / \hat{\sigma}_{DAX,t}$; rolling OLS: the slope over the previous 250 days

TSA_ch14_hedging

Interpreting the hedge ratios

- The static ratio 0.51 reflects 2005–2014, which includes the 2008 crisis; after 2015 the DCC ratio averages 0.29, from 0.07 to 1.00 (20 December 2018)
- Hedging effectiveness out of sample (2 893 days): static 13.6%, rolling OLS 18.0%, DCC 19.1%
  - the dynamic ratios remove more risk, and DCC does so without choosing a window
- All effectiveness values are modest: the BET–DAX DCC correlation averages 0.32 since 2015, so ρ^2 is small
  - a daily DCC ratio also implies daily trading costs, which are not counted here

Recap: Hedging

- $h_t^* = h_{sf,t}/h_{ff,t} = \rho_t \sigma_{s,t}/\sigma_{f,t}$; $HE = 1 - \text{Var(hedged)}/\text{Var(unhedged)}$
- BET with DAX: DCC 19.1% against static 13.6% out of sample
- Judge a hedge out of sample and after trading costs

Links with Chapters 5, 6 and 7

- **Chapter 5**
  - every MGARCH model contains univariate GARCH models
  - in DCC they are literally step 1
- **Chapter 6**
  - a VAR models spillovers in the **mean** (e.g. whether yesterday's DAX return helps to forecast today's BET return)
  - BEKK with a full **A** models spillovers in the **variance**
  - the two combine: a VAR for $\boldsymbol{\mu}_t$ and an MGARCH for $\mathbf{H}_t$; spillover indices from variance decompositions Diebold and Yilmaz (2009)
- **Chapter 7**
  - cointegration is a **long-run** link between price levels
  - correlation is a **short-run** link between returns
  - two prices can be highly correlated day by day and drift apart for ever, or cointegrated with low daily correlation
  - a VECM can carry a DCC error term: equilibrium in the mean, dynamic correlation in the shocks

Limits and extensions



Frankfurt Stock Exchange

Photo: Ank Kumar (2015); CC BY-SA 4.0; Wikimedia Commons

- One (a, b) for all pairs; no spillovers between volatilities in DCC
- Normal QML underestimates the tails: use Student t innovations or FHS for risk
- Asymmetry: correlations react more to joint falls (asymmetric DCC, Cappiello, Engle and Sheppard (2006))
- Hundreds of assets: factor models, composite likelihood and shrinkage of $\bar{\mathbf{Q}}$
- Correlation is linear dependence; for joint extremes see copulas and tail dependence Longin and Solnik (2001)

Possible contribution of AI

- **Code**
  - a first draft of the DCC likelihood, of a BEKK filter, of a Kupiec test
- **Explanation**
  - a second explanation of $\mathbf{D}_t\mathbf{R}_t\mathbf{D}_t$, of correlation targeting, of the hedge-ratio formula
- **Exploration**
  - DCC for many pairs (CEE markets, sectors, currencies) and their behaviour in crises
- Example prompt
  - `Write Python code that fits a GARCH(1,1) to daily BET and DAX log returns on common trading days, estimates a DCC(1,1) model by two-step maximum likelihood with correlation targeting, plots the dynamic hedge ratio and compares its out-of-sample hedging effectiveness with static OLS.`

Checks you must run

- The library: `arch` has no DCC model; an AI answer that calls `arch_model(..., vol='DCC')` has invented it
- The timing: $\mathbf{Q}_t$ must use $\mathbf{z}_{t-1}$, not $\mathbf{z}_t$; otherwise the VaR and the hedge use tomorrow's information
- The rescaling: $\rho_{ij,t} = q_{ij,t} / \sqrt{q_{ii,t} q_{jj,t}}$, and every $\mathbf{R}_t$ must have ones on the diagonal and $|\rho| < 1$
- The calendars: align prices on common trading days before computing returns; watch asynchronous closes
- The evaluation: estimate on one period and test on another; simulate a DCC with known (a, b) and check that the code recovers it
- Every cited reference: it must exist; check the DOI

Key takeaways

- Portfolio risk, hedging and contagion depend on time-varying covariances $\mathbf{H}_t$
- VEC is general but has too many parameters; BEKK is positive definite by construction and measures spillovers
- CCC and DCC split $\mathbf{H}_t = \mathbf{D}_t \mathbf{R}_t \mathbf{D}_t$ and are estimated in two steps; DCC needs only (a, b) in step 2
- Real correlations move: integration, crises, new assets; constant correlation is rejected
- Applications: portfolio VaR 1% with backtesting; dynamic hedge ratios judged out of sample

Key formulas

Quantity	Formula
Correlation	$\rho_{ij,t} = h_{ij,t} / \sqrt{h_{ii,t} h_{jj,t}}$
BEKK(1,1)	$\mathbf{H}_t = \mathbf{C}\mathbf{C}^\top + \mathbf{A}^\top \boldsymbol{\varepsilon}_{t-\mathbf{1}} \boldsymbol{\varepsilon}_{t-\mathbf{1}}^\top \mathbf{A} + \mathbf{B}^\top \mathbf{H}_{t-\mathbf{1}} \mathbf{B}$
CCC, DCC	$\mathbf{H}_t = \mathbf{D}_t \mathbf{R}_t \mathbf{D}_t, \quad \mathbf{D}_t = \operatorname{diag}(\sigma_{i,t})$
DCC	$\mathbf{Q}_t = (1 - a - b)\bar{\mathbf{Q}} + a\, \mathbf{z}_{t-\mathbf{1}} \mathbf{z}_{t-\mathbf{1}}^\top + b\, \mathbf{Q}_{t-\mathbf{1}}, \quad \rho_{ij,t} = q_{ij,t} / \sqrt{q_{ii,t} q_{jj,t}}$
Half-life	$\ln 0.5 / \ln(a + b)$
Portfolio VaR 1%	$-(\mathbf{w}^\top \boldsymbol{\mu} + q_{\mathbf{0.01}} \sqrt{\mathbf{w}^\top \mathbf{H}_t \mathbf{w}})$
Hedge ratio	$h_t^* = h_{sf,t} / h_{ff,t}, \quad HE = 1 - \operatorname{Var}(r_s - h r_f) / \operatorname{Var}(r_s)$

Self-assessment (1/2)

- **Question**: how many parameters does a diagonal BEKK(1,1) have for $N = 3$?
  - **Answer**: $N(N+1)/2 + 2N = 6 + 6 = 12$
- **Question**: $h_{\mathbf{11}} = 1$, $h_{\mathbf{22}} = 4$, $h_{\mathbf{12}} = 1.2$. What is the correlation, and is the matrix positive definite?
  - **Answer**: $\rho = 1.2/\sqrt{4} = 0.6$; yes, $1 \cdot 4 - 1.44 > 0$
- **Question**: DCC with $a = 0.02$, $b = 0.97$, $\bar{q}_{\mathbf{12}} = 0.4$, $q_{\mathbf{11},t-\mathbf{1}} = q_{\mathbf{22},t-\mathbf{1}} = 1$, $q_{\mathbf{12},t-\mathbf{1}} = 0.4$ and $z_{t-\mathbf{1}} = (1, -1)$. What is $q_{\mathbf{12},t}$?
  - **Answer**: $0.01 \cdot 0.4 + 0.02 \cdot (-1) + 0.97 \cdot 0.4 = 0.372$; opposite moves lower the correlation
- **Question**: what is the half-life of correlation shocks when $a + b = 0.95$?
  - **Answer**: $\ln 0.5 / \ln 0.95 \approx 13.5$ days

Self-assessment (2/2)

- **Question**: $\sigma_s = 2\%$, $\sigma_f = 1\%$, $\rho = 0.5$. What are the hedge ratio and the hedging effectiveness?
  - **Answer**: $h^* = 0.5 \cdot 2/1 = 1$; $HE = 0.25$
- **Question**: a VaR 1% has 26 violations in 2 500 days, 10 of them in one month. Is it a good VaR?
  - **Answer**: the frequency is fine (expected 25), but the violations cluster: it fails the independence test and does not adapt to crises
- **Question**: why can the daily correlation of the S&P 500 and the BET be lower than the weekly one?
  - **Answer**: asynchronous trading: news from New York reaches Bucharest the next day, so daily returns of the same day do not match
- **Question**: does a higher correlation in a crisis prove contagion?
  - **Answer**: not by itself: higher volatility raises the sample correlation mechanically Forbes and Rigobon (2002); compare correlations of standardised residuals
- Next: Chapter 15, review and exam preparation; check yourself with the quiz of this chapter

References (1/2)

- Aielli, G. P. (2013). Dynamic conditional correlation: On properties and estimation. *Journal of Business & Economic Statistics*, 31(3), 282–299. doi:10.1080/07350015.2013.771027
- Bauwens, L., Laurent, S., & Rombouts, J. V. K. (2006). Multivariate GARCH models: A survey. *Journal of Applied Econometrics*, 21(1), 79–109. doi:10.1002/jae.842
- Bollerslev, T. (1986). Generalized autoregressive conditional heteroskedasticity. *Journal of Econometrics*, 31(3), 307–327. doi:10.1016/0304-4076(86)90063-1
- Bollerslev, T. (1990). Modelling the coherence in short-run nominal exchange rates: A multivariate generalized ARCH model. *The Review of Economics and Statistics*, 72(3), 498–505. doi:10.2307/2109358
- 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. doi:10.1086/261527
- Cappiello, L., Engle, R. F., & Sheppard, K. (2006). Asymmetric dynamics in the correlations of global equity and bond returns. *Journal of Financial Econometrics*, 4(4), 537–572. doi:10.1093/jjfinec/nbl005
- Christoffersen, P. F. (1998). Evaluating interval forecasts. *International Economic Review*, 39(4), 841–862. doi:10.2307/2527341
- Diebold, F. X., & Yilmaz, K. (2009). Measuring financial asset return and volatility spillovers, with application to global equity markets. *The Economic Journal*, 119(534), 158–171. doi:10.1111/j.1468-0297.2008.02208.x
- Ederington, L. H. (1979). The hedging performance of the new futures markets. *The Journal of Finance*, 34(1), 157–170. doi:10.1111/j.1540-6261.1979.tb02077.x
- 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. doi:10.1198/073500102288618487

References (2/2)

- Engle, R. F., & Kroner, K. F. (1995). Multivariate simultaneous generalized ARCH. *Econometric Theory*, 11(1), 122–150. doi:10.1017/S0266466600009063
- Forbes, K. J., & Rigobon, R. (2002). No contagion, only interdependence: Measuring stock market comovements. *The Journal of Finance*, 57(5), 2223–2261. doi:10.1111/0022-1082.00494
- Huang, C., & Petukhina, A. (2022). *Applied Time Series Analysis and Forecasting with Python*. Springer. doi:10.1007/978-3-031-13584-2
- Kroner, K. F., & Sultan, J. (1993). Time-varying distributions and dynamic hedging with foreign currency futures. *Journal of Financial and Quantitative Analysis*, 28(4), 535–551. doi:10.2307/2331164
- Kupiec, P. H. (1995). Techniques for verifying the accuracy of risk measurement models. *The Journal of Derivatives*, 3(2), 73–84. doi:10.3905/jod.1995.407942
- Longin, F., & Solnik, B. (2001). Extreme correlation of international equity markets. *The Journal of Finance*, 56(2), 649–676. doi:10.1111/0022-1082.00340
- Markowitz, H. (1952). Portfolio selection. *The Journal of Finance*, 7(1), 77–91. doi:10.1111/j.1540-6261.1952.tb01525.x
- Silvennoinen, A., & Teräsvirta, T. (2009). Multivariate GARCH models. In T. G. Andersen et al. (Eds.), *Handbook of Financial Time Series* (pp. 201–229). Springer. doi:10.1007/978-3-540-71297-8_9
- Tse, Y. K. (2000). A test for constant correlations in a multivariate GARCH model. *Journal of Econometrics*, 98(1), 107–127. doi:10.1016/S0304-4076(99)00080-9