








Modelling Financial Markets

Seminar 5: Conditional Volatility — GARCH Models



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

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

Can We Trust a Volatility Forecast?

- Question of the seminar
  - **can we trust a volatility estimate, its uncertainty and its forecasts?**
- Deliverables
  - the A1 table and forecasts in class; A2 at home; A3–A6 if chosen as extensions
  - one figure and one table per B exercise
  - the corrected C1 code (an answer generated by artificial intelligence, AI)
  - one interpretation sentence per exercise
- Notebook for today: chapter5_seminar_notebook.ipynb (Colab)
  - each task names its notebook section
- Exercises, types and order: the Exercise Map
- Seminar: Prof. Daniel Traian Pele, danpele@ase.ro; Antoaneta Amza, PhD candidate, antoaneta.amza@csie.ase.ro

Exercise Map

Exercise	Question	Type	Model
A1	how does the variance update, and what is the 5-day risk?	Solved	–
B1	does GARCH-t fit the S&P 500, and do its residuals pass?	Solved	–
B2	does bad news raise S&P 500 volatility more?	Solved	–
B3	the same question for the BET	Proposed	B2
C1	what is wrong in an AI answer?	Proposed	A1, B1
After lecture 5, at home: the delta method, the parametric bootstrap and forecast losses			
A2	how precise are long-run volatility and half-life?	Solved	–
B4	is $\alpha + \beta = 1$? a bootstrap and a boundary test	Solved	–
B5	does anything forecast better than GARCH-t?	Solved	–
Homework	one transfer exercise: B6, B7 or B8	Proposed	B1, B4, B5
Extension packs	A3, A4, A5, A6, B5 Extended, C2, C3 (optional)	Solved or Proposed	marked on each task

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

What Units Do the Data Use?

- A simple price fall of 12 percent means a price ratio of 0.88. The corresponding log return is about minus 12.8 percent. Here we use log returns expressed in percent, not simple percentage changes.
- VIX (Cboe Volatility Index) is an annualised percentage quotation of 30-day implied volatility for the S&P 500 (Standard & Poor's). Its square is not yet a daily variance: the exercise states the conversion used.
- BTC/USD means a Bitcoin price in United States dollars. BET (Bucharest Exchange Trading) is the benchmark index of the Bucharest Stock Exchange (Bursa de Valori București, BVB); United States (US) and Romanian holidays need not coincide.
- Annualisation uses the observed number of returns per year, separately for each series. Always distinguish a daily variance, an annual variance and annual volatility.

Data Used

Series	Price used	Sample	Calendar, filter	Returns
S&P 500	close	2000-01-03 to 2026-09-18	weekdays; 0 non-trading days dropped	6 717
BET	close	2000-01-05 to 2026-09-18	weekdays; 0 non-trading days dropped	6 670
Bitcoin	close (BTC/USD)	2014-09-17 to 2026-09-18	7 days a week, no filter	4 384
VIX	close	C3 only	weekdays	–
11 US stocks	adjusted close	C3 only, 1995–2019	aligned with the VIX on common days	–

- Daily prices of the course data, one series per asset (field used: close / adjusted close); returns $r_t = 100\,(\ln P_t - \ln P_{t-\mathbf{1}})$ in %, each series on its own calendar
- Non-trading days, a heuristic rule: a weekday that repeats the previous close and shows no trading (reported zero volume, open = high = low = close, or a copy of the previous row); a missing volume does not count as zero
- The reason for each exclusion is kept and checked against the exchange calendar; a genuine unchanged close (a zero return) is kept; no missing values remain after the filter
- Annualisation with the observed frequency A (returns per year): 251.5 (S&P 500), 249.8 (BET), 365.3 (Bitcoin) returns per year

Setup: Run and Verify

```
# notebook section "Data" defines
# load_close(name) and pct_returns(name)
r = pct_returns('sp500')    # 100 * diff(log P)
print(len(r), r.index[0], r.index[-1])
print(r.head(3).round(4))

from arch import arch_model
res = arch_model(r, p=1, q=1,
                 dist='t').fit(disp='off')
print(res.params)
```

- Expected: 6 717 returns, 2000-01-04 to 2026-09-18
- First three returns (%): -3.9099, 0.1920, 0.0955
- GARCH-t (Generalised AutoRegressive Conditional Heteroskedasticity with Student-t innovations): $\hat{\alpha} = 0.1225$, $\hat{\beta} = 0.8723$, $\hat{\nu} = 6.25$
- If a count or a value differs, stop and check the file, the field, the dates and the filter
- Packages: `numpy`, `pandas`, `scipy`, `statsmodels`, `matplotlib`, `arch`

Returns and the Conditional Variance

- **Daily log return** in %, from the closing prices P_t

$$r_t = 100\,(\ln P_t - \ln P_{t-1})$$

- **Decomposition**: a constant mean plus a shock whose size changes over time; i.i.d. means independent and identically distributed

$$r_t = \mu + \varepsilon_t, \qquad \varepsilon_t = \sigma_t z_t, \qquad z_t \text{ i.i.d.}, \ \mathbb{E}z_t = 0, \ \mathrm{Var}z_t = 1$$

  - μ: mean return; ε_t: shock; z_t: innovation (Normal or Student-t, scaled to variance 1)

- **Conditional variance**: the variance of today's return given yesterday's information

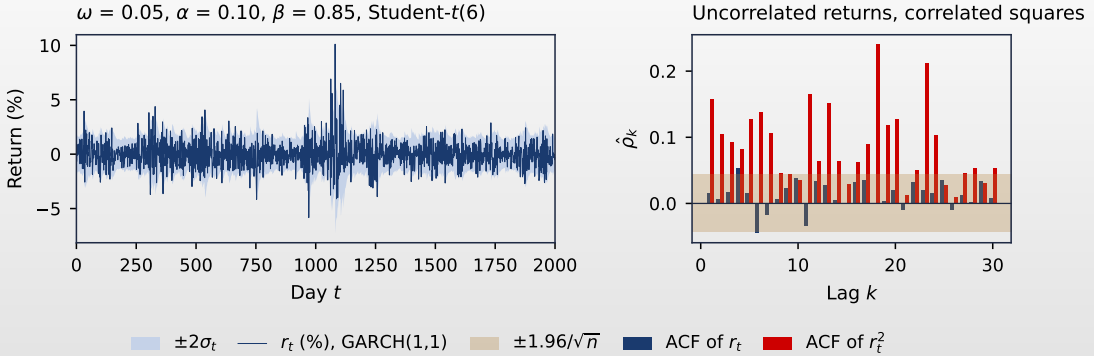
$$\sigma_t^2 = \mathrm{Var}(r_t \mid \mathcal{F}_{t-1}) = \mathbb{E}\left[\varepsilon_t^2 \mid \mathcal{F}_{t-1}\right]$$

  - $\mathcal{F}_{t-1}$: all prices up to the close of day $t-1$; σ_t is known one day ahead

What Does Volatility Clustering Look Like?

- Volatility clustering means that large moves tend to follow large moves and small moves tend to follow small moves. The signs can remain hard to predict while the size of the next move is partly predictable.
- **What do you think**: on the simulated path of the next slide, $\hat{\rho}_1 = 0.016$ for r_t and 0.158 for r_t^2; does this dependence contradict weak-form efficiency?
  - **Answer**: volatility clustering is compatible with a model with martingale-difference shocks, $\mathbb{E}[\varepsilon_{t+1} \mid \mathcal{F}_t] = 0$; on its own it does not refute weak-form efficiency
  - In GARCH, the conditional variance σ_{t+1}^2 is predictable, not the mean

Volatility Clustering in a Simulated Path



- Left: 2 000 simulated days; calm and turbulent periods alternate, and $\pm 2\sigma_t$ follows them (5.9% of the days fall outside)
- Right: sample autocorrelation $\hat{\rho}_k$ of r_t and of r_t^2 at lag k (Chapter 1)
  - $\hat{\rho}_1 = 0.016$ for r_t, inside the band $\pm 1.96/\sqrt{n}$; $\hat{\rho}_1 = 0.158$ for r_t^2: the size of the moves is predictable
  - A small autocorrelation of r_t does not rule out nonlinear predictability

ARCH and GARCH

- **ARCH(1)** Engle (1982): today's variance reacts to yesterday's squared shock

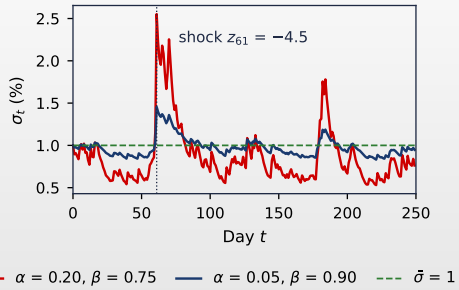
$$\sigma_t^2 = \omega + \alpha \varepsilon_{t-1}^2$$

- **GARCH(1,1)** Bollerslev (1986) adds memory of yesterday's variance

$$\sigma_t^2 = \omega + \alpha \varepsilon_{t-1}^2 + \beta \sigma_{t-1}^2$$

  - $\omega > 0$: level; $\alpha \geq 0$: reaction to news; $\beta \geq 0$: memory
  - the constraints keep $\sigma_t^2 > 0$

- Chart: the same shocks, the same $\alpha + \beta = 0.95$
  - large α: spiky volatility, peak 2.55% after $z = -4.5$
  - large β: smooth volatility, peak 1.46%



Persistence, Long-Run Variance and Half-Life

- **Persistence** $p = \alpha + \beta$; taking expectations, $\mathbb{E}\sigma_t^2 = \omega + p\,\mathbb{E}\sigma_{t-1}^2$, so for $p < 1$

$$\bar{\sigma}^2 = \mathbb{E}\sigma_t^2 = \frac{\omega}{1-p}$$

- **Forecasts** revert to $\bar{\sigma}^2$; the gap shrinks by p each day

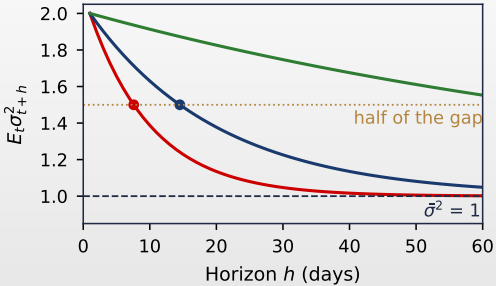
$$\mathbb{E}_t\sigma_{t+h}^2 = \bar{\sigma}^2 + p^{\,h-1}\left(\sigma_{t+1}^2 - \bar{\sigma}^2\right)$$

  - $\mathbb{E}_t = \mathbb{E}[\cdot \mid \mathcal{F}_t]$; h: horizon in days; σ_{t+1}^2 known at the close of t

- **Half-life**: days until half of the gap has gone, $p^h = 1/2$

$$h_{1/2} = \frac{\ln 0.5}{\ln p}$$

  - $p = 0.95$: $-0.693/(-0.0513) = 13.5$ days; $h_{1/2} \to \infty$ as $p \to 1$



$p = 0.90$, $h_{1/2} = 6.6$ days
$p = 0.95$, $h_{1/2} = 13.5$ days
$p = 0.99$, $h_{1/2} = 69.0$ days

Units and Multi-Day Variance

- Returns in %, so variances in $\%^2$; A: number of returns per year

$$\text{annual volatility} = \sqrt{A\,\bar{\sigma}^2}$$

  - a square root of a variance, never $A\,\bar{\sigma}^2$; $\bar{\sigma}^2 = 1$ and $A = 252$ give $\sqrt{252} = 15.9\%$ a year

- **Variance of an H-day return**: the shocks are uncorrelated (a martingale difference, $\mathbb{E}_t \varepsilon_{t+h} = 0$), so the variances add up

$$V_H = \text{Var}_t\Big(\sum_{h=1}^{H} \varepsilon_{t+h}\Big) = \sum_{h=1}^{H} \mathbb{E}_t \sigma^2_{t+h}$$

  - square-root-of-time rule: $V_H \approx H\sigma^2_{t+1}$; it holds exactly only if $\sigma^2_{t+1} = \bar{\sigma}^2$ (A1)

Testing for ARCH Effects

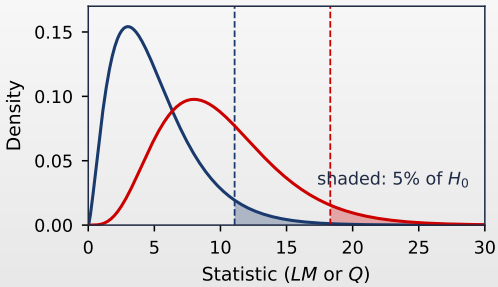
- H_0: no ARCH effects, constant conditional variance
- **Lagrange Multiplier (LM) test** Engle (1982): regress the squared demeaned return on its q lags

$$\hat{\varepsilon}_t^2 = c_0 + c_1 \hat{\varepsilon}_{t-1}^2 + \cdots + c_q \hat{\varepsilon}_{t-q}^2 + u_t, \quad LM = nR^2$$

  - $\hat{\varepsilon}_t = r_t - \bar{r}$; n: observations; R^2: coefficient of determination; $LM \sim \chi_q^2$ under H_0
- **Ljung–Box** on r_t^2: joint size of m autocorrelations

$$Q(m) = n(n+2) \sum_{k=1}^{m} \frac{\hat{\rho}_k^2}{n-k}$$

  - $\hat{\rho}_k$: autocorrelation of r_t^2 at lag k; $Q(m) \sim \chi_m^2$ under H_0
- Reject if the statistic exceeds the 5% critical value (shaded area): 11.07 for $q = 5$, 18.31 for $m = 10$
  - The χ^2 laws for r_t^2 require $\mathbb{E} r_t^4 < \infty$
  - On $\hat{z}_t^2$ (after GARCH) they require $\mathbb{E} z_t^4 < \infty$ and a correction for the estimated parameters



Estimating GARCH by Maximum Likelihood

- Maximum Likelihood Estimation (MLE): run the recursion for $\sigma_t^2(\theta)$, then choose θ to maximise the log-likelihood

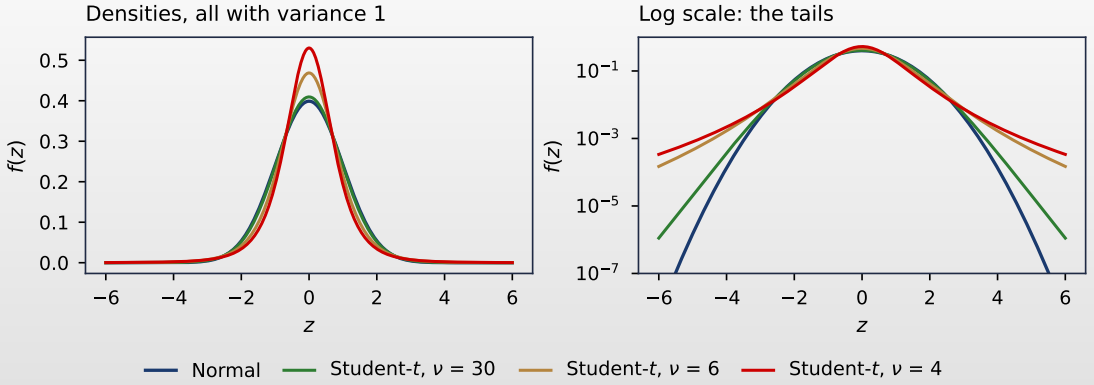
$$\ell(\theta) = \sum_{t=1}^{n} \left[\ln f\left(\frac{\varepsilon_t}{\sigma_t}\right) - \ln \sigma_t \right], \qquad \hat{\theta} = \arg\max_{\theta} \ell(\theta)$$

  - $\theta = (\mu, \omega, \alpha, \beta, \nu)$: all parameters; $\varepsilon_t = r_t - \mu$; f: density of z_t
  - $-\ln \sigma_t$: the Jacobian of $z_t = \varepsilon_t / \sigma_t$
- Normal z_t: $\ln f(z) = -\frac{1}{2}\ln(2\pi) - \frac{1}{2}z^2$; Student-t with $\nu > 2$ degrees of freedom, scaled to variance 1

$$f(z) = \frac{\Gamma\left(\frac{\nu+1}{2}\right)}{\Gamma\left(\frac{\nu}{2}\right)\sqrt{\pi(\nu-2)}} \left(1 + \frac{z^2}{\nu - 2}\right)^{-\frac{\nu+1}{2}}$$

  - Γ: gamma function; small ν: heavy tails; $\nu \to \infty$: the Normal; $\mathbb{E}z_t^4$ is finite only if $\nu > 4$

Student-t Innovations in a Chart



- All four densities have mean 0 and variance 1; a smaller ν puts more mass in the centre and in the tails
- Probability of a move beyond 4 standard deviations, $P(|z_t| > 4)$: Normal 0.0063%, $\nu = 30$: 0.026%, $\nu = 6$: 0.27%, $\nu = 4$: 0.48%

Quasi-Maximum Likelihood and Robust Standard Errors

- If the equation for σ_t^2 is correctly specified, the estimator that maximises the Normal likelihood stays consistent even for non-Normal innovations: Gaussian quasi-maximum likelihood (QMLE)
- The two matrices of the likelihood, with ℓ_t the contribution of day t

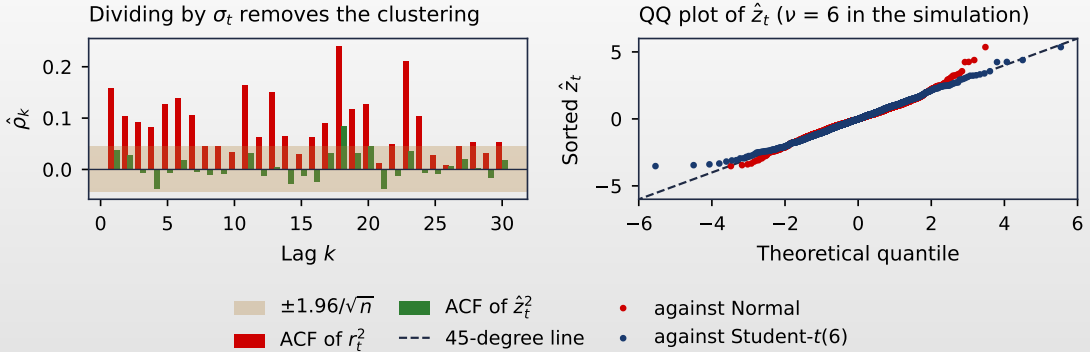
$$\mathrm{A} = -\frac{1}{n}\sum_t \frac{\partial^2 \ell_t}{\partial\theta\,\partial\theta^\top}, \qquad \mathrm{B} = \frac{1}{n}\sum_t \mathrm{s}_t\mathrm{s}_t^\top, \quad \mathrm{s}_t = \frac{\partial \ell_t}{\partial\theta}$$

  - A: mean curvature (minus the Hessian); s_t: score; B: variance of the score
- Robust (sandwich) covariance Bollerslev & Wooldridge (1992), against the classical one

$$\widehat{\mathrm{Var}}(\hat\theta) = \frac{1}{n}\mathrm{A}^{-1}\mathrm{B}\mathrm{A}^{-1} \quad \text{against} \quad \frac{1}{n}\mathrm{A}^{-1}$$

  - if f is the true density, A = B and both coincide
  - robust / classical standard-error (s.e.) ratio above 1: tails heavier than assumed
  - t-statistic of a parameter: estimate / robust s.e.

Checking the Fit: Standardised Residuals



- **Standardised residual** $\hat{z}_t = \hat{\varepsilon}_t / \hat{\sigma}_t$: i.i.d. with mean 0 and variance 1 if the model is right
  - Ljung–Box on $\hat{z}_t$: mean; on $\hat{z}_t^2$ and LM: variance (left: $\hat{\rho}_1$ falls from 0.158 to 0.038)
  - quantile–quantile (QQ) plot of $\hat{z}_t$ against the fitted f: the tails (right: S-shape against the Normal)

The Sign-Bias Test

- **Sign-bias test** Engle & Ng (1993): after fitting the model, does $\hat{z}_t^2$ still depend on the sign of yesterday's shock?

$$\hat{z}_t^2 = b_0 + b_1 S_{t-1}^- + b_2 S_{t-1}^- \hat{\varepsilon}_{t-1} + b_3 S_{t-1}^+ \hat{\varepsilon}_{t-1} + u_t$$

  - $S_{t-1}^- = 1(\hat{\varepsilon}_{t-1} < 0)$, with $1(\cdot)$ the indicator (1 if true, 0 otherwise); $S_{t-1}^+ = 1 - S_{t-1}^-$; u_t: regression error
  - b_1: sign bias (a negative day raises $\hat{z}_t^2$ by itself)
  - b_2, b_3: size bias of negative and positive shocks (the slope in $\hat{\varepsilon}_{t-1}$ on each side)
- Each slope: t-statistic $\hat{b}_j / \widehat{\text{s.e.}}(\hat{b}_j)$, compared with ± 1.96
- Joint test of $b_1 = b_2 = b_3 = 0$: $nR^2 \sim \chi_3^2$ under H_0 (no asymmetry left), 5% critical value 7.81

Asymmetric Models: GJR and EGARCH

- Asymmetric volatility (stylised fact 9, Chapter 1): bad news raises volatility more than good news of the same size
  - possible mechanisms: the leverage effect (a falling price raises the debt-to-equity ratio) and volatility feedback
- **Glosten–Jagannathan–Runkle (GJR)** Glosten, Jagannathan & Runkle (1993): negative shocks get an extra slope γ

$$\sigma_t^2 = \omega + (\alpha + \gamma I_{t-1})\varepsilon_{t-1}^2 + \beta\sigma_{t-1}^2, \qquad I_{t-1} = 1(\varepsilon_{t-1} < 0)$$

  - slope α for good news, $\alpha + \gamma$ for bad news; $\gamma > 0$: bad news weighs more; $\gamma = 0$: GARCH
  - symmetric z_t: $\mathbb{E}[z_t^2\, 1(z_t < 0)] = 1/2$, so the persistence is $p = \alpha + \gamma/2 + \beta$
- **Exponential GARCH (EGARCH)** (used in B5): the logarithm of the variance, so no sign constraints are needed

$$\ln \sigma_t^2 = \omega + \alpha\big(|z_{t-1}| - \mathbb{E}|z_{t-1}|\big) + \gamma z_{t-1} + \beta \ln \sigma_{t-1}^2$$

  - α: effect of the size of the shock; $\gamma < 0$: a negative z_{t-1} raises the variance more

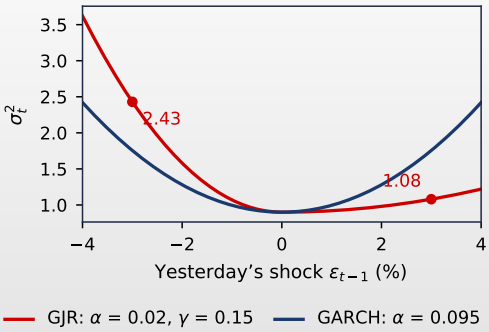
The News Impact Curve

- **News Impact Curve (NIC)** Engle & Ng (1993): tomorrow's variance as a function of today's shock, with σ^2_{t-1} held fixed

$$NIC(\varepsilon) = \omega + (\alpha + \gamma\, 1(\varepsilon < 0))\, \varepsilon^2 + \beta \bar{\sigma}^2_s$$

  - $\bar{\sigma}^2_s$: a fixed level, e.g. the sample variance
  - symmetric parabola for GARCH, steeper on the left for GJR

- Example: $\omega = 0.02$, $\alpha = 0.02$, $\gamma = 0.15$, $\beta = 0.88$, $\bar{\sigma}^2_s = 1$
  - $\varepsilon = -3$: $0.02 + 0.17 \times 9 + 0.88 = 2.43$
  - $\varepsilon = +3$: $0.02 + 0.02 \times 9 + 0.88 = 1.08$
  - GARCH with the same persistence ($\alpha = 0.095$): 1.755 for both signs



The Likelihood-Ratio Test

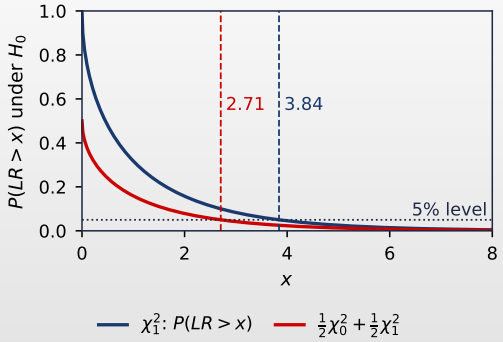
- Nested models: GARCH is GJR with $\gamma = 0$; the Likelihood Ratio (LR) statistic compares the two maximised fits

$$LR = 2\left(\ell_{\mathbf{1}} - \ell_{\mathbf{0}}\right) \geq 0$$

  - $\ell_{\mathbf{1}}$, $\ell_{\mathbf{0}}$: maximised log-likelihoods of the larger (GJR) and the smaller (GARCH) model
  - $LR \sim \chi^2_{\mathbf{1}}$ under $H_{\mathbf{0}}$ if the density is right and all other parameters are interior: critical values 3.84 (5%), 6.63 (1%)
- **Boundary**: if $H_{\mathbf{0}}$ puts a parameter on its bound (e.g. $\alpha = 0$ with $\alpha \geq 0$ in ARCH(1)), half of the time $LR = 0$ Andrews (2001)

$$LR \sim \tfrac{\mathbf{1}}{\mathbf{2}}\chi^2_{\mathbf{0}} + \tfrac{\mathbf{1}}{\mathbf{2}}\chi^2_{\mathbf{1}}, \qquad \text{5\% critical value 2.71}$$

  - $\chi^2_{\mathbf{0}}$: a point mass at 0; the chart shows the two tail probabilities



Information Criteria

- Akaike Information Criterion (AIC) and Bayesian Information Criterion (BIC): fit penalised by parameter count (Chapter 2); lower is better

$$\text{AIC} = -2\ell + 2k, \qquad \text{BIC} = -2\ell + k \ln n$$

  - ℓ: maximised log-likelihood; k: number of parameters; n: number of observations
  - BIC penalises more as soon as $\ln n > 2$, i.e. $n \geq 8$
- Link with LR, one extra parameter: $\Delta\text{AIC} = \text{AIC}_1 - \text{AIC}_0 = 2 - LR$
  - AIC prefers the larger model when $LR > 2$, BIC when $LR > \ln n$, the LR test at 5% when $LR > 3.84$
  - with $n \approx 6\,700$ days, $\ln n = 8.8$: BIC is the strictest of the three thresholds
- Depth: Chapter 5, section "Asymmetry and Model Choice"

A1 [Solved]: The Variance Recursion — Task (1/2)

- **Question**: how does GARCH(1,1) turn four shocks into four daily variances?
- **Inputs**: $\omega = 0.05$, $\alpha = 0.10$, $\beta = 0.85$ (returns in %); start $\sigma_1^2 = \bar{\sigma}^2$; shocks $\varepsilon_1, \ldots, \varepsilon_4 = -2.0, 0.5, 3.0, -1.0$; no data file
- **Tasks**
  1. Compute the persistence p and the long-run variance $\bar{\sigma}^2 = \omega/(1-p)$.
  2. Complete the worksheet below row by row: news term $\alpha\varepsilon_{t-1}^2$, memory term $\beta\sigma_{t-1}^2$ and σ_t^2 for $t = 3, 4, 5$ (row 2 is filled in as an example).
  3. Say which shock moves the variance most, and why the variance of day 5 is known at the close of day 4.
- **Report**: the completed table and one sentence
- Notebook: section A1

t	ε_{t-1}	$\alpha\varepsilon_{t-1}^2$	$\beta\sigma_{t-1}^2$	σ_t^2
1	–	–	–	1.0000
2	−2.0	0.4000	0.8500	1.3000
3	0.5	…	…	…
4	3.0	…	…	…
5	−1.0	…	…	…

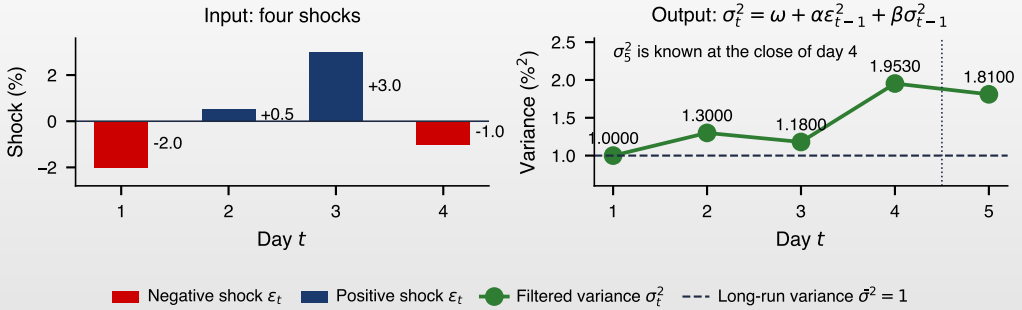
Back Exercise list

A1 [Solved]: Solution — the Filtered Variances

t	ε_{t-1}	$\alpha\varepsilon^2_{t-1}$	$\beta\sigma^2_{t-1}$	σ^2_t
1	–	–	–	1.0000
2	−2.0	0.4000	0.8500	1.3000
3	0.5	0.0250	1.1050	1.1800
4	3.0	0.9000	1.0030	1.9530
5	−1.0	0.1000	1.6600	1.8100

- $p = 0.95$, $\bar{\sigma}^2 = 0.05/0.05 = 1$; each row adds $\omega = 0.05$ to the two terms
- Row 2: $0.05 + 0.10 \times 4 + 0.85 \times 1 = 1.3000$
- The $+3\%$ shock adds $\alpha\varepsilon^2 = 0.90$: the variance jumps to 1.9530 on day 4
- Then the expected distance from $\bar{\sigma}^2$ is multiplied by $p = \alpha + \beta = 0.95$ each day
- σ^2_5 uses only ε_4 and σ^2_4, both known at the close of day 4: the one-step forecast is exact within the model
- Interpretation: the sign does not matter, the size does; one large shock raises the variance for many days

A1 [Solved]: the Recursion in a Chart



- Each bar is a shock; each point the variance it produces the next day; σ_5^2 is the one-step forecast made at the close of day 4

MFM_ch5_seminar

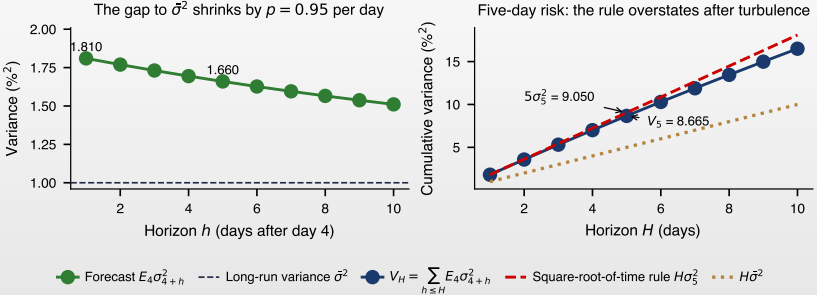
A1 [Solved]: Forecasting from Day 4 — Task (2/2)

- **Questions**
  - what is the 5-day variance after these shocks?
  - is the square-root-of-time rule right?
- **Inputs**: the model and $\sigma_5^2 = 1.8100$ from Task 1; forecast origin: the close of day 4
- **Tasks**
  1. Take $\mathbb{E}_4$ of the recursion and use $\mathbb{E}_4\varepsilon_t^2 = \mathbb{E}_4\sigma_t^2$ for $t \geq 5$ to derive $\mathbb{E}_4\sigma_{4+h}^2 = \bar{\sigma}^2 + p^{h-1}(\sigma_5^2 - \bar{\sigma}^2)$.
  2. Compute the five forecasts for $h = 1, \ldots, 5$.
  3. Explain why the 5-day return variance is the sum $V_5 = \sum_{h=1}^{5} \mathbb{E}_4\sigma_{4+h}^2$ (hint: the ε_t are uncorrelated, a martingale difference), and compute V_5 and $\sqrt{V_5}$.
  4. Compare V_5 with $5\sigma_5^2$ and derive the sign of $V_H - H\sigma_5^2$ as a function of $\sigma_5^2 - \bar{\sigma}^2$.
  5. Compute the half-life $\ln 0.5 / \ln p$.
- **Report**: the forecast formula, five forecasts, V_5 and $\sqrt{V_5}$, the half-life; interpretation: after a +3% shock, does the square-root-of-time rule overstate or understate the 5-day risk?
- Notebook: section A1

A1 [Solved]: Solution — Forecasts and Five-Day Risk

- Derivation: $\mathbb{E}_4\sigma^2_{4+h} = \omega + p\,\mathbb{E}_4\sigma^2_{4+h-1}$; subtract $\bar\sigma^2 = \omega + p\bar\sigma^2$: the gap shrinks by p per day, starting from the known σ^2_5
- Gap 0.8101; forecasts $1.8100, 1.7695, 1.7311, 1.6945, 1.6598$;
  $V_5 = 5 + 0.8101 \times (1 - 0.95^5)/0.05 = 8.665$, $\sqrt{V_5} = 2.944\%$; half-life $-0.6931/(-0.0513) = 13.5$ days
- $V_H - H\sigma^2_5 = (\bar\sigma^2 - \sigma^2_5)\sum_{j<H}(1 - p^j)$: here $5\sigma^2_5 = 9.050 > V_5$. Interpretation: after turbulence the rule overstates the 5-day risk, in calm times it understates it, because the variance reverts to $\bar\sigma^2$

A1 [Solved]: Forecasts and Five-Day Risk in a Chart



- Left: the forecasts converge to $\bar{\sigma}^2 = 1$; right: V_H stays below $H\sigma_5^2$ because the variance reverts

MFM_ch5_seminar

B1 [Solved]: GARCH-t on the S&P 500 — Task (1/2)

- **Questions**
  - is there clustering to model?
  - what does a GARCH(1,1)-t estimate say about it?
- **Inputs**: S&P 500, 6 717 daily returns (Data Used); mean $r_t = \mu + \varepsilon_t$, Student-t innovations
- **Tasks**
  1. Test H_0: no ARCH effects with Ljung–Box $Q(10)$ on r_t^2 and Engle's LM test with 5 lags on $r_t - \bar{r}$, at 5%.
  2. Estimate GARCH(1,1)-t by maximum likelihood and check that the optimiser converged.
  3. Report every parameter with its robust Bollerslev & Wooldridge (1992) and classical standard error, and plot r_t with the bands $\pm 2\hat{\sigma}_t$.
- **Report**: the two test decisions, the estimates table, the chart
- Notebook: section B1

```
from arch import arch_model
am = arch_model(r, mean='Constant',
                vol='GARCH', p=1, q=1,
                dist='t')
res = am.fit(disp='off')
print(res.convergence_flag)  # 0 = ok
rob = res.std_err            # robust
cl = am.fit(disp='off',
            cov_type='classic').std_err
sig = res.conditional_volatility
```

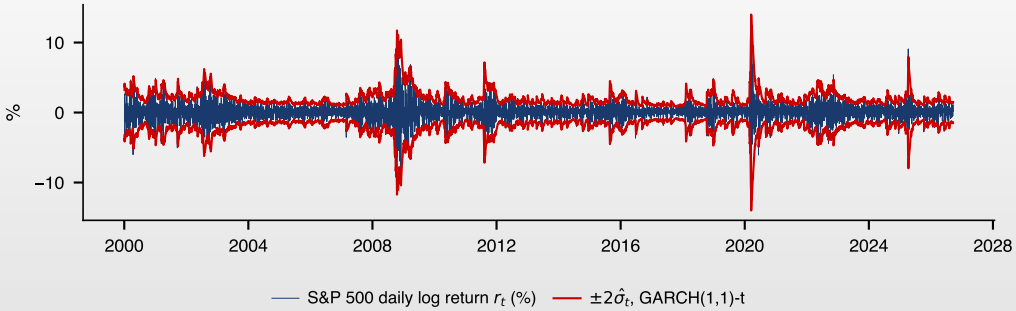
Back Exercise list

B1 [Solved]: Solution — Tests, Estimates, Fit

Parameter	Estimate	Robust s.e.	Classical s.e.
μ	0.0787	0.0089	0.0090
ω	0.0157	0.0032	0.0029
α	0.1225	0.0107	0.0101
β	0.8723	0.0103	0.0096
ν	6.25	0.48	0.48

- (1) $Q(10) = 6\,135$ and $\mathrm{LM}(5) = 1\,685$ against 18.3 and 11.1: no-ARCH is rejected ($p < 0.001$)
  - The χ^2 law requires $\mathbb{E}r_t^4 < \infty$; the fitted GARCH-t, with $\kappa_z = \mathbb{E}z_t^4$, gives $(\hat{\alpha} + \hat{\beta})^2 + (\kappa_z - 1)\hat{\alpha}^2 = 1.060$, above 1: an infinite fourth moment in the model (Chapter 5: BET 1.273, BET-TR, with reinvested dividends, 1.179, gold 1.021)
  - This does not prove that the data have an infinite moment, but the χ^2 law is read with caution
- (2) Converged after 14 iterations, $\ell = -9\,065.0$; persistence 0.9948
- (3) Robust/classical ratios 1.06 (α), 1.08 (β): the Student-t density is itself misspecified (skewness of $\hat{z}_t$: -0.58); the sandwich adjusts uncertainty, not a possible scale bias from the wrong density

B1 [Solved]: the Fitted Volatility



⊡ Interpretation: $\hat{\beta}$ carries yesterday's variance, $\hat{\alpha}$ adds the news, $\hat{\nu}$ signals heavy tails; 5.2% of the days fall outside $\pm 2\hat{\sigma}_t$

MFM_ch5_seminar

B1 [Solved]: Diagnosing the Residuals — Task (2/2)

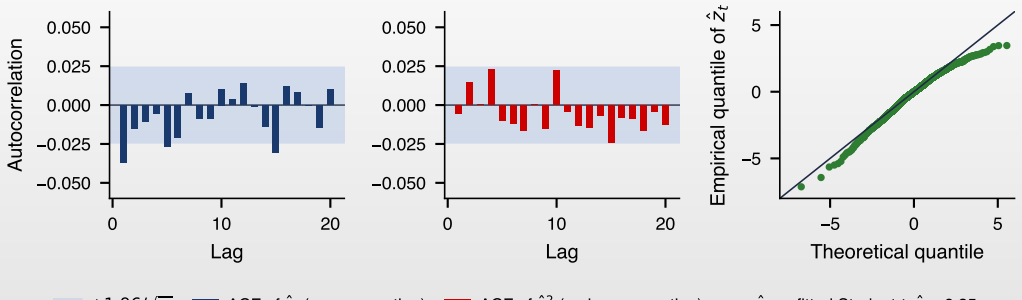
- **Question**: does the fitted GARCH-t leave any structure in $\hat{z}_t = (r_t - \hat{\mu})/\hat{\sigma}_t$?
- **Inputs**: the fit of Task 1; each test has H_0: the model captures that feature
- **Tasks**
  1. Test the mean equation with Ljung–Box $Q(10)$ on $\hat{z}_t$ (H_0: no autocorrelation up to lag 10).
  2. Test the variance equation with $Q(10)$ on $\hat{z}_t^2$ and the LM(5) test (H_0: no clustering left).
  3. Run the sign-bias regression Engle & Ng (1993) (slide: The Sign-Bias Test); report the three t-statistics and the joint $nR^2 \sim \chi^2_3$ (H_0: no asymmetry left).
  4. Plot the ACFs of $\hat{z}_t$ and $\hat{z}_t^2$ and a QQ plot of $\hat{z}_t$ against the fitted standardised Student-t.
- **Report**: a table with statistic, p-value and decision for each test, and the three panels; interpretation:
  - which part of the model fails?
  - what would you change first?
- Notebook: section B1

B1 [Solved]: Solution — Diagnostics

Part	Test	p	Decision at 5%
Mean	$Q(10)$, $\hat{z}$: 21.8	0.016	reject
Variance	$Q(10)$, $\hat{z}^2$: 13.5	0.199	not rejected
	LM(5): 5.8	0.327	not rejected
Asymmetry	joint χ^2_3: 40.6	< 0.001	reject

- Sign bias $t = 3.57$, negative size $t = 1.17$, positive size $t = -2.61$
- QQ: 16 residuals below -4 against 8.6 expected; the left tail is longer than the fitted t

B1 [Solved]: Residual ACFs and QQ Plot



- All p-values use the naive χ^2 laws, approximate with estimated parameters Li & Mak (1994). Interpretation: clustering is captured, asymmetry is not: add a sign term (B2)

MFM_ch5_seminar

B2 [Solved]: GJR versus GARCH on the S&P 500 — Task

- **Questions**
  - does bad news raise S&P 500 volatility more than good news?
  - by how much?
- **Inputs**: B1 data; `arch_model(r, p=1, o=1, q=1, dist='t')` estimates GJR-GARCH(1,1)-t, which nests GARCH-t at $\gamma = 0$; news impact at σ^2_{t-1} = sample variance 1.47
- **Tasks**
  1. Estimate GJR-GARCH(1,1)-t and report $\hat{\alpha}$, $\hat{\gamma}$ with its robust t, and $\hat{\beta}$.
  2. Build the comparison table: number of parameters k, ℓ, AIC, BIC for GARCH-t and GJR-t.
  3. Test $H_0 : \gamma = 0$ with $LR = 2(\ell_{GJR} - \ell_{GARCH})$ against χ^2_1 and state the assumptions behind this law.
  4. Compute the next-day variance after shocks of -3% and $+3\%$ in both models and plot the two news impact curves.
- **Report**: the estimates, the comparison table, the LR decision and the news-impact chart; interpretation:
  - why is $\hat{\alpha} = 0$?
  - what does it do to the χ^2_1 law?
- Notebook: section B2

Back Exercise list

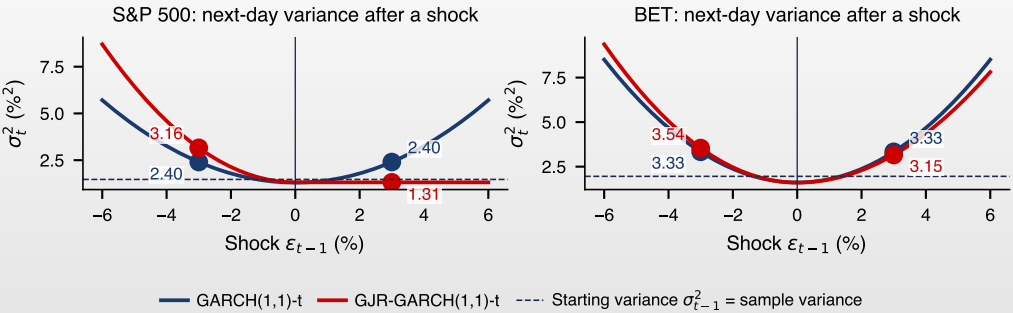
B2 [Solved]: Solution — Nested Comparison

Model	k	ℓ	AIC	BIC
GARCH(1,1)-t	5	−9 064.993	18 140.0	18 174.0
GJR-GARCH(1,1)-t	6	−8 946.905	17 905.8	17 946.7

- (1) $\hat{\alpha} = 0.000$, $\hat{\gamma} = 0.205$ (robust $t = 10.04$), $\hat{\beta} = 0.881$, $\hat{\omega} = 0.0186$, $\hat{\nu} = 6.95$: good news slope $\hat{\alpha}$, bad news $\hat{\alpha} + \hat{\gamma}$
- Why $\hat{\alpha} = 0$: once γ takes the negative shocks, positive shocks no longer raise tomorrow's variance; the estimate sits on the bound $\alpha \geq 0$; the bound affects the Wald t of $\hat{\alpha}$, not the LR test for γ
- (2) $\Delta\text{AIC} = -234.2$, $\Delta\text{BIC} = -227.4$ (GJR minus GARCH): both prefer GJR
- (3) $LR = 2(-8\,946.905 - (-9\,064.993)) = 236.17$, far above 3.84
- Assumptions of the χ^2_1 law: a correct Student-t density and true parameters under H_0 in the interior
  - The density is doubtful here: B1 found skewness
  - The boundary does not change the LR law: under the fitted GARCH-t null, $\hat{\alpha} = 0.122 > 0$; $\hat{\alpha} = 0$ in GJR is the estimate of the alternative Andrews (2001)
  - Qualified conclusion: an LR of 236.17 is some 60 times the 5% critical value; a calibrated check would simulate the LR under the fitted GARCH-t (as B4 does for Integrated GARCH (IGARCH))

MFM_ch5_seminar

B2 [Solved]: Solution — News Impact



- (4) S&P 500, GJR: -3% gives 3.16, $+3\%$ gives 1.31; GARCH: 2.40 for both signs
- With $\hat{\alpha} = 0$ a $+3\%$ shock adds no shock term: persistence below one lowers the starting variance from 1.47 to 1.31. Interpretation: in the S&P 500, only bad news raises tomorrow's variance; the BET panel is B3

MFM_ch5_seminar

B3 [Proposed]: GJR versus GARCH on the BET — Task

- **Question**: is the asymmetry of the Bucharest market as strong as that of the S&P 500?
- **Inputs**: BET, 6 670 daily returns (Data Used); the same two models and code as B2; $\ln n = 8.81$
- **Tasks**
  1. Estimate both models and fill the comparison table of B2 (k, ℓ, AIC, BIC) for the BET.
  2. Compute LR and place it against the four thresholds: 2 (AIC), 3.84 and 6.63 (χ^2_1 at 5% and 1%) and $\ln n$ (BIC); state each decision.
  3. Compute the BET news impact at $\pm 3\%$ from the sample variance, as in B2.
- **Report**: the table, the four decisions, the two news-impact values; interpretation:
  - why can LR and BIC disagree?
  - does a smaller $\hat{\gamma}$ prove that the BET is less asymmetric than the S&P 500?
- Model: B2
- Notebook: section B3

C1 [Proposed]: Critique an AI Answer — Task

- **Prompt**: "`px` holds daily S&P 500 closes. Fit a GARCH(1,1) with Student-t innovations using `arch` and give me the annualised long-run volatility and the half-life of a shock."
- **Input**: `px = load_close('sp500')`, 6 718 closes, 2000-01-03 to 2026-09-18
- **Answer's output**: `ann_vol` = 982.07, `half_life` = 0.06 days, and a `DataScaleWarning`
- **Tasks**
  1. Run the code and reproduce the output.
  2. Find the formula errors, the unit errors and the numerical risks, each with its line number.
  3. For each, give the evidence that exposes it (a warning, a number from A1 or B1) and the corrected line.
  4. Run the corrected code and record in `AI_ERRORS.md` the line, error type, evidence, corrected line and verification result.
- Model: A1, B1; annualisation: slide: Units and Multi-Day Variance and the slide "Data Used"

```
1  import numpy as np
2  from arch import arch_model
3  r = np.log(px).diff().dropna()
4  am = arch_model(r, vol='GARCH', p=1,
5                  q=1, dist='t')
6  res = am.fit(disp='off')
7  om = res.params['omega']
8  a = res.params['alpha[1]']
9  b = res.params['beta[1]']
10 lr_var = om / (1 - a - b)
11 ann_vol = lr_var * 252
12 half_life = np.log(0.5) / np.log(b)
```

Back Exercise list

The Delta Method

- **Delta method**: the standard error of a smooth function $g(\hat{\theta})$ of the estimates, by a first-order (linear) approximation

$$g(\hat{\theta}) \approx g(\theta) + \nabla g^\top (\hat{\theta} - \theta) \;\Rightarrow\; \widehat{\text{s.e.}}\, g(\hat{\theta}) = \sqrt{\nabla g^\top \, \hat{V} \, \nabla g}$$

  - ∇g: gradient of g at $\hat{\theta}$ (column of partial derivatives), $\nabla g^\top$ its transpose
  - $\hat{V}$: estimated (robust) covariance matrix of $\hat{\theta}$
- Example: the persistence $g = \alpha + \beta$ has $\nabla g = (0, 1, 1)^\top$ in the order (ω, α, β)

$$\widehat{\text{s.e.}}(\hat{p}) = \sqrt{V_{\alpha\alpha} + 2V_{\alpha\beta} + V_{\beta\beta}}$$

- Reliable only if g is nearly linear over the uncertainty range of $\hat{\theta}$; the half-life $\ln 0.5 / \ln p$ is very curved near $p = 1$ (A2)

The Parametric Bootstrap

- **Parametric bootstrap**: simulate from a fitted model, re-estimate on each simulated path; starred quantities ($\hat{p}^*$, LR^*) come from the simulated paths
- Algorithm, $b = 1, \ldots, B$
  - draw z_t^* from the fitted law, run the recursion $\sigma_t^{*\mathbf{2}} = \hat{\omega} + \hat{\alpha}\varepsilon_{t-\mathbf{1}}^{*\mathbf{2}} + \hat{\beta}\sigma_{t-\mathbf{1}}^{*\mathbf{2}}$, $\varepsilon_t^* = \sigma_t^* z_t^*$
  - re-estimate the model on the simulated path and keep $\hat{p}_b^*$ (or LR_b^*)
- Percentile interval: the 2.5% and 97.5% percentiles of $\hat{p}_{\mathbf{1}}^*, \ldots, \hat{p}_B^*$
- Bootstrap test: simulate under $H_{\mathbf{0}}$ (e.g. IGARCH, $\alpha + \beta = 1$) and count the draws at least as extreme

$$p_{\mathbf{boot}} = \frac{1 + \#\{b : LR_b^* \geq LR\}}{B + 1}$$

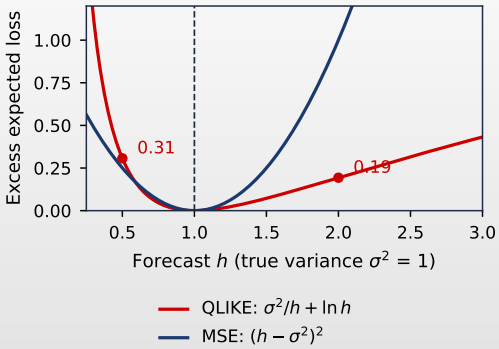
  - $\#\{\cdot\}$: number of draws that satisfy the condition; compare with the boundary law $\frac{\mathbf{1}}{\mathbf{2}}\chi_{\mathbf{0}}^{\mathbf{2}} + \frac{\mathbf{1}}{\mathbf{2}}\chi_{\mathbf{1}}^{\mathbf{2}}$ (slide: The Likelihood-Ratio Test)

The QLIKE Loss

- h_t: variance forecast for day t, made at the close of $t-1$; the true variance is never observed, so r_t^2 serves as a noisy proxy ($\mathbb{E}_{t-1} r_t^2 = \sigma_t^2$ if $\mu = 0$)
- **Quasi-likelihood (QLIKE) loss** Patton (2011)

$$L_t = \frac{r_t^2}{h_t} + \ln h_t$$

  - expected value $\sigma_t^2/h_t + \ln h_t$, lowest at $h_t = \sigma_t^2$; lower mean QLIKE is better
  - under-prediction costs more (with $\sigma_t^2 = 1$): $h = 0.5$ adds 0.31, $h = 2$ adds 0.19
  - robust to the noise of the proxy: the ranking of forecasts is the same as with the true variance



Diebold–Mariano and Mincer–Zarnowitz

- **Diebold–Mariano (DM) test** Diebold & Mariano (1995): are two forecasts equally accurate on average?

$$d_t = L_t^m - L_t^{GARCH\text{-}t}, \qquad DM = \frac{\bar{d}}{\widehat{\text{s.e.}}(\bar{d})}, \qquad \widehat{\text{s.e.}}(\bar{d})^2 = \frac{\hat{c}_0 + 2\sum_{j=1}^{J}\left(1 - \frac{j}{J+1}\right)\hat{c}_j}{n}$$

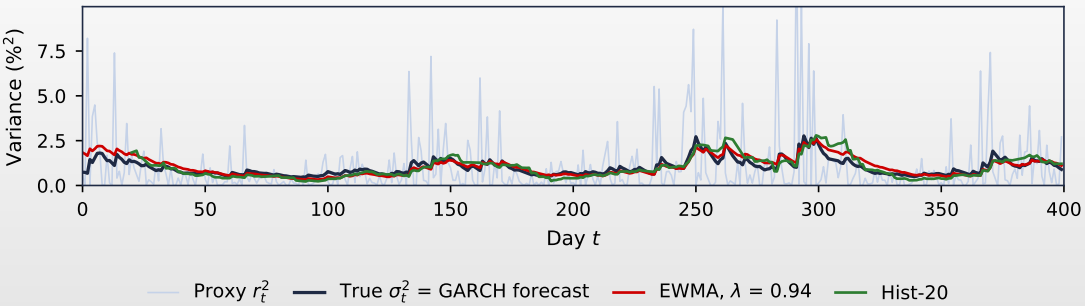
  - L_t^m: loss of model m; $\bar{d}$: mean of d_t; $\hat{c}_j$: autocovariance of d_t at lag j; a Heteroskedasticity and Autocorrelation Consistent (HAC) standard error with $J = 5$ lags
  - $DM \approx N(0,1)$ under equal accuracy; $DM < -1.96$: model m better, $DM > 1.96$: worse

- **Mincer–Zarnowitz (MZ) regression** Mincer & Zarnowitz (1969): are the forecasts unbiased in levels?

$$r_t^2 = a + b\, h_t + u_t, \qquad H_0 : a = 0,\ b = 1$$

  - Wald test $W = (\hat{\boldsymbol{\theta}} - \boldsymbol{\theta_0})^\top \hat{\mathrm{V}}^{-1} (\hat{\boldsymbol{\theta}} - \boldsymbol{\theta_0}) \sim \chi_2^2$, $\boldsymbol{\theta} = (a, b)^\top$, $\boldsymbol{\theta_0} = (0, 1)^\top$, $\hat{\mathrm{V}}$: HAC covariance
  - R^2: share of the variation of r_t^2 explained, low because the proxy is noisy

Benchmark Forecasts



- **Exponentially Weighted Moving Average (EWMA)** J.P. Morgan/Reuters (1996): $h_t = \lambda h_{t-\mathbf{1}} + (1-\lambda) r_{t-\mathbf{1}}^{\mathbf{2}}$, λ: decay factor (0.94); a GARCH with $\omega = 0$, $\alpha = 1-\lambda$, $\beta = \lambda$, no mean reversion
- **Hist-20**: $h_t = \frac{\mathbf{1}}{\mathbf{20}} \sum_{j=\mathbf{1}}^{\mathbf{20}} r_{t-j}^{\mathbf{2}}$, the mean of the last 20 squared returns
- Simulated GARCH path: mean QLIKE 0.930 (true variance), 0.959 (EWMA), 0.967 (Hist-20); Hist-20 lags behind every jump

A2 [Solved]: Long-Run Volatility and Half-Life — Task

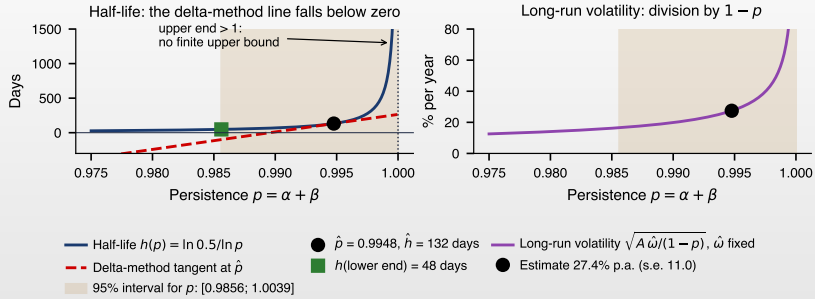
- **Question**: how precisely does the B1 fit pin down the long-run volatility and the half-life?
- **Inputs**: B1: $\hat{\omega} = 0.015669$, $\hat{\alpha} = 0.122499$, $\hat{\beta} = 0.872270$, $A = 251.5$ returns a year; robust covariance of $(\hat{\omega}, \hat{\alpha}, \hat{\beta})$, in this order: $\hat{V} = \begin{pmatrix} 9.977 \times 10^{-6} & 1.416 \times 10^{-5} & -2.234 \times 10^{-5} \\ \cdot & 1.151 \times 10^{-4} & -9.994 \times 10^{-5} \\ \cdot & \cdot & 1.066 \times 10^{-4} \end{pmatrix}$ (`res.param_cov`)
- **Tasks**
  1. Derive $\bar{\sigma}^2 = \omega/(1-\alpha-\beta)$ from the recursion and compute the long-run annual volatility $v = \sqrt{A\,\bar{\sigma}^2}$ and the half-life.
  2. Compute $\nabla\bar{\sigma}^2$, the quadratic form $\nabla^\top \hat{V} \nabla$ and, with $\partial v/\partial\bar{\sigma}^2 = A/(2v)$, the standard error of v.
  3. Compute $\widehat{\text{s.e.}}(\hat{p})$ from $\hat{V}$ and the delta-method standard error of $h(p) = \ln 0.5/\ln p$, defined for $0 < p < 1$ (here h is the half-life, not the horizon).
  4. Map the 95% interval for p into an interval for h and say what happens when it crosses 1.
- **Report**: v with its s.e., h with its s.e. and the transformed interval; interpretation: why does the linearisation fail for h?
- Notebook: section A2

Back Exercise list

A2 [Solved]: Solution — Gradients and the Half-Life Interval

- (1) $\hat{p} = 0.994769$, $\bar{\sigma}^2 = 2.995$, $v = 1.731 \times \sqrt{251.5} = 27.4\%$ (sample: 19.2%); $h = -0.6931/(-0.005245) = 132.2$ days
- (2) $\nabla\bar{\sigma}^2 = (1/(1-p), \omega/(1-p)^2, \omega/(1-p)^2) = (191.2, 572.6, 572.6)$; $\widehat{\text{s.e.}}(\bar{\sigma}^2) = 2.40$; $\widehat{\text{s.e.}}(v) = 11.0$ percentage points around 27.4%
- (3) $\widehat{\text{s.e.}}(\hat{p}) = \sqrt{V_{\alpha\alpha} + 2V_{\alpha\beta} + V_{\beta\beta}} = 0.0047$ (V_{ij}: entries of $\hat{V}$); $\partial h/\partial p = 25\,330$, so $\widehat{\text{s.e.}}(\hat{h}) = 118$ days
- (4) $[0.9856; 1.0039]$ crosses 1: $h \in [48; \infty)$, no finite upper bound within the stationary model. Interpretation: $h(p)$ is so curved near 1 that its tangent is meaningless; the question $p = 1$ needs a test (B4)

A2 [Solved]: the Half-Life Curve



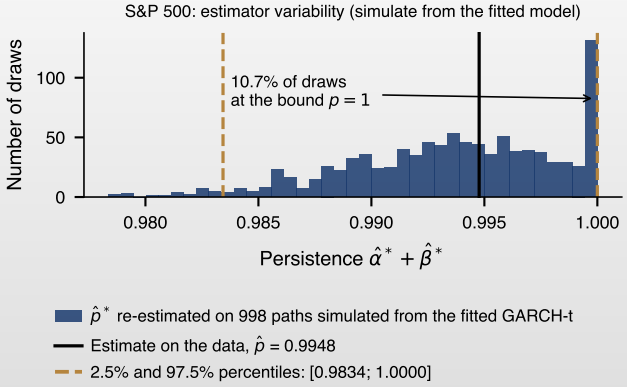
- The tangent (delta method) crosses zero inside the interval for p; the exact map $h(p)$ has no finite upper end

MFM_ch5_seminar

B4 [Solved]: Is $\alpha + \beta = 1$? Bootstrap and Boundary Test — Task

- **Questions**
  - how variable is $\hat{p} = \hat{\alpha} + \hat{\beta}$?
  - can the S&P 500 variance have a unit root?
- **Inputs**: B1 fit, $\hat{p} = 0.9948$; code `param_bootstrap`, `igarch_test` (notebook); seed 42 (43 under H_0); each path: 6 717 days after 500 discarded burn-in days, started at the sample variance; each fit imposes $\alpha + \beta \leq 1$ and keeps the best of two starting points
- **Tasks**
  1. Algorithm 1 (estimator variability): simulate $B = 999$ paths from the **fitted** GARCH-t, re-estimate p on each, and report the 2.5% and 97.5% percentiles and the share of draws at $p = 1$.
  2. Algorithm 2 (test): fit IGARCH-t ($\alpha + \beta = 1$), compute $LR = 2(\ell_{GARCH} - \ell_{IGARCH})$, simulate B paths **under** H_0, re-estimate both models on each path, and compute the bootstrap p-value (slide: The Parametric Bootstrap).
  3. Compare this p-value with the χ^2_1 and $\frac{1}{2}\chi^2_0 + \frac{1}{2}\chi^2_1$ values and decide at 5% and 10%.
- **Report**: two histograms, the interval, three p-values and the decision; interpretation: why is "the percentile interval reaches 1" not a test of $p = 1$ Andrews (2000)?
- Notebook: section B4

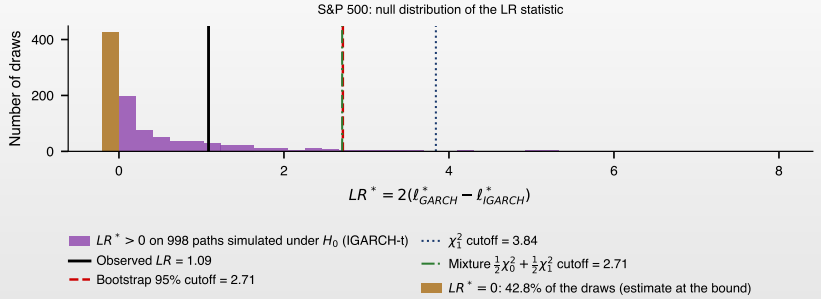
B4 [Solved]: Solution — Estimator Variability



- 95% percentile interval [0.9834; 1.0000]; bootstrap s.d. 0.0046 against the robust s.e. 0.0047: close; 10.7% of the draws sit exactly at the bound $p = 1$
- Half-life from 41 days to ∞ (delta method, A2: $[48; \infty)$, s.e. 118 days). Interpretation: near a bound the percentile interval describes spread, it is not automatically a valid confidence interval Andrews (2000)

MFM_ch5_seminar

B4 [Solved]: Solution — the Test under the Null



- IGARCH-t: $\hat{\alpha}_{\mathbf{0}} = 0.127$, $\hat{\nu}_{\mathbf{0}} = 6.03$; $LR = 1.09$; $LR^* = 0$ for 42.8% of the draws, close to the weight 1/2 of the mixture distribution; 167 of 998 draws exceed LR
- Non-converged replications, dropped: 1 of 999 under H_0, 1 of 999 from the fitted model
- p-values: 0.168 (bootstrap), 0.15 (mixture), 0.30 (naive χ^2_1): H_0 is not rejected at 5% or 10%. Interpretation: 27 years cannot tell a half-life of about 130 days from an infinite one

MFM_ch5_seminar

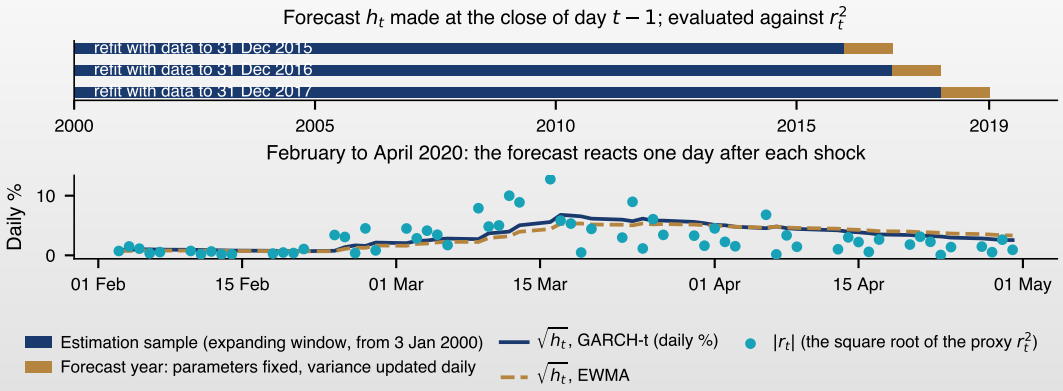
B5 [Solved]: Forecast Comparison on the S&P 500 — Task (1/2)

- **Question**: does any model forecast tomorrow's S&P 500 variance better than GARCH-t?
- **Inputs**: 2 693 forecast days, 2016-01-04 to 2026-09-18; GARCH-N (GARCH with Normal innovations), GARCH-t, GJR-t, EGARCH-t refit each January (expanding window); EWMA ($\lambda = 0.94$) and Hist-20 (slide: Benchmark Forecasts); realised variance proxied by r_t^2
- **Tasks**
  1. On the timeline chart, check that h_t uses only data up to the close of $t-1$ and that parameters are fixed within a year.
  2. Compute the QLIKE loss (slide: The QLIKE Loss) step by step for GARCH-t and EWMA on 2020-03-13, 2020-03-16 and 2020-03-17 (table below).
  3. Compute the mean QLIKE of all six models and the DM t-statistic of $d_t = L_t^m - L_t^{GARCH\text{-}t}$ (slide: Diebold–Mariano and Mincer–Zarnowitz); plot the cumulative $\sum d_t$.
- **Report**: the three QLIKE rows, a table of mean QLIKE and DM t, the loss chart; interpretation: which differences can be read as significant?
- Notebook: section B5

Day	r_t	r_t^2	h_t GARCH-t	h_t EWMA	QLIKE
2020-03-13	8.88	78.87	25.37	15.66	...
2020-03-16	−12.77	162.95	31.17	19.45	...
2020-03-17	5.82	33.90	46.04	28.06	...

Back Exercise list

B5 [Solved]: Solution — Forecast Timing



MFM_ch5_seminar

B5 [Solved]: Solution — QLIKE Step by Step

Day	L GARCH-t	L EWMA
2020-03-13	6.342	7.789
2020-03-16	8.668	11.346
2020-03-17	4.566	4.543

- e.g. 2020-03-16: $162.95/31.17 + \ln 31.17 = 8.668$; EWMA forecast lower, so a larger loss: under-prediction is punished
- First target 2016-01-04, origin 2015-12-31: refit on data to 31 December, no look-ahead bias (no future information)

MFM_ch5_seminar

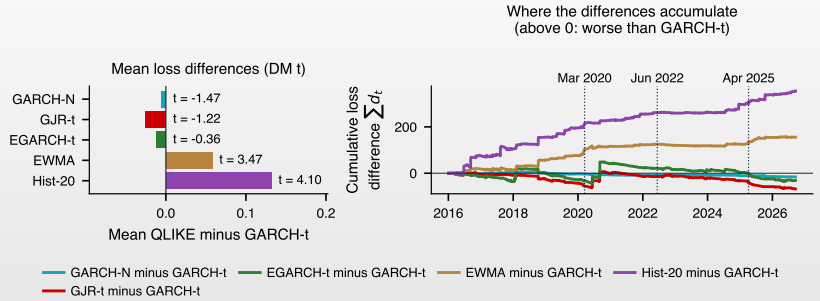
B5 [Solved]: Solution — Mean QLIKE and DM for the Six Models

Model	GARCH-N	GARCH-t	GJR-t	EGARCH-t	EWMA	Hist-20
QLIKE	0.7503	0.7562	0.7312	0.7447	0.8140	0.8879
DM t	-1.47	–	-1.22	-0.36	3.47	4.10

- S&P 500, 2 693 forecast days, 2016-01-04 to 2026-09-18; DM t of $d_t = L_t^m - L_t^{GARCH\text{-}t}$ with HAC errors (5 lags)

MFM_ch5_seminar

B5 [Solved]: Solution — Loss Differences



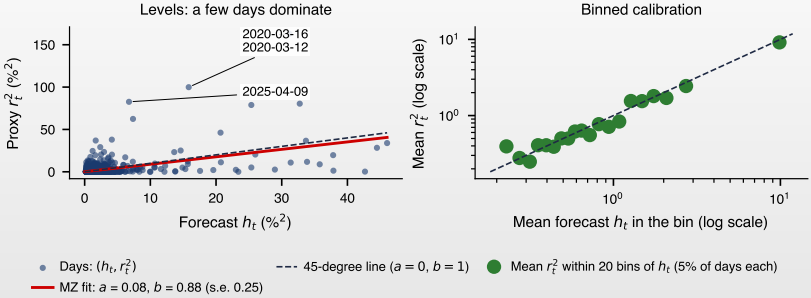
- Mean QLIKE: GJR-t 0.7312, GARCH-t 0.7562; EWMA (DM $t = 3.47$) and Hist-20 ($t = 4.10$) are worse, their gap grows in March 2020; GARCH-family t are descriptive (nested, expanding window: B5 Extended)

MFM_ch5_seminar

B5 [Solved]: Calibration — Task (2/2)

- **Questions**
  - are the GARCH-t forecasts unbiased in levels?
  - which days drive the answer?
- **Inputs**: the 2 693 pairs (h_t, r_t^2) of Task 1
- **Tasks**
  1. Estimate the Mincer–Zarnowitz regression (slide: Diebold–Mariano and Mincer–Zarnowitz) with HAC errors (5 lags) and test $a = 0$, $b = 1$ jointly.
  2. Plot r_t^2 against h_t with the fitted line and the 45-degree line, and a chart of the means over 20 bins of h_t.
  3. Re-estimate without the three largest r_t^2 and compare b.
- **Report**: a, b with s.e., R^2, the Wald p, the chart, b without the three days; interpretation: is a low R^2 a sign of a bad forecast?
- Notebook: section B5

B5 [Solved]: Solution — Calibration



- $a = 0.08$ (s.e. 0.24), $b = 0.88$ (s.e. 0.25), $R^2 = 0.238$, Wald $p = 0.802$: unbiasedness is not rejected; binned means lie close to the 45-degree line
- Without 2020-03-16, 2020-03-12, 2025-04-09: $b = 0.64$ (s.e. 0.09), Wald $p < 0.001$. Interpretation: three crisis days carry the level regression; low R^2 reflects the noisy proxy r_t^2, not a bad model Patton (2011)

MFM_ch5_seminar

Transfer Exercise: Choose B6, B7 or B8

- Choose one exercise from the extension packs and solve it with the code of its model
  - B6: innovation laws for Bitcoin (model: B1, B2)
  - B7: forecast comparison on the BET (model: B5)
  - B8: bootstrap and IGARCH test on the BET (model: B4)
- What to produce
  - The table and the chart the task asks for, the answer to the interpretation question and one limitation
- Assessment criteria
  - Reproducibility: the notebook runs top to bottom and reproduces every reported number
  - Correct inference: the right null law (boundary, nested models), uncertainty reported with every estimate
  - Interpretation: a claim supported by the numbers, with its limitation
- AI use is allowed; note the errors you actually catch, in the format of `AI_ERRORS.md`

Exit Check

Exit check (answer in two lines)

- ⊡ One justified conclusion: which model would you use for the S&P 500 and why?
- ⊡ One unresolved limitation: what can the data not decide (e.g. $p = 1$)?

Extension Packs: What Each One Needs

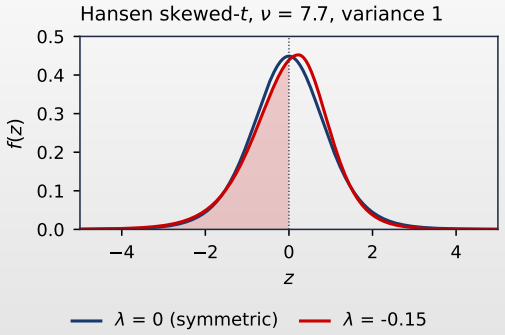
Exercise	Question	Type	Model
A3	GJR persistence when the innovations are skewed	Proposed	A1
A4	is the fourth moment finite?	Proposed	A1, A2
A5	robust versus classical variance of the Gaussian QMLE	Proposed	A2
A6	autocorrelation of squared shocks	Proposed	A1
B6	innovation laws for Bitcoin	Proposed	B1, B2
B5 Extended	Giacomini–White and the Model Confidence Set	Solved	–
B7	forecast comparison on the BET	Proposed	B5
B8	bootstrap and IGARCH test on the BET	Proposed	B4
C2	does US news drive BET volatility? (project)	Proposed	B1, B2
C3	quasi score-driven GARCH on US stocks (project)	Proposed	B1, B2

Skewed Innovations and Moments

- **Hansen skewed-t** Hansen (1994) (A3, B6): a Student-t whose two halves have different widths; mean 0, variance 1
  - $\nu > 2$: tail thickness; $\lambda \in (-1, 1)$: skewness, $\lambda < 0$ a longer left tail, $\lambda = 0$ the Student-t
- Share of the variance from negative innovations

$$m_- = \mathbb{E}\left[z_t^2\, 1(z_t < 0)\right] = \int_{-\infty}^{0} z^2 f(z)\, dz$$

  - chart: $\lambda = -0.15$, $\nu = 7.7$: $m_- = 0.549$ but $P(z_t < 0) = 0.473$; symmetric: $m_- = 0.5$
- **Kurtosis of the innovation** $\kappa_z = \mathbb{E}z_t^4$ (A4, A5): 3 for the Normal, $3(\nu-2)/(\nu-4)$ for Student-t with $\nu > 4$
- **GED** (Generalised Error Distribution) (B6): density $\propto \exp(-|z/c|^{\nu_G}/2)$, c: scale that gives variance 1; $\nu_G = 2$: Normal, $\nu_G < 2$: heavier tails



ARMA Form, PIT and Model Sets

- **AutoRegressive Moving Average (ARMA) form of GARCH(1,1)** (A6): with the surprise $v_t = \varepsilon_t^2 - \sigma_t^2$ (mean 0)

$$\varepsilon_t^2 = \omega + (\alpha + \beta)\,\varepsilon_{t-1}^2 + v_t - \beta\, v_{t-1}$$

  - ε_t^2 is an ARMA(1,1) with autoregressive (AR) coefficient $\alpha + \beta$ and moving-average (MA) coefficient $-\beta$; its ACF (AutoCorrelation Function) decays at rate $\alpha + \beta$
- **PIT** (probability integral transform; C3) Diebold, Gunther & Tay (1998): $u_t = F(\epsilon_t)$, ϵ_t the standardised innovation (z_t elsewhere), F its fitted distribution function
  - if the model is right, u_t is i.i.d. $U(0,1)$: a flat histogram, mean $1/2$, variance $1/12$
- **Model Confidence Set (MCS)** Hansen, Lunde & Nason (2011): the set of models that contains the best one with a chosen probability (90% in B5 Extended)
  - test equal accuracy of all models left; if rejected, drop the worst and repeat; the MCS p-value of a model is the largest p seen until it is dropped
- **Giacomini–White (GW)** Giacomini & White (2006): tests whether the loss difference is predictable from the information available at the forecast origin

A3 [Proposed]: GJR Persistence with Skewed Innovations — Task

- **Question**: what is the persistence of GJR-GARCH when z_t is skewed?
- **Inputs**: S&P 500, GJR-GARCH(1,1) with Hansen skewed-t innovations Hansen (1994): $\hat{\alpha} = 0.00000$, $\hat{\gamma} = 0.21111$, $\hat{\beta} = 0.87970$, $\hat{\nu} = 7.69$, $\hat{\lambda} = -0.150$; obtain the standardised density f of z_t by exponentiating the output of `SkewStudent().loglikelihood` in `arch` (unit variance, `individual=True`); λ: skewness parameter
- **Tasks**
  1. Take expectations in the recursion and show that the persistence is $\alpha + \gamma\, m_- + \beta$, with $m_- = \mathbb{E}[z_t^2 \mathbf{1}(z_t < 0)]$.
  2. Show that $m_- = 1/2$ for a symmetric z_t and explain why m_- differs from $P(z_t < 0)$ in general.
  3. Integrate numerically `quad(f, -inf, 0)` and `quad(lambda z: z**2*f(z), -inf, 0)`; check $\int f = 1$ and $\int z^2 f = 1$ over the real line.
  4. Compute the persistence and the half-life with m_- and with $1/2$.
- **Report**: m_-, $P(z < 0)$, both persistences and half-lives; interpretation: what goes wrong if $\alpha + \gamma/2 + \beta$ is used with skewed innovations?
- Model: A1
- Notebook: section A3

Back Exercise list

A4 [Proposed]: Is the Fourth Moment Finite? — Task

- **Questions**
  - do the fitted S&P 500 models have a finite kurtosis?
  - can the sample kurtosis validate them?
- **Inputs**: GARCH-N: $\hat{\alpha} = 0.120487$, $\hat{\beta} = 0.860062$; GARCH-t: $\hat{\alpha} = 0.1225$, $\hat{\beta} = 0.8723$, $\hat{\nu} = 6.246$; sample kurtosis of the returns 13.7 (total, not excess)
- **Tasks**
  1. With $\kappa_z = \mathbb{E}z_t^4$ (the kurtosis of z_t), show that $\mathbb{E}\varepsilon_t^4 = \kappa_z \mathbb{E}\sigma_t^4$ and $\mathbb{E}\sigma_t^4 < \infty$ iff $\mathbb{E}(\alpha z_t^2 + \beta)^2 = (\alpha+\beta)^2 + (\kappa_z - 1)\alpha^2 < 1$.
  2. For GARCH-N ($\kappa_z = 3$) compute the condition and the total kurtosis $K = 3[1-(\alpha+\beta)^2]/[1-(\alpha+\beta)^2 - 2\alpha^2]$, with the denominator shown; repeat with estimates rounded to 3 decimals.
  3. For GARCH-t compute $\kappa_z = 3(\nu-2)/(\nu-4)$ and the condition.
- **Report**: a worksheet row per model: κ_z, the condition, finite or not, K; interpretation: is agreement with the sample kurtosis a test of the model?
- Model: A1, A2
- Notebook: section A4

A5 [Proposed]: Robust versus Classical Variance, ARCH(1) — Task

- **Question**: by how much does the inverse Hessian misstate the variance of the Gaussian QMLE when z_t is not Normal?
- **Inputs**: ARCH(1), $\sigma_t^2 = \omega + \alpha\varepsilon_{t-1}^2$, ω known; $\ell_t = -\frac{1}{2}(\ln\sigma_t^2 + \varepsilon_t^2/\sigma_t^2)$; $x_t = \varepsilon_{t-1}^2/\sigma_t^2$ is known at $t-1$ and independent of z_t; identities $\partial\sigma_t^2/\partial\alpha = \varepsilon_{t-1}^2$, $\varepsilon_t^2/\sigma_t^2 = z_t^2$
- **Tasks**
  1. Show that the score is $s_t = \frac{1}{2}(z_t^2 - 1)x_t$ and that $B = \mathbb{E}s_t^2 = \frac{1}{4}(\kappa_z - 1)\mathbb{E}x_t^2$.
  2. Show that $\partial^2\ell_t/\partial\alpha^2 = \frac{1}{2}x_t^2(1 - 2z_t^2)$ and $A = \frac{1}{2}\mathbb{E}x_t^2$.
  3. Show that the asymptotic variance of $\sqrt{n}(\hat{\alpha} - \alpha)$ is $A^{-1}BA^{-1} = \frac{\kappa_z - 1}{2}A^{-1}$; the finite-sample variance divides it by n.
  4. Compute the s.e. ratio $\sqrt{(\kappa_z - 1)/2}$ for Student-t with $\nu = 6.25$ and for $\hat{\kappa}_z = 4.73$ (kurtosis of the GARCH-N residuals $\hat{z}_t$ on the S&P 500).
- **Report**: s_t, A, B, the ratio, two numbers; interpretation: what happens to the $\sqrt{n}$ theory when $\kappa_z = \infty$?
- Model: A2
- Notebook: section A5

A6 [Proposed]: Autocorrelation of Squared Shocks — Task

- **Questions**
  - does the ACF of r_t^2 implied by GARCH-N match the data?
  - how precise is the sample ACF?
- **Inputs**: the ARMA(1,1) form of ε_t^2 (slide: ARMA Form, PIT and Model Sets); ARMA(1,1) with AR coefficient ϕ and MA coefficient θ: $\rho_1 = (1+\phi\theta)(\phi+\theta)/(1+2\phi\theta+\theta^2)$, $\rho_k = \phi^{k-1}\rho_1$; S&P 500 GARCH-N: $\hat{\omega} = 0.0253$, $\hat{\alpha} = 0.1205$, $\hat{\beta} = 0.8601$
- **Tasks**
  1. Substitute $\phi = \alpha + \beta$, $\theta = -\beta$ and derive $\rho_1(\varepsilon^2)$.
  2. Compute ρ_1, ρ_{10} and plot the theoretical ACF against the sample ACF of $(r_t - \bar{r})^2$, lags 1–50.
  3. Simulate 200 GARCH-N paths of the sample length (500 burn-in days; Normal innovations from `np.random.default_rng(44)`, initial variance $\omega/(1-\alpha-\beta)$) and add the 95% envelope of their sample ACFs.
- **Report**: the formula, ρ_1, ρ_{10}, the chart; interpretation: is matching the ACF of r_t^2 a strong or a weak check?
- Model: A1
- Notebook: section A6

Back Exercise list

B6 [Proposed]: Innovation Laws for Bitcoin — Task

- **Questions**
  - which innovation law describes Bitcoin?
  - is Gaussian-QMLE inference valid there?
- **Inputs**: Bitcoin closing prices (BTC/USD) from 2014-09-17; returns 2014-09-18 to 2026-09-18, 4 384 returns on calendar days; GARCH(1,1), constant mean
- **Tasks**
  1. Fit Normal and Student-t innovations; compare ℓ and BIC and draw the QQ plots of $\hat{z}_t$ against each fitted law.
  2. Test Normal ($1/\nu = 0$) against Student-t by LR with the boundary law $\frac{1}{2}\chi^2_0 + \frac{1}{2}\chi^2_1$ (5% critical value 2.71).
  3. Optional: add skewed-t (ν, λ) and GED (shape ν_G); test $\lambda = 0$ by LR; report every shape parameter with its s.e. and flag active bounds.
- **Report**: a table of ℓ, BIC, persistence and shape parameters, the QQ chart, the LR tests; interpretation:
  - with $\hat{\nu} < 4$, is Gaussian-QMLE inference valid Hall & Yao (2003)?
  - how does it differ for the Student-t MLE?
- Model: B1, B2
- Notebook: section B6

Back Exercise list

B5 Extended [Solved]: Giacomini–White and the Model Confidence Set — Task

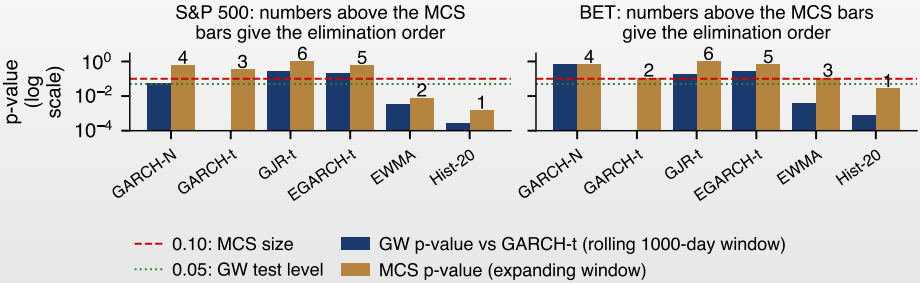
- **Question**: which comparisons of B5 survive once estimated parameters, nested models and six competitors are taken into account?
- **Inputs**: S&P 500 and BET; rolling window of $W = 1\,000$ days refit each January (GW); the expanding-window QLIKE losses of B5 (MCS)
- **Tasks**
  1. For the S&P 500 and each model m, form $d_{t+1} = L^{m}_{t+1} - L^{GARCH\text{-}t}_{t+1}$ and $Z_{t+1} = (1, d_t)'\, d_{t+1}$; compute $GW = n\bar{Z}'\hat{\Omega}^{-1}\bar{Z}$ with $\hat{\Omega} = n^{-1}\sum Z_{t+1}Z'_{t+1}$ (n: number of forecasts, $\bar{Z}$: mean of Z_{t+1}) and compare with χ^2_2 (5.99) Giacomini & White (2006).
     - for one-step forecasts under the conditional null, Z_{t+1} has zero conditional mean, so $\hat{\Omega}$ is the uncentred second moment, without HAC lags
  2. Compute the 90% Model Confidence Set Hansen, Lunde & Nason (2011): statistic T_R (the largest pairwise $|t|$ of loss differences), stationary bootstrap with mean block 10, 10 000 draws, seed 42 (`arch.bootstrap.MCS`).
  3. Report the elimination order and the models kept, for both markets.
- **Report**: GW statistics and p-values, MCS p-values, the chart; interpretation:
  - why is DM not valid for GJR-t versus GARCH-t West (1996), Clark & McCracken (2001)?
  - do DM and MCS agree?
- Notebook: section B5

B5 Extended [Solved]: Solution

Model	*GW* S&P	p	MCS p S&P	MCS p BET
GARCH-N	5.70	0.058	0.588	0.658
GARCH-t	–	–	0.341	0.108
GJR-t	2.51	0.285	1.000	1.000
EGARCH-t	3.19	0.202	0.588	0.658
EWMA	11.31	0.004	0.008	0.108
Hist-20	16.35	< 0.001	0.002	0.026

- S&P 500: GW rejects EWMA and Hist-20 only; elimination order Hist-20, EWMA, GARCH-t, GARCH-N, EGARCH-t, GJR-t; kept: the four GARCH models. BET: order Hist-20, GARCH-t, EWMA, GARCH-N, EGARCH-t, GJR-t; the MCS also keeps EWMA ($p = 0.108$), close to the 0.10 threshold: inclusion depends on the block length and the re-estimation frequency
- On the S&P 500, DM and MCS agree: EWMA and Hist-20 are worse in DM ($t = 3.47$; 4.10) and are the first models the MCS drops; the GARCH models stay in the MCS

B5 Extended [Solved]: GW and MCS in a Chart



- Under H_0 the DM differential (nested models, expanding window) degenerates
- GW (fixed-length rolling window) tests the method, under the conditions of the theorem: finite moments, limited dependence, a non-singular Ω
- The MCS controls the error over the six models under its assumptions; on expanding-window losses, the same caveat as DM

MFM_ch5_seminar

B7 [Proposed]: Forecast Comparison on the BET — Task

- **Question**: do the B5 conclusions transfer to the Bucharest market?
- **Inputs**: BET, 2 675 forecast days, 2016–2026; the same six models, expanding window, proxy and code as B5
- **Tasks**
  1. Compute the mean QLIKE, the DM t against GARCH-t and the cumulative $\sum d_t$.
  2. Estimate the MZ (Mincer–Zarnowitz, regression) regression for GARCH-t with HAC errors; report a, b, their s.e., R^2 and the Wald p, and plot the calibration chart.
  3. Put the S&P 500 and BET results side by side, labelling the estimation window of each number.
- **Report**: two tables and two charts; interpretation:
  - are the BET forecasts calibrated?
  - what would you test next?
- Model: B5
- Notebook: section B7

Back Exercise list

B8 [Proposed]: Bootstrap and IGARCH Test for the BET — Task

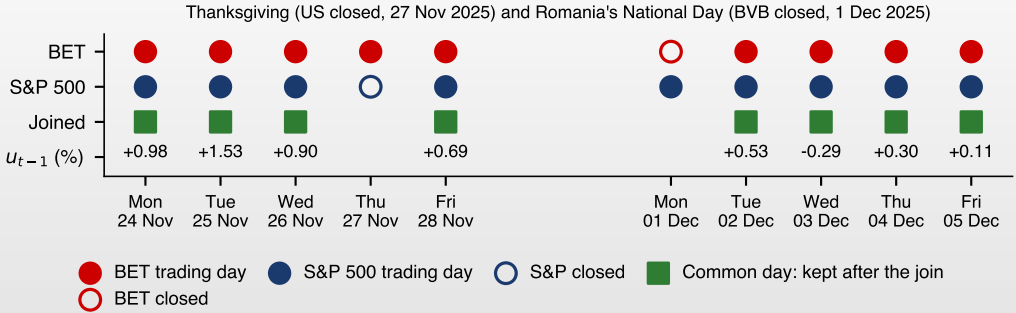
- **Questions**
  - how precise is the BET persistence?
  - can its variance have a unit root?
- **Inputs**: BET GARCH-t, $\hat{p} = 0.9907$ (half-life 74 days); the two B4 algorithms and functions, $B = 999$; paths of the BET sample length after 500 discarded burn-in days, started at the BET sample variance; seeds 42/43
- **Tasks**
  1. Run Algorithm 1 and report the percentile interval for p, the half-life interval and the share of draws at $p = 1$.
  2. Compare the bootstrap s.d. of $\hat{p}^*$ with the robust s.e. of $\hat{\alpha} + \hat{\beta}$.
  3. Run Algorithm 2 and compare the bootstrap, mixture and χ^2_1 p-values at 5%.
- **Report**: the intervals, the two s.e., the test and the two-panel chart; interpretation: do the separately computed BET and S&P 500 intervals show which half-life is longer?
- Model: B4
- Notebook: section B8

C2 [Proposed]: Does US News Drive BET Volatility? — Task (1/2): Data and Timing

- **Questions**
  - does yesterday's S&P 500 shock raise today's BET variance?
  - by how much?
- **Inputs**: BET and S&P 500 closing prices, 2000-01-07 to 2026-09-18; align the *prices* on common days, then compute returns; t indexes the common days
- **Tasks**
  1. Build the sample of common days and $u_{t-\mathbf{1}}$, the S&P 500 return of the previous common day (Wall Street closes after Bucharest).
  2. Check the holiday week below: on 2 December 2025, $u_{t-\mathbf{1}}$ is the S&P 500 return from 26 to 28 November (0.69%); say which US session is missing and where it goes.
- **Report**: the number of common days and the timing check
- Model: B1, B2
- Notebook: section C2

Back Exercise list

C2 [Proposed]: the Common-Day Calendar around a Holiday



- Filled markers: trading days; squares: days kept after aligning the prices; the numbers are u_{t-1} for each common day

MFM_ch5_seminar

C2 [Proposed]: Task (2/2): Model, Test, Contribution

- **Questions**: the same questions, now as an estimated model
- **Inputs**: $r_t = \mu + \phi u_{t-1} + \varepsilon_t$, $\sigma_t^2 = \omega + \alpha\varepsilon_{t-1}^2 + \beta\sigma_{t-1}^2 + \delta u_{t-1}^2$, $\varepsilon_t = \sigma_t z_t$ with standardised Student-t z_t; notebook functions `spill_negll`, `fit_spill`
- **Tasks**
  1. Write the log-likelihood step by step (start σ_1^2 = sample variance) and estimate the unrestricted model and the model with $\delta = 0$.
  2. Test $\delta = 0$ by LR; since $\delta \geq 0$, use $\frac{1}{2}\chi_0^2 + \frac{1}{2}\chi_1^2$; repeat on 2000–2012 and 2013–2026.
  3. Compute the accumulated US term $s_t = \beta s_{t-1} + \delta u_{t-1}^2$ with $s_1 = 0$, the share s_t/σ_t^2 and the volatility without it, $\sqrt{\sigma_t^2 - s_t}$; plot both.
- **Report**: estimates for three samples, the LR tests, the contribution chart; a Quantlet folder (README, script/notebook, data source, figure and metainfo) and one slide with the conclusion and its main limitation; interpretation: is the link causal or predictive?
- **Assessment criteria**: correct inference (boundary, LR), a reproducible Quantlet folder, the main limitation stated

C3 [Proposed]: Quasi Score-Driven (QSD) GARCH on US Stocks — Task (1/2)

- **Question**: is the tail thickness of returns the same thing as the downweighting of large shocks in the variance update Blasques, Francq & Laurent (2023)?
- **Inputs**: start with JPMorgan (adjusted closing prices) and the VIX (closing values), aligned on common days, 3 January 1995 – 28 February 2019; $r_t = 100\,\Delta \ln P_t$, $X_t = \mathrm{VIX}_t^2/252$ (the VIX of day t enters f_{t+1}); model (the estimated recursion): $y_t = r_t - \mu - \phi r_{t-1} = \sqrt{f_t}\,\epsilon_t$, ϵ_t standardised Student-t with $1/\xi$ degrees of freedom; $f_{t+1} = \omega + \varpi X_t + \alpha\,\Psi\big(\frac{1+\zeta}{1-2\zeta+\zeta\epsilon_t^2}\big)\epsilon_t^2 f_t + \beta f_t$, $\Psi(x) = x\tanh(cx/2)$, $c = 1\,000$, $\zeta \in (-1, 1/2)$, f_1 = variance of y_t; ϖ: weight of the VIX term X_t; parameters: ξ = tail thickness, ζ = downweighting; $\zeta = 0$: GARCH(1,1)-t; $\zeta = \xi$: Beta-t-GARCH
- **Tasks**
  1. Estimate by MLE GARCH(1,1)-t ($\zeta = 0$), Beta-t-GARCH ($\zeta = \xi$) and QSD (ζ free) with `qsd_negll` and `fit_qsd` (notebook of `MFM_ch5_quasi_score_driven`).
  2. Test $\zeta = 0$ and $\xi = \zeta$ by LR against χ^2_1 ($\zeta = 0$ is interior thanks to Ψ).
- **Report**: estimates, log-likelihoods, the two LR tests
- Model: B1, B2

Back Exercise list

C3 [Proposed]: Task (2/2): Diagnostics and Extension

- **Questions**
  - do the three models fit?
  - does the answer hold beyond one stock?
- **Inputs**: the fits of Task 1; the PIT $u_t = F(\epsilon_t)$ (slide: ARMA Form, PIT and Model Sets); T: sample length
- **Tasks**
  1. Plot $\sqrt{f_t}$ of the three models around the largest shock and the PIT histograms.
  2. Compute the naive statistic $(\sum_t u_t - T/2)/\sqrt{T/12}$ and explain why it ignores estimation error (the correction of Pierce (1982) lowers the variance below $T/12$).
  3. Optional: extend to the 11 course stocks and to 18 September 2026; report rejection frequencies with a binomial interval that assumes, hypothetically, independent rejections.
- **Report**: a Quantlet folder (README, script/notebook, data source, figure and metainfo) and one slide with the table, the tests and the main limitation; interpretation: why can 11 selected survivors not confirm the paper's 89.5% (share of stocks with $\zeta = 0$ rejected)?
- **Assessment criteria**: correct inference (boundary, LR), a reproducible Quantlet folder, the main limitation stated

Acronyms used in this material (1/2)

- **ACF**: AutoCorrelation Function
- **AI**: Artificial Intelligence
- **AIC**: Akaike Information Criterion
- **AR**: AutoRegressive (model)
- **ARCH**: AutoRegressive Conditional Heteroskedasticity
- **ARMA**: AutoRegressive Moving Average (model)
- **BET**: Bucharest Exchange Trading (index)
- **BIC**: Bayesian Information Criterion
- **BVB**: Bursa de Valori București — *the Bucharest Stock Exchange*
- **DM**: Diebold–Mariano (test)
- **EGARCH**: Exponential GARCH
- **EWMA**: Exponentially Weighted Moving Average
- **GARCH**: Generalised AutoRegressive Conditional Heteroskedasticity
- **GARCH-N**: GARCH with Normal innovations
- **GED**: Generalised Error Distribution
- **GJR**: Glosten–Jagannathan–Runkle (asymmetric GARCH)

Acronyms used in this material (2/2)

- **GJR-GARCH**: Glosten–Jagannathan–Runkle GARCH
- **GW**: Giacomini–White (test)
- **HAC**: Heteroskedasticity and Autocorrelation Consistent (standard errors)
- **IGARCH**: Integrated GARCH
- **LM**: Lagrange Multiplier
- **LR**: Likelihood Ratio (test)
- **MA**: Moving Average
- **MCS**: Model Confidence Set
- **MLE**: Maximum Likelihood Estimation
- **MZ**: Mincer–Zarnowitz (regression)
- **NIC**: News Impact Curve
- **PIT**: Probability Integral Transform
- **QLIKE**: Quasi-LIKElihood (loss function)
- **QMLE**: Quasi-Maximum Likelihood Estimator
- **QQ**: Quantile–Quantile (plot)
- **QSD**: Quasi Score-Driven (model)
- **US**: United States
- **VIX**: CBOE Volatility IndeX

References (1/3)

- Andrews, D.W.K. (2000). Inconsistency of the bootstrap when a parameter is on the boundary of the parameter space. *Econometrica*, 68(2), 399–405.
- Andrews, D.W.K. (2001). Testing when a parameter is on the boundary of the maintained hypothesis. *Econometrica*, 69(3), 683–734.
- Blasques, F., Francq, C., Laurent, S. (2023). Quasi score-driven models. *Journal of Econometrics*, 234(1), 251–275.
- Bollerslev, T. (1986). Generalized autoregressive conditional heteroskedasticity. *Journal of Econometrics*, 31(3), 307–327.
- Bollerslev, T., Wooldridge, J.M. (1992). Quasi-maximum likelihood estimation and inference in dynamic models with time-varying covariances. *Econometric Reviews*, 11(2), 143–172.
- Clark, T.E., McCracken, M.W. (2001). Tests of equal forecast accuracy and encompassing for nested models. *Journal of Econometrics*, 105(1), 85–110.
- Diebold, F.X., Gunther, T.A., Tay, A.S. (1998). Evaluating density forecasts with applications to financial risk management. *International Economic Review*, 39(4), 863–883.

References (2/3)

- Diebold, F.X., Mariano, R.S. (1995). Comparing predictive accuracy. *Journal of Business & Economic Statistics*, 13(3), 253–263.
- Engle, R.F., Ng, V.K. (1993). Measuring and testing the impact of news on volatility. *Journal of Finance*, 48(5), 1749–1778.
- Franke, J., Härdle, W.K., Hafner, C.M. (2019). Statistics of Financial Markets: An Introduction, 5th ed. Springer.
- Giacomini, R., White, H. (2006). Tests of conditional predictive ability. *Econometrica*, 74(6), 1545–1578.
- Glosten, L.R., Jagannathan, R., Runkle, D.E. (1993). On the relation between the expected value and the volatility of the nominal excess return on stocks. *Journal of Finance*, 48(5), 1779–1801.
- Hall, P., Yao, Q. (2003). Inference in ARCH and GARCH models with heavy-tailed errors. *Econometrica*, 71(1), 285–317.
- Hansen, B.E. (1994). Autoregressive conditional density estimation. *International Economic Review*, 35(3), 705–730.

References (3/3)

- Hansen, P.R., Lunde, A., Nason, J.M. (2011). The model confidence set. *Econometrica*, 79(2), 453–497.
- J.P. Morgan/Reuters (1996). RiskMetrics — Technical Document, 4th ed. New York.
- Li, W.K., Mak, T.K. (1994). On the squared residual autocorrelations in non-linear time series with conditional heteroskedasticity. *Journal of Time Series Analysis*, 15(6), 627–636.
- Mincer, J., Zarnowitz, V. (1969). The evaluation of economic forecasts. In J. Mincer (ed.), *Economic Forecasts and Expectations: Analysis of Forecasting Behavior and Performance*, NBER.
- Patton, A.J. (2011). Volatility forecast comparison using imperfect volatility proxies. *Journal of Econometrics*, 160(1), 246–256.
- Pierce, D.A. (1982). The asymptotic effect of substituting estimators for parameters in certain types of statistics. *The Annals of Statistics*, 10(2), 475–478.
- West, K.D. (1996). Asymptotic inference about predictive ability. *Econometrica*, 64(5), 1067–1084.