Master in Financial Engineering (EPFL)

Financial Econometrics

Exercises: GARCH models

Exercise 1: Let $(X_t) \sim \texttt{GARCH}(p,q)$ such that

$$\begin{aligned} X_t &= \sigma_t Z_t \\ Z_t &\sim \text{ i.i.d. } \mathcal{N}(0,1) \\ \sigma_t^2 &= \omega + \sum_{i=1}^{p} \alpha_i X_{t-i}^2 + \sum_{j=1}^{q} \beta_j \sigma_{t-j}^2 \end{aligned}$$

where the usual conditions on the parameters hold true.

1. Let η_t denote $\eta_t = \sigma_t^2(Z_t^2 - 1)$. Show that $\mathbb{E}\left[\eta_t \mid I_{t-1}\right] = 0$ where I_{t-1} is the information set at time $t-1$.

2. Using the previous result and the representation of a $\texttt{GARCH}(p,q)$, show that the stochastic process (X_t^2) has the following representation:

$$X_t^2 = \omega + \sum_{i=1}^{r} (\alpha_i + \beta_i) X_{t-i}^2 + \eta_t - \sum_{j=1}^{q} \beta_j \eta_{t-j}.$$

   where $r = \max\{p,q\}$, $\alpha_i = 0$ if $i > p$, and $\beta_j = 0$ if $j > q$.

3. Interpret the previous result.

Exercise 2: Let $(X_t) \sim \texttt{GARCH}(1,1)$ such that

$$\begin{aligned} X_t &= \sigma_t Z_t \\ \sigma_t^2 &= \omega + \alpha X_{t-1}^2 + \beta \sigma_{t-1}^2 \end{aligned}$$

where $Z_t \sim \texttt{i.i.d.}\mathcal{N}(0,1)$, $\omega > 0$, $\alpha \geq 0$ and $\beta \geq 0$ (with $\alpha + \beta < 1$). We further assume that $1 - 3\alpha^2 - \beta^2 - 2\alpha\beta > 0$.

1. Find the autocovariance and autocorrelation function of (X_t^2).

Exercise 3: Daily log returns of the S&P500 index (during the period 1993-2013), denoted $r_1,\cdots,r_T$ with $T = 5248$, are assumed to follow an ARCH(1) specification:

$$\begin{aligned} r_t &= \sigma_t z_t \quad t \geq 2 \\ \sigma_t^2 &= 0.0001 + \alpha_1 r_{t-1}^2 \end{aligned}$$

where $z_t \sim \texttt{i.i.d.}\mathcal{N}(0,1)$. The maximum likelihood estimate of α_1 is 0.3577. In addition,

$$\sum_{t=2}^{T} \frac{r_{t-1}^4}{2\sigma_t^4} = 1757.807$$

$$\sum_{t=2}^{T} \frac{r_t^2 r_{t-1}^4}{2\sigma_t^6} = 3363.839$$

1. How would you test the presence of an ARCH(1) effect?
2. Using the available information, estimate the asymptotic variance associated with the ML estimate of α_1 and provide a test statistic. What is your conclusion?
3. The last observation available in the dataset is $r_T = -0.00066$. Using Appendix 1, provide the following VaR forecasts at date $T+1$:
   - $\texttt{VaR}_{90\%}^{T+1}$
   - $\texttt{VaR}_{95\%}^{T+1}$
   - $\texttt{VaR}_{99\%}^{T+1}$

Exercise 4: Daily log returns of the S&P500 index (during the period 1993-2013), denoted $r_1,\cdots,r_T$ with $T = 5248$, are assumed to follow a GARCH(1,1) specification:

$$\begin{aligned} r_t &= \sigma_t z_t \quad t \geq 2 \\ \sigma_t^2 &= 0.0000013 + 0.08184 r_{t-1}^2 + 0.90897 \sigma_{t-1}^2 \end{aligned}$$

where $z_t \sim \texttt{i.i.d.}\mathcal{N}(0,1)$. The last observation available in the dataset is $r_T = -0.00066$ and $\sigma_T^2 = 0.000039$.

1. Using Appendix 1, provide the following GARCH(1,1)-based VaR forecasts at date $T+1$:
   - $\texttt{VaR}_{90\%}^{T+1}$
   - $\texttt{VaR}_{95\%}^{T+1}$
   - $\texttt{VaR}_{99\%}^{T+1}$
2. Taking that the last 20 day log-returns available in the dataset are (in increasing order):

   $$\begin{aligned} &-0.01272, -0.00351, -0.00317, -0.00313, -0.00218, -0.0023, -0.00100, -0.00066 \\ &0.00017, 0.00028, 0.00239, 0.00243, 0.00351, 0.00434, 0.00498, 0.00499, 0.00509 \\ &0.00799, 0.00804, 0.01340 \end{aligned}$$

   Provide the 90% value-at-risk forecast at date $T+1$ using the 20-day historical simulation method.

Exercise 5: Let $(\epsilon_t) \sim \texttt{GARCH}(1,1)$ such that

$$\begin{aligned} \epsilon_t &= \sigma_t Z_t \\ \sigma_t^2 &= \omega + \alpha\epsilon_{t-1}^2 + \beta\sigma_{t-1}^2 \end{aligned}$$

where $Z_t \sim \texttt{i.i.d.}\mathcal{N}(0,1)$, $\omega > 0$, $\alpha \geq 0$ and $\beta \geq 0$ (with $\alpha + \beta < 1$).

1. Show that the h-step ahead forecast of the conditional variance (at time t), denoted $\sigma_t^2(h)$, is given by (see also Case Study 2):

   $$\sigma_t^2(h) = \frac{\omega\left(1-(\alpha+\beta)^{h-1}\right)}{1-(\alpha+\beta)} + (\alpha+\beta)^{h-1}\,\sigma_t^2(1).$$

   where $\sigma_t^2(1)$ is the one-step ahead forecast of the conditional variance.

2. What happens as $h \to \infty$? Interpret (see also Case Study 2).

3. Suppose now that one considers an $\texttt{AR}(1)$ process with $\texttt{GARCH}(1,1)$ error terms:

   $$\begin{aligned} X_t &= \phi X_{t-1} + \epsilon_t \\ \epsilon_t &= \sigma_t Z_t \\ \sigma_t^2 &= \omega + \alpha\epsilon_{t-1}^2 + \beta\sigma_{t-1}^2 \end{aligned}$$

   where $|\phi| < 1$.

   3.1. Find $\mathbb{V}\left[X_{t+h} \mid X_u, u < t\right]$ if $\phi^2 \neq \alpha + \beta$ (respectively, $\phi^2 = \alpha + \beta$)).

   3.2. Interpret the previous result.

   3.3. What would happen if $\mid \phi \mid = 1$ and the process is initialized at 0 (i.e., all the variables at negative dates are initialized at zero)? Explain.

Exercise 6: Standard GARCH models have the symmetry property. Indeed, if the law of η_t is symmetric (and as long as the $\texttt{GARCH}(p,q)$ is second-order stationary), there is absence of covariance between σ_t and ϵ_{t-h}, $\mathbb{C}ov(\sigma_t, \epsilon_{t-h}) = 0$ for $h > 0$. Furthermore, let ϵ_t^+ and ϵ_t^- denote the positive and negative components of ϵ_t (for all t), respectively:

$$\epsilon_t^+ = \max\{\epsilon_t, 0\} \quad \text{and} \quad \epsilon_t^- = \min\{\epsilon_t, 0\}.$$

It can be shown that the property, $\mathbb{C}ov(\sigma_t, \epsilon_{t-h}) = 0$ for $h > 0$, is true if and only if $\mathbb{C}ov(\epsilon_t^+, \epsilon_{t-h}) = \mathbb{C}ov(\epsilon_t^-, \epsilon_{t-h}) = 0$ for $h > 0$. However, this characterization of the symmetry property using the autocovariance function is in general rejected on financial series. In this respect, asymmetric GARCH models have been proposed in the literature. Based on (weekly) CAC40 index returns, one estimates the following symmetric and asymmetric GARCH models (see Lecture notes):

- $\texttt{GARCH}(1,1)$ model:

$$\begin{aligned} r_t &= \mu + \epsilon_t \\ \epsilon_t &= \sigma_t z_t, z_t \sim \texttt{i.i.d}\mathcal{N}(0,1) \\ \sigma_t^2 &= \omega + \alpha\epsilon_{t-1}^2 + \beta\sigma_{t-1}^2 \end{aligned}$$

- $\texttt{EGARCH}(1,1)$ model:

$$\begin{aligned} r_t &= \mu + \epsilon_t \\ \epsilon_t &= \sigma_t z_t, z_t \sim \texttt{i.i.d}\mathcal{N}(0,1) \\ \log(\sigma_t^2) &= \omega + \alpha\left(\theta z_t + \gamma\left[|z_t| - \sqrt{\frac{2}{\pi}}\right]\right) + \beta\log(\sigma_{t-1}^2) \end{aligned}$$

- $\texttt{QGARCH}(1,1)$ model:

$$\begin{aligned} r_t &= \mu + \epsilon_t \\ \epsilon_t &= \sigma_t z_t, z_t \sim \texttt{i.i.d}\mathcal{N}(0,1) \\ \sigma_t^2 &= \omega + \alpha\epsilon_{t-1}^2 + \phi\epsilon_{t-1} + \beta\sigma_{t-1}^2 \end{aligned}$$

- $\texttt{GJR-GARCH}(1,1)$ model:

$$\begin{aligned} r_t &= \mu + \epsilon_t \\ \epsilon_t &= \sigma_t z_t, z_t \sim \texttt{i.i.d}\mathcal{N}(0,1) \\ \sigma_t^2 &= \omega + \left[\alpha\epsilon_{t-1}^2 + \gamma\Pi_{t-1}^{-}\epsilon_{t-1}^2\right] + \beta\sigma_{t-1}^2 \end{aligned}$$

  where Π_t^- equals one if $\epsilon_t < 0$, and 0 otherwise.

- $\texttt{TGARCH}(1,1)$ model:

$$\begin{aligned} r_t &= \mu + \epsilon_t \\ \epsilon_t &= \sigma_t z_t, z_t \sim \texttt{i.i.d}\mathcal{N}(0,1) \\ \sigma_t &= \omega + \left[\alpha_1\epsilon_{t-1}^{-} + \alpha_2\epsilon_{t-1}^{+}\right] + \beta\sigma_{t-1} \end{aligned}$$

  Note that the last equation of a TGARCH model can also be written as:

$$\sigma_t = \omega + \left[\gamma_1|\epsilon_{t-1}| + \gamma_2\Pi_{t-1}^{-}|\epsilon_{t-1}|\right] + \beta_1\sigma_{t-1}$$

  where Π_t^- equals one of $\epsilon_t < 0$, and 0 otherwise.

On obtains the following estimations. All coefficients are statistically different from zero.

Table 1: Symmetric and asymmetric GARCH models

	GARCH	EGARCH	QGARCH	GJR-GARCH	TGARCH
μ	5×10^{-4}	4×10^{-4}	3×10^{-4}	4×10^{-4}	4×10^{-4}
ω	8×10^{-6}	-0.640	9×10^{-6}	1×10^{-5}	8×10^{-4}
α	0.090	0.149	0.070	0.130	
β	0.840	0.850	0.850	0.840	0.870
θ		-0.530			
γ		1		-0.100	
ϕ			-8×10^{-4}		
α_1					-0.120
α_2					0.030

1. Comment the different specifications.
2. Provide an interpretation of the estimated coefficients.
3. The value of the log-likelihood function is given in Table 2. Which model would you suggest to use? Explain carefully.

Table 2: Log-likelihhod values of symmetric and asymmetric GARCH models

	GARCH	EGARCH	QGARCH	GJR-GARCH	TGARCH
ℓ	6392	6406	6406	6410	6408

Appendix 1: Value-at-Risk

Consider a portfolio of n assets with fixed allocation $a = (a_1, \cdots, a_n)^\top$ between t and $t+h$. At date t, the investor's endowment: $W_t(a) = a^\top p_t$ is allocated between portfolio choice (at time $t+h$) and a reserve, R_t. Especially, R_t is chosen such that the global position (portfolio value + reserve) corresponds to a loss with a predetermined (small) probability at date $t+h$:

$$\mathbb{P}_{\mathbf{t}}\left[W_{t+h}(a) + R_t < \mathbf{0}\right] = \alpha$$

$\Rightarrow -R_t$ is the α-quantile of the conditional distribution of future portfolio value (P&L distribution). In this respect, the required capital is (theoretically) the Value-at-Risk (VaR) denoted by

$$\texttt{VaR}_t = W_t(a) + R_t$$

and is characterized by the condition

$$\mathbb{P}_t\left[W_{t+h}(a) - W_t(a) + \texttt{VaR}_t < 0\right] = \alpha.$$

Therefore, the `VaR` depends on the information available at time t, the horizon h, the portfolio (set of assets and allocation), and the loss probability α:

$$\texttt{VaR}_t = \texttt{VaR}(I_t, a, h, \alpha)$$

and has (at least) two objectives: (1) To measure the risk and (2) To determine the capital reserve.

Example 1: Interpretation of `VaR`

- Suppose the daily `VaR` of a trading portfolio is 5'000'000 at the 99% confidence level.

- There is 1 chance in 100 that a loss > 5m (USD) will occur the next day (under normal market conditions).

- The `VaR` is the 1%-quantile of the probability distribution of the position. □

Example 2: Gaussian (conditional) `VaR`

- Suppose $h = 1$ and $\Delta p_t \sim \mathcal{N}(\mu_t, \Omega_t)$ where $\mu_t = \mathbb{E}_t[\Delta p_t]$ and $\Omega_t = \mathbb{V}_t[\Delta p_t]$.

- The (conditional) $\texttt{VaR}_t$ satisfies (for a loss probability α):

  $$\mathbb{P}_t\left[a^\top \Delta p_t < -\texttt{VaR}_t\right] = \alpha$$

  or, equivalently,

  $$\mathbb{P}\left[Z_t < \frac{-\texttt{VaR}_t - a^\top \mu_t}{(a^\top \Omega_t a)^{1/2}}\right] = \alpha.$$

  where $Z_t \sim \mathcal{N}(0, 1)$.

- Therefore

$$
\begin{aligned}
\texttt{VaR}_t &= -a^\top \mu_t - \left(a^\top \Omega_t a\right)^{1/2} z_\alpha \\
&= -a^\top \mu_t + \left(a^\top \Omega_t a\right)^{1/2} z_{1-\alpha}
\end{aligned}
$$

where $z_{1-\alpha}$ is the $(1-\alpha)$-quantile of a standard normal distribution. □

Different methods have been proposed to estimate the conditional VaR. We focus here on parametric methods to provide value-at-risk forecasts at date $T+1$. Suppose that the information set is given by $I_t = \{x_t, x_{t-1}, \cdots\} \overset{\text{notation}}{=} \underline{x_{t-1}}$ and

$$
x_t = m(\underline{x_{t-1}}, \theta) + \sigma(\underline{x_{t-1}}, \theta) v_t
$$

where $\mathbb{E}_{t-1}[x_t] = m(\underline{x_{t-1}}, \theta)$ and $\mathbb{V}_{t-1}[x_t] = \sigma^2(\underline{x_{t-1}}, \theta)$ depend on parameter θ, and v_t denote some i.i.d. error terms. The vector of parameters, denoted θ, can be estimated by maximum likelihood:

$$
\hat{\theta}_T = \underset{\theta}{\texttt{Argmax}} - \frac{1}{2} \sum_{t=1}^{T} \left\{ \ln \sigma^2(\underline{x_{t-1}}, \theta) + \frac{\left(x_t - m(\underline{x_{t-1}}, \theta)\right)^2}{\sigma^2(\underline{x_{t-1}}, \theta)} \right\}.
$$

Taking the estimate of θ, one can find approximation (estimates) of the conditional drift, $\hat{m}_T = m(., \hat{\theta}_T)$, the conditional variance $\hat{\sigma}^2_T = \sigma^2(., \widehat{\theta}_T)$ and the standardized residuals, $\widehat{v}_t$. Especially,

- At horizon 1, the (theoretical) conditional VaR is defined by

$$
\mathbb{P}_t\left[x_{t+1} < \texttt{VaR}_t(a, \alpha, 1)\right] = \alpha
$$

or, equivalently,

$$
\mathbb{P}_t\left[v_t < -\frac{\texttt{VaR}(a, \alpha, 1) - m(\underline{x_{t-1}}, \theta)}{\sigma(\underline{x_{t-1}}, \theta)}\right] = \alpha.
$$

An estimate of the conditional VaR is then

$$
\widehat{\texttt{VaR}}_t(a, \alpha, 1) = -\widehat{m}_T + \widehat{\sigma}_T F^{-1}(1-\alpha)
$$

where F is the cdf of (i.i.d.) error terms v_t.

Example 3: Suppose that the log returns follow an ARCH(1) process (with Gaussian error terms)

$$
\begin{aligned}
r_t &= \sigma_t \epsilon_t \\
\sigma_t^2 &= \omega + \alpha_1 r_{t-1}^2
\end{aligned}
$$

and $x_t = a r_t \equiv a \Delta p_t$ with

$$
x_t \sim \mathcal{N}\left(0, a^2\left(\omega + \alpha_1 r_{t-1}^2\right)\right).
$$

The conditional VaR is minus the α-quantile of a $\mathcal{N}\left(0, a^2\left(\omega + \alpha_1 r_{t-1}^2\right)\right)$ or its $(1-\alpha)$-quantile:

$$\begin{aligned} \mathtt{VaR}(a, \alpha, 1) &= -a\sqrt{\omega + \alpha_1 r_{t-1}^2}\, z_\alpha \\ &= a\sqrt{\omega + \alpha_1 r_{t-1}^2}\, z_{1-\alpha} \end{aligned}$$

where z_α is the α-quantile, $F^{-1}(\alpha)$, of a $\mathcal{N}(0,1)$ (e.g., $z_{0.1} = -1.282$, $z_{0.05} = -1.645$, or $z_{0.01} = -2.326$).