AMATH546/ECON589 HW4

Eric Zivot
University of Washington, Seattle

Due Thursday, May 16th, 2013

1 Reading

- FRF chapter 3.
- QRM chapter 3, sections 1-3.
- FMUND chapter 6.

2 Data and Programs

Download daily adjusted closing prices on Microsoft and the S&P 500 over the period 2000-01-03 to 2012-04-10 and compute the continuously compound returns. You will find my R scripts on the class webpage helpful for this assignment.

3 Univariate Filtered Historical Simulation

An interesting semi-parametric way to compute VaR and ES combines historical simulation with GARCH estimates of volatility. This approach is called *filtered historical simulation*. Here, filtering refers to using the estimated GARCH volatility to create standardized innovations from which we compute empirical quantiles for VaR and ES. The following exercises explain how the method works.

1. Fit a normal-GARCH(1,1) model

$$\begin{aligned} r_t &= \mu + \varepsilon_t,\ \varepsilon_t = \sigma_t z_t \\ \sigma_t^2 &= \omega + \alpha\varepsilon_{t-1}^2 + \beta\sigma_{t-1}^2 \end{aligned}$$

   to the full sample of daily cc returns on MSFT. Extract the estimated standardized residuals $\hat{z}_t = \hat{\varepsilon}_t/\hat{\sigma}_t$. Show a normal-QQ plot for $\hat{z}_t$. Does it look normally distributed?

2. Compute the 1-day ahead 99% VaR based on the normal GARCH(1,1) model. Recall,

$$\widehat{VaR}^{GARCH}_{T+1|T,.99} = \hat{\mu} + \hat{\sigma}_{T+1|T} \times q^{Z}_{.01} \quad (1)$$

   where $\hat{\sigma}_{T+1|T}$ is the 1-day ahead GARCH(1,1) predicted conditional volatility

3. Compute the unconditional 1-day ahead 99% VaR based on historical simulation. Verify, that you can also compute HS using the formula

$$\widehat{VaR}^{HS}_{.99} = \hat{\mu} + \hat{\sigma} \times \hat{q}^{\hat{Z}}_{.01}$$

   where $\hat{\mu}$ and $\hat{\sigma}$ are the sample mean and standard deviation and $\hat{q}^{\hat{Z}}_{.01}$ is the 1% quantile of the standardized returns $\hat{z}_t = (r - \hat{\mu})/\hat{\sigma}$.

4. Compute the conditional 1-day ahead 99% VaR based on filtered historical simulation using the equation

$$\widehat{VaR}^{FHS}_{T+1|T,.99} = \hat{\mu} + \hat{\sigma}_{T+1|T} \times \hat{q}^{\hat{Z}}_{.01}$$

   where $\hat{\sigma}_{T+1|T}$ is the 1-day ahead GARCH(1,1) predicted conditional volatility and $\hat{q}^{\hat{Z}}_{.01}$ is the 1% quantile of the GARCH(1,1) standardized residuals $\hat{z}_t = \hat{\varepsilon}_t/\hat{\sigma}_t$. This differs from the normal-GARCH(1,1) VaR because we use the empirical quantile of $\hat{z}_t$ and not the normal quantile.

4 Rolling Covariances and Correlations

1. Using a 20-day moving window, compute and plot rolling covariances and correlations. Briefly comment on what you see. Hint: Use the function `rollapply()` from the zoo package. Note, you will have to coerce your data to a "`zoo`" object in order for `rollapply` to work correctly.

5 EWMA Covariances and Correlations

1. Let $\mathbf{r}_t = (r_{MSFT,t}, r_{GSPC,t})'$. Use the `covEWMA()` function on the class webpage to compute the EWMA covariance matrix estimates

$$\begin{aligned} \hat{\boldsymbol{\Sigma}}_t^{EWMA} &= (1-\lambda)\boldsymbol{\varepsilon}_t\boldsymbol{\varepsilon}_t' + \lambda\hat{\boldsymbol{\Sigma}}_{t-1}^{EWMA} \\ \boldsymbol{\varepsilon}_t &= \mathbf{r}_t - \hat{\boldsymbol{\mu}} \end{aligned}$$

   using $\lambda = 0.94$ (the magic number from RiskMetrics). The `covEWMA()` function will return a $T \times 2 \times 2$ array of covariance matrices at each date.

2. Extract the conditional covariance estimates $\hat{\sigma}_{12,t}^{EWMA}$ and the conditional correlation estimates $\rho_{12,t}^{EWMA} = \hat{\sigma}_{12,t}^{EWMA}/\hat{\sigma}_{1,t}^{EWMA} \times \hat{\sigma}_{2,t}^{EWMA}$ and plot them. Compare your results to the rolling estimates. Hint: To compute the conditional correlations, use the functions `lapply()` and `cov2cor()`.

6 DCC Covariances and Correlations

1. Let $\mathbf{r}_t = (r_{MSFT,t}, r_{GSPC,t})'$. Using the `ddcfit()` function from the rmgarch package, Estimate the normal-DCC(1,1) model

$$\begin{aligned} \mathbf{r}_t &= \boldsymbol{\mu} + \boldsymbol{\varepsilon}_t,\ \boldsymbol{\varepsilon}_t = \boldsymbol{\Sigma}_t^{1/2}\mathbf{z}_t,\ \mathbf{z}_t \sim N(\mathbf{0}, \mathbf{I}_2) \\ \boldsymbol{\Sigma}_t &= \mathbf{D}_t\mathbf{R}_t\mathbf{D}_t \\ D_t &= \begin{pmatrix} \sigma_{1t} & 0 \\ 0 & \sigma_{2t} \end{pmatrix},\ R_t = \begin{pmatrix} 1 & \rho_{12,t} \\ \rho_{12,t} & 1 \end{pmatrix} \end{aligned}$$

   where σ_{1t} and σ_{2t} are univariate GARCH(1,1) volatilities and $\rho_{12,t}$ is the GARCH(1,1) conditional correlation between the univariate GARCH(1,1) standardized residuals. Briefly comment on the estimated coefficients and the fit of the model.

2. Plot the estimated in-sample conditional covariances and correlations. Compare the EWMA and rolling estimates.

7 Forecasting Correlations

1. Using the Estimated DCC(1,1) model, compute (using `dccforecast()` function) and plot the first 100 h-step ahead forecasts of conditional covariance and correlation.

2. What are the h-step ahead forecasts of conditional covariance and correlation from the EWMA model?