AMATH546/ECON589 HW3

Eric Zivot
University of Washington, Seattle

Due Thursday, April 25th, 2013

1 Reading

- FRF chapter 2; chapter 5 sections 5-6.
- QRM chapter 2, section 3.
- SDAFE chapter 18 and chapter 19
- FMUND, chapter 3 sections 1-6.

1. Zivot, E. (2008). "Practical Issues in the Analysis of Univariate GARCH Models," *Handbook of Financial Time Series* (available on class webpage).

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. (e.g., use `getSymbols()` from the quantmod package and `Return.calculate()` from the PerformanceAnalytics package). You will find my R scripts on the class webpage helpful for this assignment.

3 MLE of Bivariate Return Distributions and Monte Carlo Estimates of VaR and ES

In this exercise, you will find the maximum likelihood estimates (mles) of the parameters of two bivariate distributions for asset returns. You will also use Monte Carlo simulation to find VaR and ES for a portfolio of Microsoft and the S&P 500. For simulating multivariate normal and multivariate t distributions, you will use the mvtnorm package.

1. Fit a bivariate normal distribution to the returns on Micrsoft and the S&P 500 using the sample moments. Simulate n observations from this fitted distribution, where n is the sample size of your data, and create a scatterplot showing the observed data together with the simulated normal data. How well does the normal data capture the scatter of the observed data?

2. Fit a bivariate Student's distribution to the returns on Micrsoft and the S&P 500 using `cov.trob()` function from the MASS library (profile MLE). Compare the estimates of the mean and correlation matrix to those from the bivariate normal distribution. Simulate n observations from this fitted distribution, where n is the sample size of your data, and create a scatterplot showing the observed data together with the simulated Student's t data. How well does the Student's t data capture the scatter of the observed data compared to the bivariate normal?

3. Now consider a portfolio of 40% Microsoft and 60% S&P 500 Using the fitted bivariate Student's t distribution found above, simulate 10,000 observations and compute nonparametric estimates of 95% and 99% VaR and ES as well as nonparametric estimates of the asset contributions (marginal, component and percent) to portfolio ES.

4 Analytics for the GARCH(1,1) Process

Consider the GARCH(1,1) process

$$\begin{aligned} r_t - \mu &= \varepsilon_t = \sigma_t z_t, \ z_t \sim iid \ N(0,1) \\ \sigma_t^2 &= a_0 + a_1 \varepsilon_{t-1}^2 + b_1 \sigma_{t-1}^2 \end{aligned}$$

Derive the following results:

1. $E[\varepsilon_t] = 0, E[\varepsilon_t^2] = E[\sigma_t^2] = a_0/(1 - a_1 - b_1) = \bar{\sigma}^2$
2. $E[\varepsilon_t | I_{t-1}] = 0,\ E[\varepsilon_t^2 | I_{t-1}] = \sigma_t^2$
3. ε_t^2 has an ARMA(1,1) representation of the form
$$\varepsilon_t^2 = a_0 + (a_1 + b_1)\varepsilon_{t-1}^2 + v_t - b_1 v_{t-1}$$
where $v_t = \varepsilon_t^2 - \sigma_t^2$ is a MDS.
4. σ_t^2 has an ARCH(∞) representation with $a_i = a_1 b_1^{i-1}$
5. Write out the GARCH(1,1) log-likelihood function $\ln L(\mu, a_0, a_1, b_1)$based on a sample $\{r_1, \ldots, r_T\}$ of size T. Specify initial values for ε_0^2 and σ_0^2 and the parameters a_0, a_1 and b_1.
6. σ_{T+k}^2 has the forecasting equation
$$E_T[\sigma_{T+k}^2] - \bar{\sigma}^2 = (a_1 + b_1)^{k-1}(E[\sigma_{T+1}^2] - \bar{\sigma}^2).$$

5 Estimating GARCH Models and Computing Volatility Forecasts

In this exercise you will analyze univariate ARCH and GARCH models for the Microsoft and the S&P 500 returns. Univariate GARCH estimation in R is most easily performed using the functions in the rugarch package. Be sure to read the package documentation (in the doc directory of the package)

1. In this exercise, you will test for ARCH effects in the daily returns on Microsoft and the S&P 500. For each series, compute Engle's LM test for ARCH effects using 5 and 10 lags.
2. In this exercise, you will estimate ARCH and GARCH models for the daily returns on Microsoft and the S&P 500.
   (a) First, estimate an ARCH(5) model for the each series. What is the sum of the ARCH coefficients?
   (b) Plot the in-sample estimates of the conditional volatility and compare with the absolute returns. What is the estimate of the unconditional volatility? Compare this with the sample standard deviation of returns.

(c) Next, estimate a GARCH(1,1) model for each series. What is the sum of the ARCH and GARCH coefficients? Plot the in-sample estimates of the conditional volatility and compare with the absolute returns. Do the conditional volatility estimates look similar to the ARCH(5) estimates? What is the estimate of the unconditional volatility? Compare this with the sample standard deviation of returns and the estimate from the ARCH(5) model.

3. Using the estimated ARCH(5) and GARCH(1,1) models, compute h-step-ahead volatility forecasts for $h = 1, \ldots, 100$ days and plot these along with the estimate of unconditional volatility. Which forecasts converge faster to the unconditional volatility?

6 Rolling Estimation and VaR with GARCH

The **ugarchroll()** function allows you to perform a rolling estimation and forecasting of a model. Once we specify the model with **ugarchspec()**, we can apply **ugarchroll()** for rolling estimation. Take a look at the examples in the rugarch vignette and read the help file for **ugarchroll**() so you understand how it works.

1. Suppose the continuously compounded return r_t follows a normal GARCH(1,1). Write down the formulas for 1 day conditional VaR and h day conditional VaR.

2. Estimate normal GARCH(1,1) models for S&P 500 and Microsoft (using all of the data) and calculate 1 day ahead VaR and 10 day ahead VaR.

3. Compute 1000, 1-period ahead rolling forecasts of a normal GARCH(1,1) for S&P 500 and Microsoft, refitting the GARCH model every 20 observations (this is a monthly rolling window). In ugarchroll(), `n.ahead=1` means compute 1-step ahead forecasts, `forecast.length=1000` means set aside 1000 observations for the forecast, `refit.every=20` means we will refit our model every 20 days for calculation speed, and `refit.window="moving"` is the rolling forecast. Make sure to calculate 1% VaR when you do the rolling estimation. What is the MSE and MAE of the forecast?. Plot the forecast conditional volatility.

4. You can use the rolling GARCH results to backtesting VaR models. What is the expected 1% VaR violation and what is the actual violation for S&P 500 and Microsoft? Is the model rejected by unconditional coverage test? Is the model reject conditional coverage test for independence?