Eco 4306 Economic and Business Forecasting

Lecture 28

Chapter 15: Financial Applications of Time-Varying Volatility

Motivation

- investors and financial institutions allocate capital among different assets with different amount of risk
- some of the applications of modeling and forecasting the time-varying conditional variance: risk management, portfolio allocation, asset pricing, and option pricing

15.1 Risk Management

- main issue in risk management: assessment of losses in a probabilistic fashion
- various approaches to risk evaluation, offering complementary views of risk
- we will analyze two of these measures: **value-at-risk** and **expected shortfall**

15.1.1 Value-at-Risk (VaR)

- suppose that you are managing a portfolio of assets and you have a long position (you are a buyer of assets)
- a negative scenario for your portfolio: prices of the assets go down, positive scenario: prices go up
- potential maximum loss: all assets in your portfolio become worthless, resulting in 100% capital loss
- but what is the *probability* of such an event?
- more generally, we may wish to assess the probability of a 40%, 30%, or 10% loss
- or, equivalently, we may want to determine how much capital would be lost if a low-probability negative event were to happen
- these are the fundamental questions behind value-at-risk (VaR)

15.1.1 Value-at-Risk (VaR)

- VaR calculations are very prominent among financial institutions
- U.S. banking institutions need to maintain minimum capital requirements, which regulatory agency monitors periodically
- Basle Accord endorses the VaR methodology to assess and monitor market risk capital requirements
- regulators require the institution to calculate the 1% VaR for a 10-day horizon, and to hold enough capital to cover the potential losses assessed by the VaR measure

15.1.1 Value-at-Risk (VaR)

- **value-at-risk (VaR)**: for a random variable r_t, e.g. portfolio return, we define the α-VaR, denoted as $r_t^{VaR(\alpha)}$, as the value of r_t such that the probability of obtaining an equal or smaller value than this is α%

$$P(r_t \leq r_t^{VaR(\alpha)}) = \alpha$$

- we are thus essentially interested in the quantiles of a random variable r_t: using cummulative distribution function F for random variable r_t we have

$$r_t^{VaR(\alpha)} = F^{-1}(\alpha)$$

- note that VaR is the *minimum* loss that occurs for a given probability of tail event

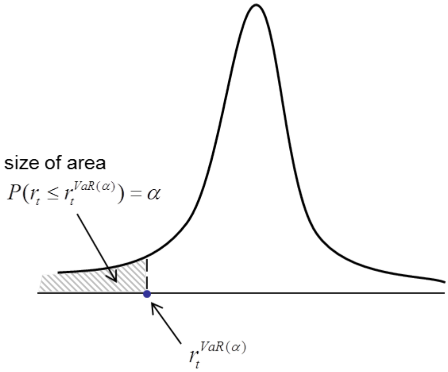

15.1.1 Value-at-Risk (VaR)

- consider the stochastic process of returns to a portfolio of assets

$$r_t = \mu_{t|t-1} + \sigma_{t|t-1} z_t$$

where $z_t \sim N(0,1)$ is iid white noise, and $\mu_{t|t-1}$ and $\sigma_{t|t-1}$ are the conditional mean and conditional standard deviation

- by applying the definition of VaR and standardizing the random variable r_t we get

$$\begin{aligned}
\alpha &= P(r_t \leq r_t^{VaR(\alpha)}) \\
&= P\left(\frac{r_t - \mu_{t|t-1}}{\sigma_{t|t-1}} \leq \frac{r_t^{VaR(\alpha)} - \mu_{t|t-1}}{\sigma_{t|t-1}}\right) \\
&= P\left(z_t \leq \frac{r_t^{VaR(\alpha)} - \mu_{t|t-1}}{\sigma_{t|t-1}}\right) = \Phi\left(\frac{r_t^{VaR(\alpha)} - \mu_{t|t-1}}{\sigma_{t|t-1}}\right)
\end{aligned}$$

where Φ is the cdf of a standard normal distribution

- for the α-VaR we thus have

$$r_t^{VaR(\alpha)} = \mu_{t|t-1} + \Phi^{-1}(\alpha)\sigma_{t|t-1}$$

where Φ^{-1} is the inverse of the cdf of a standard normal distribution (so a normal quantile function)

- since $\Phi^{-1}(0.05) = -1.645$, the 5% VaR is $r_t^{VaR(0.05)} = \mu_{t|t-1} - 1.645\sigma_{t|t-1}$
- since $\Phi^{-1}(0.01) = -2.326$, the 1% VaR is $r_t^{VaR(0.01)} = \mu_{t|t-1} - 2.326\sigma_{t|t-1}$

15.1.1 Value-at-Risk (VaR)

- consider the GARCH(1,1) model for S&P 500 daily returns from 1/2/1998 to 7/25/2008 that we estimated last time

Dependent Variable: R
Method: ML ARCH - Normal distribution (OPG - BHHH / Marquardt steps)
Date: 05/05/19 Time: 16:52
Sample: 1/02/1998 7/25/2008
Included observations: 2657
Convergence achieved after 54 iterations
Coefficient covariance computed using outer product of gradients
Presample variance: unconditional
GARCH = C(2) + C(3)*RESID(-1)^2 + C(4)*GARCH(-1)

Variable	Coefficient	Std. Error	z-Statistic	Prob.
C	0.036082	0.018125	1.990718	0.0465
Variance Equation				
C	0.010592	0.001835	5.772041	0.0000
RESID(-1)^2	0.066302	0.006754	9.816720	0.0000
GARCH(-1)	0.926581	0.007286	127.1660	0.0000

R-squared	-0.000527	Mean dependent var	0.009761
Adjusted R-squared	-0.000527	S.D. dependent var	1.146761
S.E. of regression	1.147063	Akaike info criterion	2.888312
Sum squared resid	3494.644	Schwarz criterion	2.897172
Log likelihood	-3833.123	Hannan-Quinn criter.	2.891519
Durbin-Watson stat	2.079153		

- the estimated GARCH(1,1) model with normally dstributed innovations is

$$
\begin{aligned}
r_t &= 0.036 + \varepsilon_t \\
\varepsilon_t &= \sigma_{t|t-1} z_t \qquad z_t \sim N(0,1) \\
\sigma^2_{t|t-1} &= 0.010 + 0.065\varepsilon^2_{t-1} + 0.927\sigma^2_{t-1|t-2}
\end{aligned}
$$

- we can use this model to construct the forecast for 1-step-ahead conditional mean $\mu_{t+1|t}$ and standard deviation $\sigma_{t+1|t}$ to calculate the 1-step-ahead VaR

15.1.1 Value-at-Risk (VaR)

to calculate the 1-step-ahead 1%-VaR after estimating the GARCH model in EViews:

- click on **Forecast** button, enter **r_f** into "Forecast name" box and **sigmasq_f** into "GARCH (optional)" box", change"Method" to "Static forecast"
- after that select **Object** → **Generate Series** and enter the following **VaR_1pct = r_f + @qnorm(0.01)*sigmasq_f^0.5**
- note: **@qnorm(0.01)** calculates the 1% quantile of the standard normal distribution, if we wanted to construct the 5% VaR we would need to change this into **@qnorm(0.05)**
- finally, to create an indicator whether the actual return is below the 1% VaR, select **Object** → **Generate Series** and enter **x = (r < VaR_1pct)**

15.1.1 Value-at-Risk (VaR)

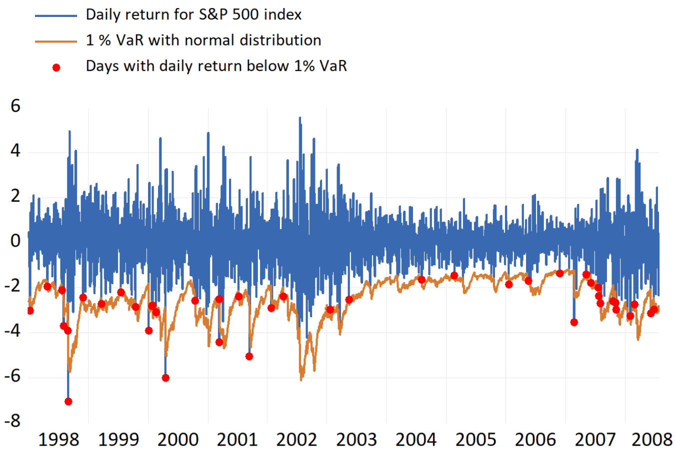

15.1.1 Value-at-Risk (VaR)

- for instance, for April 2, 2008, the 1-day-ahead conditional mean is $\mu_{t|t-1} = 0.036$, the 1-day-ahead conditional standard deviation $\sigma_{t|t-1} = 1.785$, so that the 1-day-ahead 1% VaR is $0.036 - 2.326 \times 1.785 = -4.117\%$
- thus if on April 1, we have a portfolio of $100,000, there is 1% chance that we could lose at least $4,117 on April 2
- observe that, over the time series plot, there are some violations of the 1% boundary - these are the days in which the actual returns are below the VaR
- theoretically, since there are 2657 observations in the sample, 1% of these observations, so about 26, should be below the 1% VaR
- the actual number of violations is 42, which is noticebly higher and represents 1.58% of observations

15.1.1 Value-at-Risk (VaR)

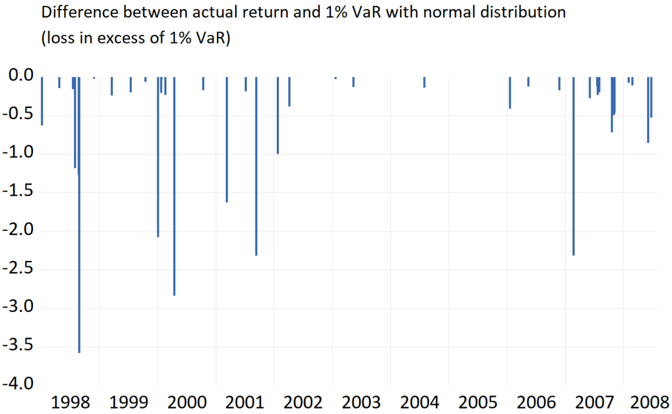

15.1.1 Value-at-Risk (VaR)

- normal density is not well suited to account for excess kurtosis that most financial time series exhibit
- it is thus more common to use Student-t distribution or Generalized Error Distribution (GED) for ARCH/GARCH models because they have fatter tails than normal distribution

15.1.1 Value-at-Risk (VaR)

in EViews, to estimate a GARCH(1,1) with Student-t innovations enter the following information in the specification window:

- estimation settings: choose “ARCH - Autoregressive Conditional Heteroscedasticity”" instead of “LS - Least Squares”
- mean equation: r c
- variance and distribution specification: ARCH 1, GARCH 1
- error distribution: Student's t

15.1.1 Value-at-Risk (VaR)

- the estimated GARCH(1,1) model with Student-t innovations is

$$
\begin{aligned}
r_t &= 0.045 + \varepsilon_t \\
\varepsilon_t &= \sigma_{t|t-1} z_t \qquad z_t \sim t(9.22) \\
\sigma^2_{t|t-1} &= 0.006 + 0.063\varepsilon^2_{t-1} + 0.933\sigma^2_{t-1|t-2}
\end{aligned}
$$

- degrees of freedom paramater is estimated as $\hat{\nu} = 9.22$; some econometricians would round this down to the closest integer, but this is not crucial

Dependent Variable: R
Method: ML ARCH - Student's t distribution (OPG - BHHH / Marquardt steps)
Date: 05/05/19 Time: 16:52
Sample (adjusted): 1/02/1998 7/25/2008
Included observations: 2657 after adjustments
Convergence achieved after 21 iterations
Coefficient covariance computed using outer product of gradients
Presample variance: backcast (parameter = 0.7)
GARCH = C(2) + C(3)*RESID(-1)^2 + C(4)*GARCH(-1)

Variable	Coefficient	Std. Error	z-Statistic	Prob.
C	0.045172	0.016959	2.663592	0.0077
Variance Equation				
C	0.006595	0.002681	2.459674	0.0139
RESID(-1)^2	0.063880	0.009078	7.036415	0.0000
GARCH(-1)	0.933170	0.009115	102.3810	0.0000
T-DIST. DOF	9.224482	1.335634	6.906442	0.0000

R-squared	-0.000954	Mean dependent var	0.009761
Adjusted R-squared	-0.000954	S.D. dependent var	1.146761
S.E. of regression	1.147308	Akaike info criterion	2.864795
Sum squared resid	3496.135	Schwarz criterion	2.875869
Log likelihood	-3800.880	Hannan-Quinn criter.	2.868803
Durbin-Watson stat	2.078266		

15.1.1 Value-at-Risk (VaR)

- for the Student-t distribution with ν degrees of freedom, the 1% VaR is calculated as

$$r_t^{VaR(0.01)} = \mu_{t|t-1} + F_\nu^{-1}(0.01)\sqrt{\frac{\nu-2}{\nu}}\sigma_{t|t-1}$$

where ν is the parameter for degrees of freedon of the distribution, and F_ν^{-1} is the inverse of the Student-t cdf function with ν degrees of freedom

15.1.1 Value-at-Risk (VaR)

in EViews, to calculate the 1-step-ahead 1%-VaR after estimating the GARCH model with Student-t innovations:

- click on **Forecast** button, enter **r_f** into "Forecast name" box and **sigmasq_f** into "GARCH (optional)" box", change"Method" to "Static forecast"
- after that select **Object** $\rightarrow$ **Generate Series** and enter the following **VaR_1pct = r_f + @qtdist(0.01,9.22)*(7.22/9.22)^0.5*sigmasq_f^0.5**
- note: **@qtdist(0.01,9.22)** calculates the 1% quantile of the Student-t distribution with 9.22 degrees of freedom, if we wanted to construct the 5% VaR we would need to change this into **@tdist(0.05,9.22)**
- finally, to create an indicator whether the actual return is below the 1% VaR, select **Object** $\rightarrow$ **Generate Series** and enter **x = (r < VaR_1pct)**

15.1.1 Value-at-Risk (VaR)

- for example, for April 2, 2008, the 1-day-ahead conditional mean is $\mu_{t|t-1} = 0.045$, the 1-day-ahead conditional standard deviation $\sigma_{t|t-1} = 1.802$, so that the 1-day-ahead 1% VaR with 9 df is $0.045 - 2.821 \times \sqrt{7/9} \times 1.802 = -4.440$%
- the 1% VaR is larger in magnitude compared to -4.16% obtained for normal distribution, because of the fat-tail property of the Student-t
- number of violations now is 30 so 1.13% of the sample, which is considerably closer to the theoretical value of 26 than 42 violations under Normal distribution

15.1.1 Value-at-Risk (VaR)

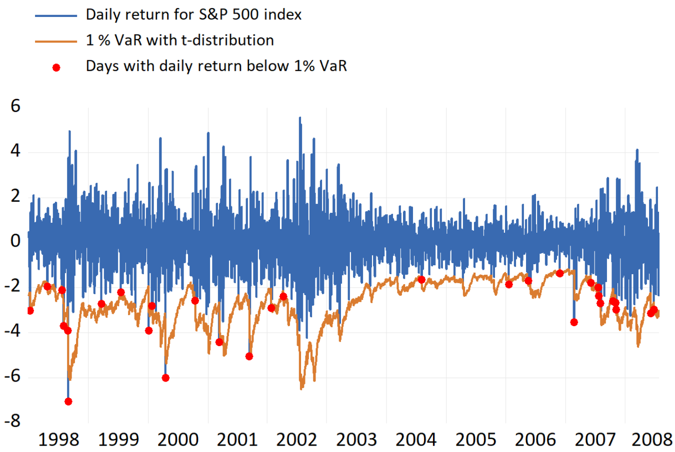

15.1.1 Value-at-Risk (VaR)

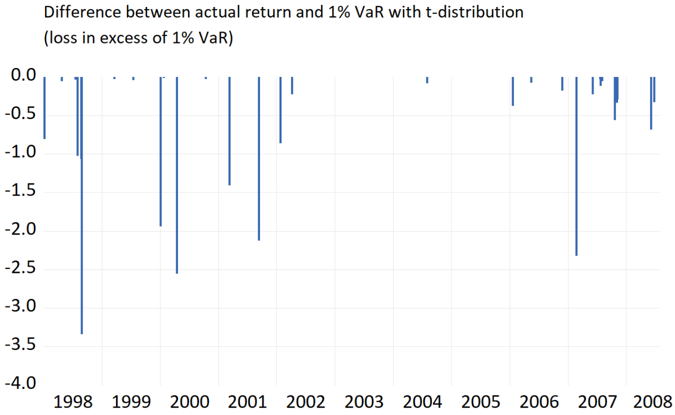

15.1.1 Value-at-Risk (VaR)

- regulators usually require the calculation of VaR for a 10-day horizon
- this is accomplished by using the rule of "square root to time" to extend the daily VaR forecasts to horizons with multiple trading days
- if we are interested in a 10-day horizon, we multiply the daily forecast by $\sqrt{10}$
- thus, on April 2, 2008, the 10-day-ahead 1% VaR under normality will be $\sqrt{10} \times 4.117\% = -13.01\%$
- thus if on April 2 we have a portfolio of $100,000, there is 1% chance that 10 days later, on April 12, we could face a loss of at least $13,010
- under Student-t with $\nu = 9$, the 10-day-ahead 1% VaR will be -14.04%, which means that we could lose at least $14,040 in our $100,000 portfolio

15.1.2 Expected Shortfall (ES)

- VaR is the *minimum* loss that we should expect with $\alpha\%$ probability, but actual losses could be higher
- it is of interest to have a measure of the *average* loss within the observations contained in the $\alpha\%$ region
- that is: the expected value of r_t for only those values where $r_t < r_t^{VaR(\alpha)}$

$$ES(\alpha) = E(r_t | r_t < r_t^{VaR(\alpha)})$$

- this measure is called the **expected shortfall (ES)**
- expected shortfall is also referred to as **Conditional Value at Risk (CVaR)** and **expected tail loss (ETL)**

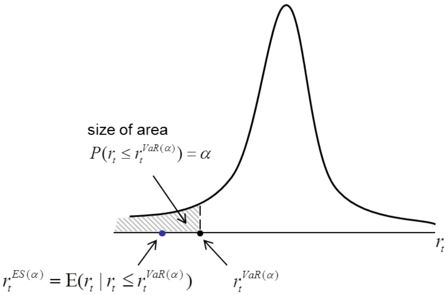

15.1.2 Expected Shortfall (ES)

- with innovation z_t drawn from normal distribution density, the formula to compute the expected shortfall as follows
- for a standard normal random variable z we have

$$E(z|z < z_\alpha) = -\frac{1}{\alpha}\frac{1}{\sqrt{2\pi}}e^{-\frac{z_\alpha^2}{2}}$$

  where z_α is the α% quantile of the standard normal distribution
- thus since $r_t = \mu_{t|t-1} + \sigma_{t|t-1} z_t$ we get

$$ES(\alpha) = E(r_t|r_t < r_t^{VaR(\alpha)}) = \mu_{t|t-1} + E(z|z < z_\alpha)\sigma_{t|t-1}$$

- for $\alpha = 0.05$ we have $z_\alpha = -1.645$, thus $ES(0.05) = \mu_{t|t-1} - 2.0622\sigma_{t|t-1}$
- for $\alpha = 0.01$ we have $z_\alpha = -2.326$, thus $ES(0.05) = \mu_{t|t-1} - 2.6426\sigma_{t|t-1}$

15.1.2 Expected Shortfall (ES)

- for instance, for April 2, 2008, the 1-day-ahead conditional mean is $\mu_{t|t-1} = 0.036$, the 1-day-ahead conditional standard deviation $\sigma_{t|t-1} = 1.785$,
- we calculated the 1-day-ahead 1% VaR ased on GARCH(1,1) with normal innovations to be $0.036 - 2.326 \times 1.785 = -4.117$%
- corresponding expected shortfall for the 1% VaR is $0.036 - 2.0622 \times 1.785 = -4.681$%, which is the average of the values of r_t within the interval $(-\infty, -4.117)$
- so if on April 1 we have a portfolio of $100,000, there is 1% chance that on April 2 we would have a minimum loss of $4,117 and an average loss of $4,681.