Estimating volatility and correlations
Introduction
Estimating volatility: EWMA and GARCH(1,1)
Maximum Likelihood methods
Using GARCH (1, 1) model to forecast volatility
Correlations
Extensions of GARCH
References

Lecture Quantitative Finance
Spring Term 2015

Prof. Dr. Erich Walter Farkas

Lecture 12: May 21, 2015

Estimating volatility and correlations
Introduction
Estimating volatility: EWMA and GARCH(1,1)
Maximum Likelihood methods
Using GARCH (1, 1) model to forecast volatility
Correlations
Extensions of GARCH

References

Outline

1 Estimating volatility and correlations
- Introduction
- Estimating volatility: EWMA and GARCH(1,1)
- Maximum Likelihood methods
- Using GARCH $(1, 1)$ model to forecast volatility
- Correlations
- Extensions of GARCH

Estimating volatility and correlations
Introduction
Estimating volatility: EWMA and GARCH(1,1)
Maximum Likelihood methods
Using GARCH (1, 1) model to forecast volatility
Correlations
Extensions of GARCH

References

Outline of the Presentation

1. Estimating volatility and correlations
   - Introduction
   - Estimating volatility: EWMA and GARCH(1,1)
   - Maximum Likelihood methods
   - Using GARCH $(1, 1)$ model to forecast volatility
   - Correlations
   - Extensions of GARCH

Introduction

- The goal of this chapter is to explain how historical data can be used to produce estimates of the current and future levels of volatilities and correlations.
- This problem is relevant both for the calculation of risk measures (such as Value-at-Risk) and for the valuation of derivatives.
- We consider the following three models:
  - (i) the exponentially weighted moving average (EWMA) model;
  - (ii) the autoregressive conditional heteroscedascity (ARCH) model;
  - (iii) the generalized ARCH (GARCH) model.
- The distinctive feature is that these models recognize that volatilities and correlations are not constant.
- During some periods, a particular volatility or correlation may be relatively low, whereas during other periods it may be relatively high.
- The models attempt to keep track of the variations in the volatility or correlation through time.

Introduction

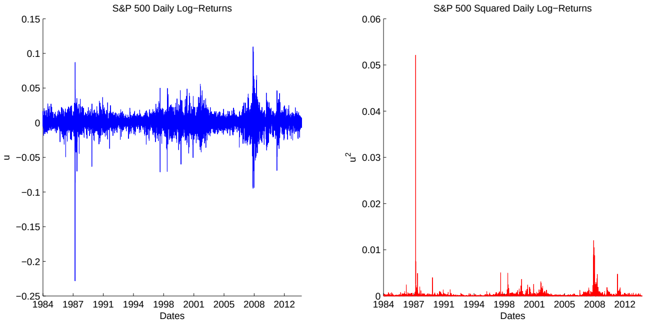

Introduction

- To estimate the volatility of a stock from (empirical) data, the price of the stock is observed at fixed intervals of time (e.g. every day, week, or month).
- Consider

$$\begin{array}{rcl} n+1 & : & \text{number of observations} \\ S_i & : & \text{stock price at the end of the } i^{th} \text{ interval, with } i = 0, 1, \ldots n \\ \tau & : & \text{length of the time intervals in years} \end{array}$$

and let

$$u_i = \log\left(\frac{S_i}{S_{i-1}}\right) \quad i = 1, 2, \ldots, n.$$

- The usual estimate s of the standard deviation of the u_i's is given by

$$s = \sqrt{\frac{1}{n-1}\sum_{i=1}^{n}(u_i - \overline{u})^2},$$

where $\overline{u}$ is the sample mean of u_i.
- The annualized volatility σ can be estimated as $\widehat{\sigma} = \frac{s}{\sqrt{\tau}}$
- The standard error of this estimate can be shown to be approximately $\widehat{\sigma}/(\sqrt{2n})$.

Introduction

- Choosing an appropriate value for n is not easy.
- More data generally leads to more accuracy, but σ does change over time and data that is too old may not be relevant for predicting the future volatility.
- A compromise that seems to work reasonably well is to use closing prices from daily data over the most recent 90 to 180 days.
- An often used rule of thumb is to set n equal to the number of days to which the volatility is applied.
- Thus, if the volatility estimate is used to value a 2-year option, daily data for the last 2 years are used.

Introduction: Example

A sequence of stock prices during 21 consecutive trading days:

Day	Closing price	S_i/S_{i-1}	$\log(S_i/S_{i-1})$
0	20.00		
1	20.10	1.00500	0.00499
2	19.90	0.99005	0.01000
3	20.00	1.00503	0.00501
4	20.50	1.02500	0.02469
5	20.25	0.98780	-0.01227
6	20.90	1.03210	0.03159
7	20.90	1.00000	0.00000
8	20.90	1.00000	0.00000
9	20.75	0.99282	-0.00720
10	20.75	1.00000	0.00000
11	21.00	1.01205	0.01198
12	21.10	1.00476	0.00475
13	20.90	0.99052	-0.00952
14	20.90	1.00000	0.00000
15	21.25	1.01675	0.01661
16	21.40	1.00706	0.00703
17	21.40	1.00000	0.00000
18	21.25	0.99299	-0.00703
19	21.75	1.02353	0.02326
20	22.00	1.01149	0.01143

In this case

$$\sum u_i = 0.09531 \quad \text{and} \quad \sum u_i^2 = 0.00326$$

Introduction: Example

- The estimate of the standard deviation of daily returns is

$$\sqrt{\frac{0.00326}{19} - \frac{0.09531^2}{20 \cdot 19}} = 0.01216 \quad \text{or} \quad 1.216\%.$$

- Assuming that there are 252 trading days per year, i.e., $\tau = 1/252$, an estimate for the volatility per annum is

$$0.01216 \times \sqrt{252} = 0.193 \quad \text{or} \quad 19.3\%.$$

- The standard error of this estimate is

$$\frac{0.193}{\sqrt{2 \times 20}} = 0.031 \quad \text{or} \quad 3.1\% \quad \text{per annum.}$$

- With dividend paying stocks: the return u_i during a time interval that includes an ex-dividend day is given by

$$u_i = \log \frac{S_i + D_i}{S_{i-1}}$$

where D_i is the amount of the dividend at time i.

Introduction: Trading days vs. calendar days

- An important issue is whether time should be measured in calendar days or trading days when volatility parameters are being estimated and used.
- Practitioners tend to ignore days on which the exchange is closed when estimating volatility from historical data (and when calculating the life of an option).
- The volatility per annum is calculated from the volatility per trading day using the formula

$$\text{Vol per annum} = \text{Vol per tr. day} \times \sqrt{\text{nr. of tr. days per annum}}.$$

Introduction: What causes volatility?

- It is natural to assume that the volatility of a stock is caused by new information reaching the market.
- New information causes people to revise their opinions about the value of the stock: the price of the stock changes and volatility results.
- With several years of daily stock price date researchers have calculated:
  - (i) the variance of the stock price returns between the close of trading on one day and the close of trading on the next day when there are no intervening non-trading days (in fact a variance of returns over a 1-day period);
  - (ii) the variance of the stock price returns between the close of trading on Friday and the close of trading on Monday (in fact a variance of returns over 3-day period).

Introduction: What causes volatility?

- We might be tempted to expect the second variance to be three times as great as the first variance but this is not the case (Fama 1965, French 1980, French and Roll 1980: second variance is, respectively, only 22%, 19% and 10.7% higher than the first variance).
- It seems that volatility is to a large extent caused by trading itself.

Outline of the Presentation

1 Estimating volatility and correlations
- Introduction
- Estimating volatility: EWMA and GARCH(1,1)
- Maximum Likelihood methods
- Using GARCH (1, 1) model to forecast volatility
- Correlations
- Extensions of GARCH

Estimating volatility and correlations
Introduction
Estimating volatility: EWMA and GARCH(1,1)
Maximum Likelihood methods
Using GARCH (1, 1) model to forecast volatility
Correlations
Extensions of GARCH
References

Estimating volatility

- Define σ_n the volatility of a market variable on day n, as estimated at the end of day $n-1$.
- The square of the volatility σ_n^2 on day n is the variance rate.
- Recall that the variable u_i is defined as the continuously compounded return between the end of day $i-1$ and the end of day i:

$$u_i = \log \frac{S_i}{S_{i-1}}.$$

- An unbiased estimate of the variance rate per day, σ_n^2, using the most recent m observations on the u_i is

$$\sigma_n^2 = \frac{1}{m-1} \sum_{i=1}^{m} (u_{n-i} - \overline{u})^2$$

where the mean $\overline{u}$ is given by

$$\overline{u} = \frac{1}{m} \sum_{i=1}^{m} u_{n-i}.$$

Estimating volatility

- For the purposes of monitoring daily volatility the last formula is usually changed in a number of ways
  (i) u_i is defined as the percentage change in the market variable between the end of day $i-1$ and the end of day i, so that
  $$u_i = \frac{S_i - S_{i-1}}{S_{i-1}};$$
  (ii) $\overline{u}$ is assumed to be zero;
  (iii) $m-1$ is replaced by m.
- These three changes make very little difference to the estimates that are calculated but they allow us to simplify the formula for the variance rate to
$$\sigma_n^2 = \frac{1}{m}\sum_{i=1}^{m} u_{n-i}^2.$$
- The last expression gives equal weight to $u_{n-1}^2, u_{n-2}^2, \ldots, u_{n-m}^2$.
- Our objective is to estimate the current level of volatility σ_n, therefore it makes sense to give more weight to recent data.

Estimating volatility

- We can accomplish this with a model that sets:

$$\sigma_n^2 = \sum_{i=1}^{m} \alpha_i u_{n-i}^2. \qquad (1)$$

- The coefficient $\alpha_i > 0$ is the weight given to the observation i days ago.
- If we choose them so that $\alpha_i < \alpha_j$ when $i > j$, less weight is given to older observations.
- The weights must sum up to unity, so we have

$$\sum_{i=1}^{m} \alpha_i = 1.$$

Estimating volatility

- An extension of the idea in Eq. (1) is to assume that there is a long-run average variance rate and that this should be given some weight.
- This leads to a model that takes the form

$$\sigma_n^2 = \gamma V_L + \sum_{i=1}^{m} \alpha_i u_{n-i}^2, \qquad (2)$$

  where V_L is the long-run variance rate and γ is the weight assigned to V_L.
- Because the weights must sum to unity, we have

$$\gamma + \sum_{i=1}^{m} \alpha_i = 1.$$

Estimating volatility

- This is known as an ARCH(m) model and it was first suggested by Robert Engle in 1982.
- The estimate of the variance is based on a long-run average variance and m observations: the older an observation, the less weight it is given.
- Defining $\omega = \gamma V_L$ the model in Eq. (2) can be written as

$$\sigma_n^2 = \omega + \sum_{i=1}^{m} \alpha_i u_{n-i}^2.$$

- This is the version of the model used when the parameters are estimated.

The EWMA model

- The exponentially weighted moving average (EWMA) model is a particular case of the model in Eq. (1) where the weights α_i decrease exponentially as we move back through time.
- Specifically $\alpha_{i+1} = \lambda\alpha_i$ where λ is a constant between 0 and 1.
- It turns out that this weighting scheme leads to a particularly simple formula for updating volatility estimates:

$$\sigma_n^2 = \lambda\sigma_{n-1}^2 + (1-\lambda)u_{n-1}^2. \qquad (3)$$

- The estimate σ_n is the volatility for day n (made at the end of day $n-1$) is calculated from σ_{n-1} (the estimate that was made at the end of day $n-2$ of the volatility for day $n-1$) and u_{n-1} (the most recent percentage change).

The EWMA model

- To understand why Eq. (3) corresponds to weights that decrease exponentially, we substitute for σ_{n-1}^2 to get

$$\sigma_n^2 = \lambda[\lambda\sigma_{n-2}^2 + (1-\lambda)u_{n-2}^2] + (1-\lambda)u_{n-1}^2,$$

or

$$\sigma_n^2 = (1-\lambda)(u_{n-1}^2 + \lambda u_{n-2}^2) + \lambda^2\sigma_{n-2}^2.$$

- Substituting in a similar way for σ_{n-2}^2 gives

$$\sigma_n^2 = (1-\lambda)(u_{n-1}^2 + \lambda u_{n-2}^2 + \lambda^2 u_{n-3}^2) + \lambda^3\sigma_{n-3}^2.$$

- Continuing in this way we see that

$$\sigma_n^2 = (1-\lambda)\sum_{i=1}^{m}\lambda^{i-1}u_{n-i}^2 + \lambda^m\sigma_{n-m}^2.$$

The EWMA model

- Recall

$$\sigma_n^2 = (1-\lambda)\sum_{i=1}^{m} \lambda^{i-1} u_{n-i}^2 + \lambda^m \sigma_{n-m}^2. \qquad (4)$$

- For large m the term $\lambda^m \sigma_{n-m}^2$ is sufficiently small to be ignored so that Eq. (4) is the same as Eq. (1) with $\alpha_i = (1-\lambda)\lambda^{i-1}$.
- The weights for the u_i decline at rate λ as we move back through time; each weight is λ times the previous weight.

The EWMA model: Example

- Suppose that $\lambda = 0.90$, the volatility estimated for a market variable for day $n-1$ is 1% per day and during day $n-1$ the market variable increased by 2%.
- This means $\sigma_{n-1}^2 = 0.01^2 = 0.0001$ and $u_{n-1}^2 = 0.02^2 = 0.0004$.
- Eq. (3) gives

$$\sigma_n^2 = 0.9 \times 0.0001 + 0.1 \times 0.0004 = 0.00013.$$

- The estimate of the volatility σ_n for day n is therefore $\sqrt{0.00013}$ or 1.14% per day.
- Note that the expected value of u_{n-1}^2 is σ_{n-1}^2 or 0.0001.
- In this example, the realized value of u_{n-1}^2 is greater than the expected value and as a result our volatility estimate increases.
- If the realized value of u_{n-1}^2 had been less than its expected value, our estimate of the volatility would have decreased.

The EWMA model

- The EWMA approach has the attractive feature that relatively little data needs to be stored.
- At any given time we need to remember only the current estimate of the variance rate and the most recent observation on the value of the market variable.
- When we get a new observation on the value of the market variable, we calculate a new daily percentage change and use Eq. (3) to update our estimate of the variance rate.
- The old estimate of the variance rate and the old value of the market variable can then be discarded.
- The EWMA approach is designed to track changes in the volatility.
- The Risk Metrics database, which was originally created by J. P. Morgan and made publicly available in 1994, uses the EWMA model with $\lambda = 0.94$ for updating daily volatility estimates.
- The company found that, across a range of different market variables, this value of λ gives forecasts of the variance rate that come closest to the realized variance rate.

The GARCH(1,1) model

- Proposed by T. Bollerslev in 1986.
- The difference between GARCH(1,1) and EWMA is analogous to the difference between Eq. (1) and Eq. (2).
- In GARCH(1,1) σ_n^2 is calculated from the long-run average variance rate V_L as well as from σ_{n-1} and u_{n-1}.
- The equation for GARCH(1,1) is

$$\sigma_n^2 = \gamma V_L + \alpha u_{n-1}^2 + \beta \sigma_{n-1}^2,$$

  where γ is the weight assigned to V_L, α is the weight assigned to u_{n-1}^2 and β is the weight assigned to σ_{n-1}^2.
- The weights sum up to one

$$\gamma + \alpha + \beta = 1.$$

The GARCH(1,1) model

- The EWMA model is a particular case of GARCH(1,1) where $\gamma = 0$, $\alpha = 1 - \lambda$, $\beta = \lambda$.
- The (1,1) in GARCH(1,1) indicates that σ_n^2 is based on the most recent observation of u^2 and the most recent estimate of the variance rate.
- The more general GARCH(p,q) model calculates σ_n^2 from the most recent p observations of u^2 and the most recent q estimates of the variance rate; GARCH(1,1) is by far the most popular of the GARCH models.
- Setting $\omega = \gamma V_L$, the GARCH(1,1) model can also be written as

$$\sigma_n^2 = \omega + \alpha u_{n-1}^2 + \beta \sigma_{n-1}^2. \qquad (5)$$

- The last form is usually used for the purposes of estimating the parameters.
- For a stable GARCH(1,1) process we require $\alpha + \beta < 1$, otherwise the weight term applied to the long-term variance is negative.

The GARCH(1,1) model: Example

- Suppose that a GARCH(1,1) model is estimated from daily data as

$$\sigma_n^2 = 0.000002 + 0.13u_{n-1}^2 + 0.86\sigma_{n-1}^2.$$

- This corresponds to $\alpha = 0.13$, $\beta = 0.86$ and $\omega = 0.000002$.
- Because $\gamma = 1 - \alpha - \beta$ it follows that $\gamma = 0.01$.
- Because $\omega = \gamma V_L$ it follows that $V_L = 0.0002$.
- In other words, the long-run average variance per day implied by the model is 0.0002.
- This corresponds to a volatility of $\sqrt{0.0002} = 0.014$ or 1.4% per day.
- Suppose that the estimate of the volatility on day $n-1$ is 1.6% per day, so that $\sigma_{n-1}^2 = 0.016^2 = 0.000256$ and that on day $n-1$ the market variable decreased by 1% so that $u_{n-1}^2 = 0.01^2 = 0.0001$.
- Then

$$\sigma_n^2 = 0.000002 + 0.13 \times 0.0001 + 0.86 \times 0.000256 = 0.00023516.$$

- The new estimate of the volatility is therefore $\sqrt{0.00023516} = 0.0153$ or 1.53%.

The GARCH(1,1) model: The weights

- Substituting for σ^2_{n-1} and, afterwards, for σ^2_{n-2} in Eq. (5), we get

$$\sigma^2_n = \omega + \beta\omega + \beta^2\omega + \alpha u^2_{n-1} + \alpha\beta u^2_{n-2} + \alpha\beta^2 u^2_{n-3} + \beta^3\sigma^2_{n-3}.$$

- Continuing in this was we see that the weight applied to u^2_{n-i} is $\alpha\beta^{i-1}$.
- The weights decline exponentially at rate β.
- The parameter β can be interpreted as decay rate; it is similar to the λ in the EWMA model.
- The GARCH(1,1) model is similar to the EWMA model except that, in addition to assign weights that decline exponentially to past u^2 it also assigns some weight to the long-run average volatility.

Estimating volatility and correlations
Introduction
Estimating volatility: EWMA and GARCH(1,1)
Maximum Likelihood methods
Using GARCH (1, 1) model to forecast volatility
Correlations
Extensions of GARCH
References

The GARCH(1,1) model: Mean reversion (optional part)

- The GARCH(1,1) model recognizes that over time the variance tends to get pulled back to a long-run average level of V_L.
- The amount of weight assigned to V_L is $\gamma = 1 - \alpha - \beta$.
- The GARCH(1,1) is a equivalent to a model where the variance V follows the stochastic process

$$dV = a(V_L - V)dt + \xi V dz$$

  where time is measured in days, $a = 1 - \alpha - \beta$ and $\xi = \alpha\sqrt{2}$; this is the mean reverting model.
- The variance has a drift that pulls it back to V_L at rate a.
- When $V > V_L$, the variance has a negative drift; when $V < V_L$ it has a positive drift.

Choosing between the models

- In practice variance rates tend to be mean reverting.
- The GARCH(1,1) model incorporates mean reversion, whereas the EWMA model does not.
- GARCH(1,1) is therefore more appealing than the EWMA model.
- A question that needs to be discussed is how best-fit parameters ω, α, β in GARCH(1,1) can be estimated.
- When the parameter ω is zero, the GARCH(1,1) reduces to EWMA.
- In circumstances where the best-fit value of ω turns out to be negative, the GARCH(1,1) model is not stable and it makes sense to switch to the EWMA model.

Outline of the Presentation

1. Estimating volatility and correlations
   - Introduction
   - Estimating volatility: EWMA and GARCH(1,1)
   - Maximum Likelihood methods
   - Using GARCH $(1,1)$ model to forecast volatility
   - Correlations
   - Extensions of GARCH

Maximum Likelihood (ML) methods

- How are the parameters estimated from historical data in the models we have been considering?
- A commonly applied approach is known as the maximum likelihood (ML) method.
- It involves choosing values for the parameter that maximize the chance (or likelihood) of the data occurring.

In General

- Suppose we have a sample $x_1, x_2, \ldots, x_N$ of N i.i.d. random variables, coming from a parametric model.
- The joint density function of the observations is

$$f(x_1, x_2, \ldots, x_N|\theta) = f_1(x_1|\theta) \cdot f_2(x_2|\theta) \cdot \ldots \cdot f_n(x_N|\theta),$$

  where θ summarises the model parameters.
- The idea of the maximum likelihood (ML) method is to chose θ such that the joint density function is maximised, given the observed sample of data.
- A natural tool to this end is the likelihood function, which we define as

$$\mathcal{L}(\theta|x_1, x_2, \ldots, x_N) := f(x_1, x_2, \ldots, x_N|\theta) = \prod_{i=1}^{N} f_i(x_i|\theta).$$

In General

- To convert the product to summation (which is easier to handle on a computer), we take the logarithm. The result is called the log-likelihood:

$$\log \mathcal{L}(\theta | x_1, x_2, \ldots, x_n) = \sum_{i=1}^{n} \log(f_i(x_i|\theta)).$$

- The ML method estimates θ by finding a value for θ that maximises $\log \mathcal{L}(\theta | x_1, x_2, \ldots, x_n)$, i.e.,

$$\hat{\theta}_{\text{mle}} := \arg\max_{\theta} \log \mathcal{L}(\theta | x_1, x_2, \ldots, x_n).$$

Estimating a constant variance

- Problem: estimate the variance of a variable X from m observations on X when the underlying distribution is normal with zero mean.
- Let $u_1, u_2, \ldots$ denote the sample of m observations.
- Denote the unknown variance parameter by v.
- The likelihood of u_i being observed is the probability density function for X when $X = u_i$

$$\frac{1}{\sqrt{2\pi v}} \exp\left(\frac{-u_i^2}{2v}\right).$$

- The likelihood of m observations occuring in order in which they are observed is

$$\prod_{i=1}^{m}\left[\frac{1}{\sqrt{2\pi v}} \exp\left(\frac{-u_i^2}{2v}\right)\right].$$

- Using the maximum likelihood method, the best estimate of v is the value that maximizes this expression.

Estimating a constant variance

- Maximizing an expression is equivalent to maximizing the logarithm of the expression.
- Taking logarithms and ignoring constant multiplicative factors, it can be seen that we wish to maximize

$$\sum_{i=1}^{m}\left[-\log v-\frac{u_i^2}{v}\right].$$

- Differentiating this expression with respect to v and setting the result equation to zero, we see that the maximum likelihood estimator of v is

$$\frac{1}{m}\sum_{i=1}^{m}u_i^2.$$

Estimating volatility and correlations
Introduction
Estimating volatility: EWMA and GARCH(1,1)
Maximum Likelihood methods
Using GARCH (1, 1) model to forecast volatility
Correlations
Extensions of GARCH
References

Estimating GARCH(1,1) parameters

- How can the likelihood method be used to estimate the parameters when the GARCH(1,1) method or some other volatility update scheme is used?
- Define $v_i = \sigma_i^2$ as the variance estimated for day i.
- We assume that the probability distribution of u_i conditional on the variance is normal.
- A similar analysis to the one just given shows that the best parameters are the ones that maximize

$$\prod_{i=1}^{m} \left[\frac{1}{\sqrt{2\pi v_i}} \exp\left(\frac{-u_i^2}{2v_i} \right) \right].$$

- This is equivalent (taking logarithms) to maximizing

$$\sum_{i=1}^{m} \left[-\log v_i - \frac{u_i^2}{v_i} \right]. \qquad (6)$$

- This is the same expression as above, except that v is replaced by v_i.
- We search iteratively to find the parameters of the model that maximize the expression in (6).

Estimation of GARCH(1,1) parameters: Example

- The data below concern the Japanese yen exchange rate between January 6, 1988 and August 15, 1997.

Date	Day i	S_i	u_i	$v_i = \sigma_i^2$	$-\log(v_i) - u_i^2/v_i$
06-Jan-88	1	0.007728			
07-Jan-88	2	0.007779	0.006599		
08-Jan-88	3	0.007746	-0.004242	0.00004355	9.6283
11-Jan-88	4	0.007816	0.009037	0.00004198	8.1329
12-Jan-88	5	0.007837	0.002687	0.00004455	9.8568
13-Jan-88	6	0.007924	0.011101	0.00004220	7.1529
...	...	...	...	...	...
13-Aug-97	2421	0.008643	0.003374	0.00007626	9.3321
14-Aug-97	2422	0.008493	-0.017309	0.00007092	5.3294
15-Aug-97	2423	0.008495	0.000144	0.00008417	9.3824
					$\sum = 22063.5763$

Estimation of GARCH(1,1) parameters: Example

- The fifth column shows the estimate of the variance rate $v_i = \sigma_i^2$ for day i made at the end of day $i-1$.
- On day 3 we start things off by setting the variance equal to u_2^2.
- On subsequent days, we use equation

$$\sigma_n^2 = \omega + \alpha u_{n-1}^2 + \beta \sigma_{n-1}^2.$$

- The sixth column tabulates the likelihood measure $-\log(v_i) - u_i^2/v_i$.
- The values in the fifth and sixth columns are based on the current trial estimates of ω, α and β: we are interested in maximizing the sum of the members in the sixth column.
- This involves an iterative search procedure.
- In our example the optimal values of the parameters turn out to be

$$\omega = 0.00000176, \quad \alpha = 0.0626, \quad \beta = 0.8976.$$

- The numbers shown in the above table were calculated on the final iteration of the search for the optimal ω, α, and β.
- The long-term variance rate V_L in our example is

$$\frac{\omega}{1-\alpha-\beta} = \frac{0.00000176}{0.0398} = 0.00004422.$$

Estimation of GARCH(1,1) parameters: Example

- When the EWMA model is used, the estimation procedure is relatively simple: we set $\omega = 0$, $\alpha = 1 - \lambda$, and $\beta = \lambda$.
- In the table above, the value of λ that maximizes the objective function is 0.9686 and the value of the objective function is 21995.8377.
- Both GARCH(1,1) and the EWMA method can be implemented by using the solver routine in Excel to search for the values of the parameters that maximize the likelihood function.

Estimating volatility and correlations
Introduction
Estimating volatility: EWMA and GARCH(1,1)
Maximum Likelihood methods
Using GARCH $(1, 1)$ model to forecast volatility
Correlations
Extensions of GARCH

References

Outline of the Presentation

1. Estimating volatility and correlations
   - Introduction
   - Estimating volatility: EWMA and GARCH(1,1)
   - Maximum Likelihood methods
   - Using GARCH $(1, 1)$ model to forecast volatility
   - Correlations
   - Extensions of GARCH

Using GARCH $(1,1)$ to forecast future volatility

- The variance rate estimated at the end of day $n-1$ for day n, when GARCH (1,1) is used, is

$$\sigma_n^2 = (1-\alpha-\beta)V_L + \alpha u_{n-1}^2 + \beta\sigma_{n-1}^2$$

so that

$$\sigma_n^2 - V_L = \alpha(u_{n-1}^2 - V_L) + \beta(\sigma_{n-1}^2 - V_L).$$

- On day $n+t$ in the future, we have

$$\sigma_{n+t}^2 - V_L = \alpha(u_{n+t-1}^2 - V_L) + \beta(\sigma_{n+t-1}^2 - V_L).$$

- The expected value of u_{n+t-1}^2 is σ_{n+t-1}^2, hence

$$\mathbb{E}[\sigma_{n+t}^2 - V_L] = (\alpha+\beta)\mathbb{E}[\sigma_{n+t-1}^2 - V_L].$$

- Using this equation repeatedly yields

$$\mathbb{E}[\sigma_{n+t}^2 - V_L] = (\alpha+\beta)^t(\sigma_n^2 - V_L)$$

or

$$\mathbb{E}[\sigma_{n+t}^2] = V_L + (\alpha+\beta)^t(\sigma_n^2 - V_L). \qquad (7)$$

- This equation forecasts the volatility on day $n+t$ using the information available at the end of the day $n-1$.

Estimating volatility and correlations
Introduction
Estimating volatility: EWMA and GARCH(1,1)
Maximum Likelihood methods
Using GARCH (1, 1) model to forecast volatility
Correlations
Extensions of GARCH
References

Using GARCH(1,1) to forecast future volatility

- In the EWMA model $\alpha + \beta = 1$ and the last equation shows that the expected future variance rate equals the current variance rate.
- When $\alpha + \beta < 1$ the final term in the equation becomes progressively smaller as t increases.
- As mentioned earlier, the variance rate exhibits mean reversion with a reversion level of V_L and a reversion rate of $1 - \alpha - \beta$.
- Our forecast of the future variance rate tends towards V_L as we look further and further ahead.
- This analysis emphasizes the point that we must have $\alpha + \beta < 1$ for a stable GARCH(1,1) process.
- When $\alpha + \beta > 1$ the weight given to the long-term average variance is negative and the process is mean fleeing rather than mean reverting.

Using GARCH(1,1) to forecast future volatility

- In the yen-dollar exchange rare example considered earlier $\alpha + \beta = 0.9602$ and $V_L = 0.00004422$.
- Suppose that our estimate of the current variance rate per day is 0.00006 (this corresponds to a volatility of 0.77% per day).
- In 10 days the expected variance rate is

$$0.00004422 + 0.9602^{10}(0.00006 - 0.00004422) = 0.00005473.$$

- The expected volatility per day is 0.0074, still well above the long-term volatility of 0.00665 per day.
- However the expected variance rate in 100 days is

$$0.00004422 + 0.9602^{100}(0.00006 - 0.00004422) = 0.00004449$$

and the expected volatility per day is 0.00667 very close to the long-term volatility.

Outline of the Presentation

1 Estimating volatility and correlations
- Introduction
- Estimating volatility: EWMA and GARCH(1,1)
- Maximum Likelihood methods
- Using GARCH $(1, 1)$ model to forecast volatility
- Correlations
- Extensions of GARCH

Correlations

- The discussion so far has centered on the estimation and forecasting of volatility.
- Correlations play a key role in the computation of VaR.
- The goal of this section is to show how correlation estimates can be updated in a similar way to volatility estimates.
- Recall that the covariance between two random variables X and Y is defined as

$$\mathbb{E}[(X - \mu_X)(Y - \mu_Y)]$$

  where μ_X and μ_Y are respectively the means of X and Y.
- The correlation between two random variables X and Y is

$$\frac{cov(X, Y)}{\sigma_X \sigma_Y}$$

  where σ_X and σ_Y are the two standard deviations of X and Y and $cov(X, Y)$ is the covariance between X and Y.

Correlations

- Define x_i and y_i as the percentage changes in X and Y between the end of the day $i-1$ and the end of day i:

$$x_i = \frac{X_i - X_{i-1}}{X_{i-1}} \quad \text{and} \quad y_i = \frac{Y_i - Y_{i-1}}{Y_{i-1}}$$

  where X_i and Y_i are the values of X and Y at the end of the day i.
- We also define

$$\begin{aligned}
\sigma_{x,n} &: \text{daily volatility of variable } X \text{ estimated for day } n; \\
\sigma_{y,n} &: \text{daily volatility of variable } Y \text{ estimated for day } n; \\
cov_n &: \text{estimate of covariance between daily changes in } X \text{ and } Y, \\
&\quad \text{calculated on day } n.
\end{aligned}$$

- Our estimate of the correlation between X and Y on day n is

$$\frac{cov_n}{\sigma_{x,n}\sigma_{y,n}}.$$

Correlations

- Using an equal-weighting scheme and assuming that the means of x_i and y_i are zero, as before we can estimate the variance rates of X and Y from the most recent m observations as

$$\sigma_{x,n}^2 = \frac{1}{m}\sum_{i=1}^{m} x_{n-i}^2 \quad \text{and} \quad \sigma_{y,n}^2 = \frac{1}{m}\sum_{i=1}^{m} y_{n-i}^2.$$

- A similar estimate for the covariance between X and Y is

$$cov_n = \frac{1}{m}\sum_{i=1}^{m} x_{n-i} y_{n-i}.$$

Estimating volatility and correlations
Introduction
Estimating volatility: EWMA and GARCH(1,1)
Maximum Likelihood methods
Using GARCH (1, 1) model to forecast volatility
Correlations
Extensions of GARCH
References

Correlations

- One alternative for updating covariances is an EWMA model as previously discussed.
- The formula for updating the covariance estimate is then

$$cov_n = \lambda cov_{n-1} + (1-\lambda) x_{n-1} y_{n-1}.$$

- A similar analysis to that presented for the EWMA volatility model shows that the weights given to observations on the x_i and y_i decline as we move back through time.
- The lower the value of λ, the greater the weight that is given to recent observations.

Estimating volatility and correlations
Introduction
Estimating volatility: EWMA and GARCH(1,1)
Maximum Likelihood methods
Using GARCH (1, 1) model to forecast volatility
Correlations
Extensions of GARCH
References

Correlations: Example

- Assume $\lambda = 0.95$ and that the estimate of the correlation between two variables X and Y on day $n-1$ is 0.6.
- Assume that the estimate of the volatilities for the X and Y on day $n-1$ are 1% and 2% respectively.
- From the relationship between correlation and covariance, the estimate for the covariance between X and Y on day $n-1$ is

$$0.6 \times 0.01 \times 0.02 = 0.00012.$$

- Suppose that the percentage changes in X and Y on day $n-1$ are 0.5% and 2.5% respectively.
- The variance and covariance for day n would be updated as follows:

$$\begin{aligned}
\sigma_{x,n}^2 &= 0.95 \times 0.01^2 + 0.05 \times 0.005^2 = 0.00009625; \\
\sigma_{y,n}^2 &= 0.95 \times 0.02^2 + 0.05 \times 0.025^2 = 0.00041125; \\
cov_n &= 0.95 \times 0.00012 + 0.05 \times 0.005 \times 0.025 = 0.00012025.
\end{aligned}$$

Correlations: Example

- The new volatility of X is $\sqrt{0.00009625} = 0.981\%$.
- The new volatility of Y is $\sqrt{0.00041125} = 2.028\%$.
- The new coefficient of correlation between X and Y is

$$\frac{0.00012025}{0.00981 \times 0.02028} = 0.6044.$$

Correlations

- GARCH models can also be used for updating covariance estimates and forecasting the future level of covariances.
- For example the GARCH(1,1) model for updating a covariance is

$$cov_n = \omega + \alpha x_{n-1} y_{n-1} + \beta cov_{n-1}$$

  and the long-term average covariance is $\omega/(1 - \alpha - \beta)$.
- Similar formulas to those discussed above can be developed for forecasting future covariances and calculating the average covariance during the life time of an option.

Consistency condition for covariances

- Once all the variances and covariances have been calculated, a variance-covariance matrix can be constructed.
- When $i \neq j$, the (i, j) element of this matrix shows the covariance between variable i and j; when $j = i$ it shows the variance of variable i.
- Not all variance-covariance matrices are internally consistent; the condition for an $N \times N$ variance-covariance matrix Ω to be internally consistent is

$$w^T \cdot \Omega \cdot w \geq 0$$

  for all $N \times 1$ vectors w, where w^T is the transpose of w; such a matrix is positive-semidefinite.
- To understand why the last condition must hold, suppose that w^T is $(w_1, ..., w_n)$; the expression $w^T \cdot \Omega \cdot w$ is the variance of $w_1 x_1 + ... + w_n x_n$ where x_i is the value of the variable i; as such it cannot be negative.

Consistency condition for covariances

- To ensure that a positive-semidefinite matrix is produced, variances and covariances should be calculated consistently.
- For example, if variances are calculated by giving equal weight to the last m data items, the same should be done for covariances.
- If variances are updated using an EWMA model with $\lambda = 0.94$ the same should be done for covariances.
- An example of a variance-covariance matrix that it is not internally consistent is

$$\begin{bmatrix} 1 & 0 & 0.9 \\ 0 & 1 & 0.9 \\ 0.9 & 0.9 & 1 \end{bmatrix}$$

- The variance of each variable is 1.0 and so the covariances are also coefficients of correlation.
- The first variable is highly correlated with the third variable and the second variable is highly correlated with the third variable.
- However there is no correlation at all between the first and the second variables; this seems strange; when we set w equal to $(1, 1, -1)$ we find that the positive semi-definiteness condition above is not satisfied proving that the matrix is not positive-semidefinite.

Outline of the Presentation

1. Estimating volatility and correlations
    - Introduction
    - Estimating volatility: EWMA and GARCH(1,1)
    - Maximum Likelihood methods
    - Using GARCH $(1, 1)$ model to forecast volatility
    - Correlations
    - Extensions of GARCH

Estimating volatility and correlations
Introduction
Estimating volatility: EWMA and GARCH(1,1)
Maximum Likelihood methods
Using GARCH (1, 1) model to forecast volatility
Correlations
Extensions of GARCH
References

Extensions of GARCH

- Exponential GARCH (EGARCH)
- Threshold GARCH (TGARCH)

Exponential GARCH (EGARCH)

- The EGARCH model is a GARCH variant that models the logarithm of the conditional variance.
- It includes a leverage term to capture the asymmetric effects between positive and negative asset returns.
- The EGARCH(1,1) model takes the following form:

$$\log \sigma_n^2 = \omega + \alpha g\left(\epsilon_{n-1}\right) + \beta \log \sigma_{n-1}^2,$$

  where $\epsilon_n = u_n/\sigma_n$ and $g(\epsilon_n) = \theta\epsilon_n + \gamma(|\epsilon_n| - E[|\epsilon_n|])$.
- Since negative returns have a more pronounced effect on volatility than positive returns of the same magnitude, the parameter θ usually takes negative values.

Threshold GARCH (TGARCH)

- The TGARCH model is a specification of conditional variance.
- Like the EGARCH model, it allows positive returns to have a larger/smaller impact on volatility than negative returns.
- The TGARCH(1,1) model has the following form:

$$\sigma_n^2 = \omega + (\alpha + \gamma N_{n-1})u_{n-1}^2 + \beta\sigma_{n-1}^2$$

where N_{n-1} is an indicator for negative u_{n-1}, that is

$$N_{n-1} = \begin{cases} 1 \text{ if } u_{n-1} < 0 \\ 0 \text{ if } u_{n-1} \geq 0 \end{cases}$$

References

Books:

- Hull, John C.,*Options, futures and other derivatives*, Sixth Edition, Prentice Hall, 2006, Chapter 19: Estimating Volatilities, pages 461 - 480.
- Tsay, Ruey S.,*Analysis of Financial Time Series*, Third Edition, Wiley, 2010, Chapter 3: Conditional Heteroscedastic Models, pages 143 - 150.