OPTIONS, FUTURES, AND OTHER DERIVATIVES

NINTH EDITION

JOHN C. HULL

CHAPTER 23 Estimating Volatilities and Correlations

In this chapter we explain how historical data can be used to produce estimates of the current and future levels of volatilities and correlations. The chapter is relevant both to the calculation of value at risk using the model-building approach and to the valuation of derivatives. When calculating value at risk, we are most interested in the current levels of volatilities and correlations because we are assessing possible changes in the value of a portfolio over a very short period of time. When valuing derivatives, forecasts of volatilities and correlations over the whole life of the derivative are usually required.

The chapter considers models with imposing names such as exponentially weighted moving average (EWMA), autoregressive conditional heteroscedasticity (ARCH), and generalized autoregressive conditional heteroscedasticity (GARCH). The distinctive feature of the models is that they 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.

23.1 ESTIMATING VOLATILITY

Define σ_n as 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*. We described the standard approach to estimating σ_n from historical data in Section 15.4. Suppose that the value of the market variable at the end of day i is S_i. The variable u_i is defined as the continuously compounded return during day i (between the end of day $i-1$ and the end of day i):

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

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

$$\sigma_n^2 = \frac{1}{m-1} \sum_{i=1}^{m} (u_{n-i} - \bar{u})^2 \tag{23.1}$$

where $\bar{u}$ is the mean of the u_is:

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

For the purposes of monitoring daily volatility, the formula in equation (23.1) is usually changed in a number of ways:

1. 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:[1]

$$u_i = \frac{S_i - S_{i-1}}{S_{i-1}} \tag{23.2}$$

2. $\bar{u}$ is assumed to be zero.[2]
3. $m-1$ is replaced by m.[3]

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

$$\sigma_n^2 = \frac{1}{m}\sum_{i=1}^{m} u_{n-i}^2 \tag{23.3}$$

where u_i is given by equation (23.2).[4]

Weighting Schemes

Equation (23.3) gives equal weight to $u_{n-1}^2, u_{n-2}^2, \dots, u_{n-m}^2$. Our objective is to estimate the current level of volatility, σ_n. It therefore makes sense to give more weight to recent data. A model that does this is

$$\sigma_n^2 = \sum_{i=1}^{m} \alpha_i u_{n-i}^2 \tag{23.4}$$

The variable α_i is the amount of weight given to the observation i days ago. The α's are positive. If we choose them so that $\alpha_i < \alpha_j$ when $i > j$, less weight is given to older observations. The weights must sum to unity, so that

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

[1] This is consistent with the point made in Section 22.3 about the way that volatility is defined for the purposes of VaR calculations.

[2] As explained in Section 22.3, this assumption usually has very little effect on estimates of the variance because the expected change in a variable in one day is very small when compared with the standard deviation of changes.

[3] Replacing $m-1$ by m moves us from an unbiased estimate of the variance to a maximum likelihood estimate. Maximum likelihood estimates are discussed later in the chapter.

[4] Note that the u's in this chapter play the same role as the Δx's in Chapter 22. Both are daily percentage changes in market variables. In the case of the u's, the subscripts count observations made on different days on the same market variable. In the case of the Δx's, they count observations made on the same day on different market variables. The use of subscripts for σ is similarly different between the two chapters. In this chapter, the subscripts refer to days; in Chapter 22 they referred to market variables.

An extension of the idea in equation (23.4) is to assume that there is a long-run average variance rate and that this should be given some weight. This leads to the model that takes the form

$$\sigma_n^2 = \gamma V_L + \sum_{i=1}^{m} \alpha_i u_{n-i}^2 \tag{23.5}$$

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

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

This is known as an ARCH(m) model. It was first suggested by Engle.[5] 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 equation (23.5) can be written

$$\sigma_n^2 = \omega + \sum_{i=1}^{m} \alpha_i u_{n-i}^2 \tag{23.6}$$

In the next two sections we discuss two important approaches to monitoring volatility using the ideas in equations (23.4) and (23.5).

23.2 THE EXPONENTIALLY WEIGHTED MOVING AVERAGE MODEL

The exponentially weighted moving average (EWMA) model is a particular case of the model in equation (23.4) 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. The formula is

$$\sigma_n^2 = \lambda\sigma_{n-1}^2 + (1-\lambda)u_{n-1}^2 \tag{23.7}$$

The estimate, σ_n, of the volatility of a variable 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 daily percentage change in the variable).

To understand why equation (23.7) 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

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

---

[5] See R. Engle "Autoregressive Conditional Heteroscedasticity with Estimates of the Variance of UK Inflation," *Econometrica*, 50 (1982): 987–1008.

Continuing in this way gives

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

For large m, the term $\lambda^m\sigma_{n-m}^2$ is sufficiently small to be ignored, so that equation (23.7) is the same as equation (23.4) 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.

Example 23.1

Suppose that λ is 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 that $\sigma_{n-1}^2 = 0.01^2 = 0.0001$ and $u_{n-1}^2 = 0.02^2 = 0.0004$. Equation (23.7) 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 approach has the attractive feature that relatively little data need be stored. At any given time, only the current estimate of the variance rate and the most recent observation on the value of the market variable need be remembered. When a new observation on the market variable is obtained, a new daily percentage change is calculated and equation (23.7) is used to update the 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. Suppose there is a big move in the market variable on day $n-1$, so that u_{n-1}^2 is large. From equation (23.7) this causes the estimate of the current volatility to move upward. The value of λ governs how responsive the estimate of the daily volatility is to the most recent daily percentage change. A low value of λ leads to a great deal of weight being given to the u_{n-1}^2 when σ_n is calculated. In this case, the estimates produced for the volatility on successive days are themselves highly volatile. A high value of λ (i.e., a value close to 1.0) produces estimates of the daily volatility that respond relatively slowly to new information provided by the daily percentage change.

The RiskMetrics database, which was originally created by JPMorgan and made publicly available in 1994, used the EWMA model with $\lambda = 0.94$ for updating daily volatility estimates. This is because 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.[6] The realized variance rate on a particular day was calculated as an equally weighted average of the u_i^2 on the subsequent 25 days (see Problem 23.19).

[6] See JPMorgan, *RiskMetrics Monitor*, Fourth Quarter, 1995. We will explain an alternative (maximum likelihood) approach to estimating parameters later in the chapter.

23.3 THE GARCH(1,1) MODEL

We now move on to discuss what is known as the GARCH(1,1) model, proposed by Bollerslev in 1986.[7] The difference between the GARCH(1,1) model and the EWMA model is analogous to the difference between equation (23.4) and equation (23.5). In GARCH(1,1), σ_n^2 is calculated from a 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

$$\sigma_n^2 = \gamma V_L + \alpha u_{n-1}^2 + \beta \sigma_{n-1}^2 \tag{23.8}$$

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. Since the weights must sum to unity, it follows that

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

The EWMA model is a particular case of GARCH(1,1) where $\gamma = 0$, $\alpha = 1 - \lambda$, and $\beta = \lambda$.

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 on u^2 and the most recent q estimates of the variance rate.[8] 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

$$\sigma_n^2 = \omega + \alpha u_{n-1}^2 + \beta \sigma_{n-1}^2 \tag{23.9}$$

This is the form of the model that is usually used for the purposes of estimating the parameters. Once ω, α, and β have been estimated, we can calculate γ as $1 - \alpha - \beta$. The long-term variance V_L can then be calculated as ω/γ. For a stable GARCH(1,1) process we require $\alpha + \beta < 1$. Otherwise the weight applied to the long-term variance is negative.

Example 23.2

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

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

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.

---

[7] See T. Bollerslev, "Generalized Autoregressive Conditional Heteroscedasticity," *Journal of Econometrics*, 31 (1986): 307–27.

[8] Other GARCH models have been proposed that incorporate asymmetric news. These models are designed so that σ_n depends on the sign of u_{n-1}. Arguably, the models are more appropriate for equities than GARCH(1,1). As mentioned in Chapter 20, the volatility of an equity's price tends to be inversely related to the price so that a negative u_{n-1} should have a bigger effect on σ_n than the same positive u_{n-1}. For a discussion of models for handling asymmetric news, see D. Nelson, "Conditional Heteroscedasticity and Asset Returns: A New Approach," *Econometrica*, 59 (1990): 347–70; R. F. Engle and V. Ng, "Measuring and Testing the Impact of News on Volatility," *Journal of Finance*, 48 (1993): 1749–78.

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%, per day.

The Weights

Substituting for σ_{n-1}^2 in equation (23.9) gives

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

or

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

Substituting for σ_{n-2}^2 gives

$$\sigma_n^2 = \omega + \beta\omega + \beta^2\omega + \alpha u_{n-1}^2 + \alpha\beta u_{n-2}^2 + \alpha\beta^2 u_{n-3}^2 + \beta^3\sigma_{n-3}^2$$

Continuing in this way, we see that the weight applied to u_{n-i}^2 is $\alpha\beta^{i-1}$. The weights decline exponentially at rate β. The parameter β can be interpreted as a "decay rate". It is similar to λ in the EWMA model. It defines the relative importance of the observations on the u's in determining the current variance rate. For example, if $\beta = 0.9$, then u_{n-2}^2 is only 90% as important as u_{n-1}^2; u_{n-3}^2 is 81% as important as u_{n-1}^2; and so on. The GARCH(1,1) model is similar to the EWMA model except that, in addition to assigning weights that decline exponentially to past u^2, it also assigns some weight to the long-run average volatility.

Mean Reversion

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 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 $\xi = \alpha\sqrt{2}$ (see Problem 23.14). This is a 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. Superimposed on the drift is a volatility ξ. Chapter 27 discusses this type of model further.

23.4 CHOOSING BETWEEN THE MODELS

In practice, variance rates do tend to be mean reverting. The GARCH(1,1) model incorporates mean reversion, whereas the EWMA model does not. GARCH (1,1) is therefore theoretically more appealing than the EWMA model.

In the next section, we will discuss how best-fit parameters ω, α, and 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.

23.5 MAXIMUM LIKELIHOOD METHODS

It is now appropriate to discuss how the parameters in the models we have been considering are estimated from historical data. The approach used is known as the *maximum likelihood method.* It involves choosing values for the parameters that maximize the chance (or likelihood) of the data occurring.

To illustrate the method, we start with a very simple example. Suppose that we sample 10 stocks at random on a certain day and find that the price of one of them declined on that day and the prices of the other nine either remained the same or increased. What is the best estimate of the probability of a stock's price declining on the day? The natural answer is 0.1. Let us see if this is what the maximum likelihood method gives.

Suppose that the probability of a price decline is p. The probability that one particular stock declines in price and the other nine do not is $p(1-p)^9$. Using the maximum likelihood approach, the best estimate of p is the one that maximizes $p(1-p)^9$. Differentiating this expression with respect to p and setting the result equal to zero, we find that $p = 0.1$ maximizes the expression. This shows that the maximum likelihood estimate of p is 0.1, as expected.

Estimating a Constant Variance

Our next example of maximum likelihood methods considers the problem of estimating the variance of a variable X from m observations on X when the underlying distribution is normal with zero mean. Assume that the observations are $u_1, u_2, \dots, u_m$. Denote the variance by v. The likelihood of u_i being observed is defined as the probability density function for X when $X = u_i$. This is

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

The likelihood of m observations occurring in the order in which they are observed is

$$\prod_{i=1}^{m}\left[\frac{1}{\sqrt{2\pi v}} \exp\left(\frac{-u_i^2}{2v}\right)\right] \tag{23.10}$$

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

Maximizing an expression is equivalent to maximizing the logarithm of the expression. Taking logarithms of the expression in equation (23.10) and ignoring constant multiplicative factors, it can be seen that we wish to maximize

$$\sum_{i=1}^{m}\left[-\ln(v) - \frac{u_i^2}{v}\right] \tag{23.11}$$

or

$$-m\ln(v) - \sum_{i=1}^{m}\frac{u_i^2}{v}$$

Differentiating this expression with respect to v and setting the resulting equation to zero, we see that the maximum likelihood estimator of v is[9]

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

Estimating EWMA or GARCH (1,1) Parameters

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

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

Taking logarithms, we see that this is equivalent to maximizing

$$\sum_{i=1}^{m}\left[-\ln(v_i) - \frac{u_i^2}{v_i}\right] \tag{23.12}$$

This is the same as the expression in equation (23.11), except that v is replaced by v_i. It is necessary to search iteratively to find the parameters in the model that maximize the expression in equation (23.12).

The spreadsheet in Table 23.1 indicates how the calculations could be organized for the GARCH(1,1) model. The table analyzes data on the S&P 500 between July 18, 2005, and August 13, 2010.[10] The first column in the table records the date. The second column counts the days. The third column shows the S&P 500, S_i, at the end of day i. The fourth column shows the proportional change in the S&P 500 between the end of day $i-1$ and the end of day i. This is $u_i = (S_i - S_{i-1})/S_{i-1}$. 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, equation (23.9) is used. The sixth column tabulates the likelihood measure, $-\ln(v_i) - u_i^2/v_i$. The values in the fifth and sixth columns are based on the current trial estimates of ω, α, and β. We are interested in choosing ω, α, and β to maximize the sum of the numbers in the sixth column. This involves an iterative search procedure.[11]

---

[9] This confirms the point made in footnote 3.

[10] The data and calculations can be found at www.rotman.utoronto.ca/~hull/OFOD/GarchExample.

[11] As discussed later, a general purpose algorithm such as Solver in Microsoft's Excel can be used. Alternatively, a special purpose algorithm, such as Levenberg–Marquardt, can be used. See, e.g., W. H. Press, B.P. Flannery, S.A. Teukolsky, and W.T. Vetterling. *Numerical Recipes in C: The Art of Scientific Computing*, Cambridge University Press, 1988.

Table 23.1 Estimation of Parameters in GARCH(1,1) Model for S&P 500 between July 18, 2005, and August 13, 2010.

Date	*Day i*	S_i	u_i	$v_i = \sigma_i^2$	$-\ln(v_i) - u_i^2/v_i$
18-Jul-2005	1	1221.13			
19-Jul-2005	2	1229.35	0.006731		
20-Jul-2005	3	1235.20	0.004759	0.00004531	9.5022
21-Jul-2005	4	1227.04	−0.006606	0.00004447	9.0393
22-Jul-2005	5	1233.68	0.005411	0.00004546	9.3545
25-Jul-2005	6	1229.03	−0.003769	0.00004517	9.6906
⋮	⋮	⋮	⋮	⋮	⋮
11-Aug-2010	1277	1089.47	−0.028179	0.00011834	2.3322
12-Aug-2010	1278	1083.61	−0.005379	0.00017527	8.4841
13-Aug-2010	1279	1079.25	−0.004024	0.00016327	8.6209
					10,228.2349

Trial estimates of GARCH parameters

$\omega = 0.0000013465$ $\alpha = 0.083394$ $\beta = 0.910116$

In our example, the optimal values of the parameters turn out to be

$$\omega = 0.0000013465, \quad \alpha = 0.083394, \quad \beta = 0.910116$$

and the maximum value of the function in equation (23.12) is 10,228.2349. The numbers shown in Table 23.1 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

$$\frac{\omega}{1-\alpha-\beta} = \frac{0.0000013465}{0.006490} = 0.0002075$$

The long-term volatility is $\sqrt{0.0002075}$, or 1.4404%, per day.

Figures 23.1 and 23.2 show the S&P 500 index and its GARCH(1,1) volatility during the 5-year period covered by the data. Most of the time, the volatility was less than 2% per day, but volatilities as high as 5% per day were experienced during the credit crisis. (Very high volatilities are also indicated by the VIX index—see Section 15.11.)

An alternative approach to estimating parameters in GARCH(1,1), which is sometimes more robust, is known as *variance targeting*.[12] This involves setting the long-run average variance rate, V_L, equal to the sample variance calculated from the data (or to some other value that is believed to be reasonable). The value of ω then equals $V_L(1-\alpha-\beta)$ and only two parameters have to be estimated. For the data in Table 23.1, the sample variance is 0.0002412, which gives a daily volatility of 1.5531%. Setting V_L equal to the sample variance, the values of α and β that maximize the objective function in equation (23.12) are 0.08445 and 0.9101, respectively. The value of the objective function is 10,228.1941, only marginally below the value of 10,228.2349 obtained using the earlier procedure.

[12] See R. Engle and J. Mezrich, "GARCH for Groups," *Risk*, August 1996: 36–40.

Figure 23.1 S&P 500 index: July 18, 2005, to August 13, 2010.

When the EWMA model is used, the estimation procedure is relatively simple. We set $\omega = 0$, $\alpha = 1 - \lambda$, and $\beta = \lambda$, and only one parameter has to be estimated. In the data in Table 23.1, the value of λ that maximizes the objective function in equation (23.12) is 0.9374 and the value of the objective function is 10,192.5104.

For both GARCH (1,1) and EWMA, we can use the Solver routine in Excel to search for the values of the parameters that maximize the likelihood function. The routine works well provided that the spreadsheet is structured so that the parameters being searched for have roughly equal values. For example, in GARCH (1,1) we could let cells

Figure 23.2 Daily volatility of S&P 500 index: July 18, 2005, to August 13, 2010.

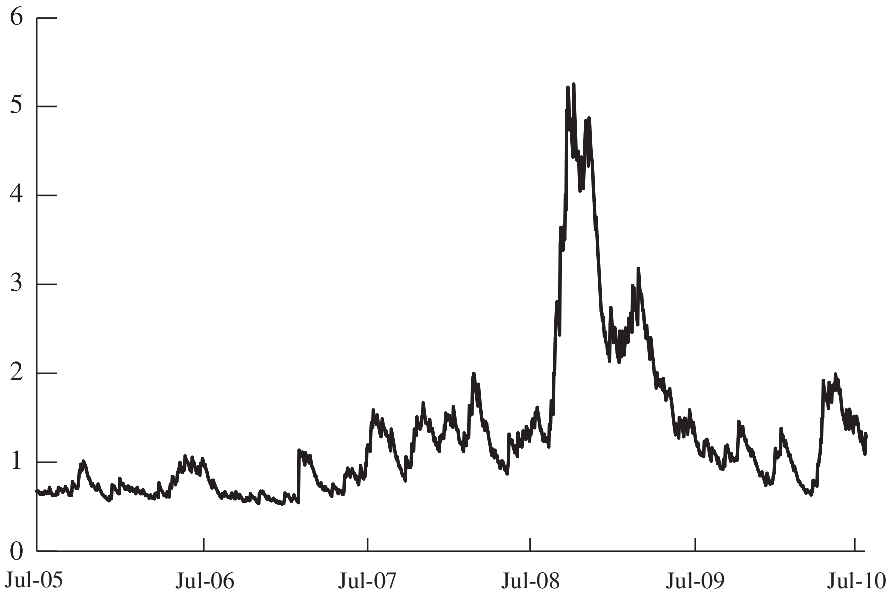

A1, A2, and A3 contain $\omega \times 10^5$, 10α, and β. We could then set B1 = A1/100,000, B2 = A2/10, and B3 = A3. We would use B1, B2, and B3 to calculate the likelihood function. We would ask Solver to calculate the values of A1, A2, and A3 that maximize the likelihood function. Occasionally Solver gives a local maximum, so testing a number of different starting values for parameters is a good idea.

How Good Is the Model?

The assumption underlying a GARCH model is that volatility changes with the passage of time. During some periods volatility is relatively high; during other periods it is relatively low. To put this another way, when u_i^2 is high, there is a tendency for u_{i+1}^2, $u_{i+2}^2, \dots$ to be high; when u_i^2 is low, there is a tendency for u_{i+1}^2, $u_{i+2}^2, \dots$ to be low. We can test how true this is by examining the autocorrelation structure of the u_i^2.

Let us assume the u_i^2 do exhibit autocorrelation. If a GARCH model is working well, it should remove the autocorrelation. We can test whether it has done so by considering the autocorrelation structure for the variables u_i^2/σ_i^2. If these show very little autocorrelation, our model for σ_i has succeeded in explaining autocorrelations in the u_i^2.

Table 23.2 shows results for the S&P 500 data used above. The first column shows the lags considered when the autocorrelation is calculated. The second shows autocorrelations for u_i^2; the third shows autocorrelations for u_i^2/σ_i^2.[13] The table shows that the autocorrelations are positive for u_i^2 for all lags between 1 and 15. In the case of u_i^2/σ_i^2, some of the autocorrelations are positive and some are negative. They are all much smaller in magnitude than the autocorrelations for u_i^2.

Table 23.2 Autocorrelations before and after the use of a GARCH model for S&P 500 data.

Time lag	*Autocorrelation for u_i^2*	*Autocorrelation for u_i^2/σ_i^2*
1	0.183	−0.063
2	0.385	−0.004
3	0.160	−0.007
4	0.301	0.022
5	0.339	0.014
6	0.308	−0.011
7	0.329	0.026
8	0.207	0.038
9	0.324	0.041
10	0.269	0.083
11	0.431	−0.007
12	0.286	0.006
13	0.224	0.001
14	0.121	0.017
15	0.222	−0.031

[13] For a series x_i, the autocorrelation with a lag of k is the coefficient of correlation between x_i and x_{i+k}.

The GARCH model appears to have done a good job in explaining the data. For a more scientific test, we can use what is known as the Ljung–Box statistic.[14] If a certain series has m observations the Ljung–Box statistic is

$$m \sum_{k=1}^{K} w_k \eta_k^2$$

where η_k is the autocorrelation for a lag of k, K is the number of lags considered, and

$$w_k = \frac{m+2}{m-k}$$

For $K = 15$, zero autocorrelation can be rejected with 95% confidence when the Ljung–Box statistic is greater than 25.

From Table 23.2, the Ljung–Box statistic for the u_i^2 series is about 1,566. This is strong evidence of autocorrelation. For the u_i^2/σ_i^2 series, the Ljung–Box statistic is 21.7, suggesting that the autocorrelation has been largely removed by the GARCH model.

23.6 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,

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

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

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

where E denotes expected value. Using this equation repeatedly yields

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

or

$$E[\sigma_{n+t}^2] = V_L + (\alpha+\beta)^t(\sigma_n^2 - V_L) \tag{23.13}$$

This equation forecasts the volatility on day $n+t$ using the information available at the end of day $n-1$. In the EWMA model, $\alpha+\beta=1$ and equation (23.13) 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. Figure 23.3 shows the expected path followed by the variance rate for situations where the current variance rate is different from V_L. 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

[14] See G. M. Ljung and G. E. P. Box, "On a Measure of Lack of Fit in Time Series Models," *Biometrica*, 65 (1978): 297–303.

Figure 23.3 Expected path for the variance rate when (a) current variance rate is above long-term variance rate and (b) current variance rate is below long-term variance rate.

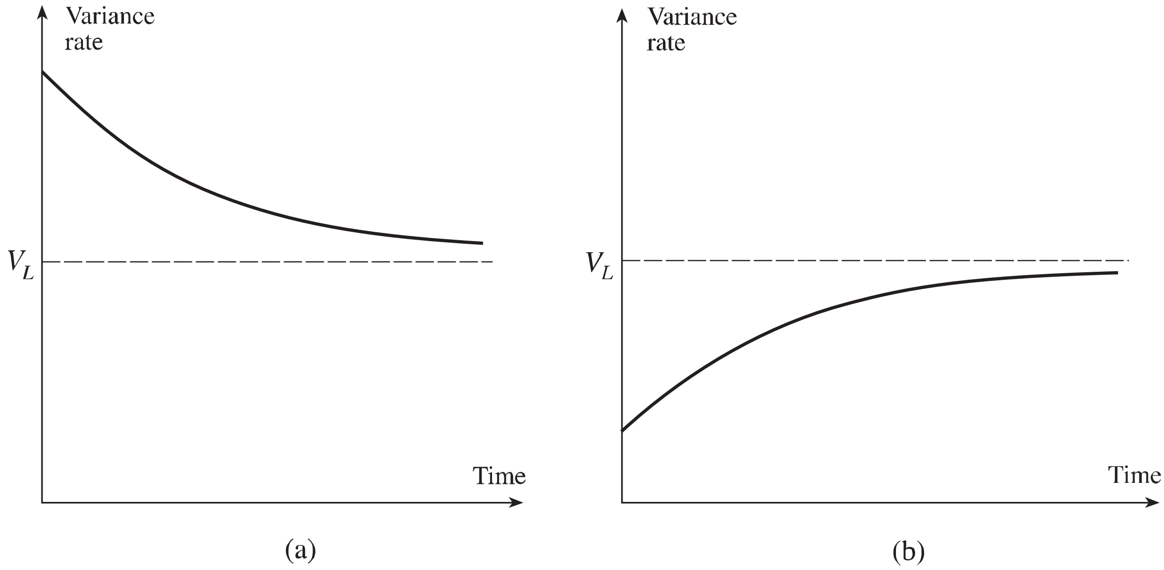

variance rate tends towards V_L as we look further and further ahead. This analysis emphasizes the point that we must have $\alpha + \beta < 1$ for a stable GARCH(1,1) process. When $\alpha + \beta > 1$, the weight given to the long-term average variance is negative and the process is "mean fleeing" rather than "mean reverting".

For the S&P 500 data considered earlier, $\alpha + \beta = 0.9935$ and $V_L = 0.0002075$. Suppose that the estimate of the current variance rate per day is 0.0003. (This corresponds to a volatility of 1.732% per day.) In 10 days, the expected variance rate is

$$0.0002075 + 0.9935^{10}(0.0003 - 0.0002075) = 0.0002942$$

The expected volatility per day is 1.72%, still well above the long-term volatility of 1.44% per day. However, the expected variance rate in 500 days is

$$0.0002075 + 0.9935^{500}(0.0003 - 0.0002075) = 0.0002110$$

and the expected volatility per day is 1.45%, very close to the long-term volatility.

Volatility Term Structures

Suppose it is day n. Define:

$$V(t) = E(\sigma_{n+t}^2)$$

and

$$a = \ln \frac{1}{\alpha + \beta}$$

so that equation (23.13) becomes

$$V(t) = V_L + e^{-at}[V(0) - V_L]$$

Here, $V(t)$ is an estimate of the instantaneous variance rate in t days. The average

Table 23.3 S&P 500 volatility term structure predicted from GARCH(1,1).

Option life (days)	10	30	50	100	500
Option volatility (% per annum)	27.36	27.10	26.87	26.35	24.32

variance rate per day between today and time T is given by

$$\frac{1}{T}\int_0^T V(t)\,dt = V_L + \frac{1-e^{-aT}}{aT}[V(0) - V_L]$$

The larger T is, the closer this is to V_L. Define $\sigma(T)$ as the volatility per annum that should be used to price a T-day option under GARCH(1,1). Assuming 252 days per year, $\sigma(T)^2$ is 252 times the average variance rate per day, so that

$$\sigma(T)^2 = 252\left(V_L + \frac{1-e^{-aT}}{aT}[V(0) - V_L]\right) \tag{23.14}$$

As discussed in Chapter 20, the market prices of different options on the same asset are often used to calculate a *volatility term structure*. This is the relationship between the implied volatilities of the options and their maturities. Equation (23.14) can be used to estimate a volatility term structure based on the GARCH(1,1) model. The estimated volatility term structure is not usually the same as the implied volatility term structure. However, as we will show, it is often used to predict the way that the implied volatility term structure will respond to volatility changes.

When the current volatility is above the long-term volatility, the GARCH(1,1) model estimates a downward-sloping volatility term structure. When the current volatility is below the long-term volatility, it estimates an upward-sloping volatility term structure. In the case of the S&P 500 data, $a = \ln(1/0.99351) = 0.006511$ and $V_L = 0.0002075$. Suppose that the current variance rate per day, $V(0)$, is estimated as 0.0003 per day. It follows from equation (23.14) that

$$\sigma(T)^2 = 252\left(0.0002075 + \frac{1-e^{-0.006511T}}{0.006511T}(0.0003 - 0.0002075)\right)$$

where T is measured in days. Table 23.3 shows the volatility per year for different values of T.

Impact of Volatility Changes

Equation (23.14) can be written

$$\sigma(T)^2 = 252\left[V_L + \frac{1-e^{-aT}}{aT}\left(\frac{\sigma(0)^2}{252} - V_L\right)\right]$$

When $\sigma(0)$ changes by $\Delta\sigma(0)$, $\sigma(T)$ changes by approximately

$$\frac{1-e^{-aT}}{aT}\,\frac{\sigma(0)}{\sigma(T)}\,\Delta\sigma(0) \tag{23.15}$$

Table 23.4 Impact of 1% change in the instantaneous volatility predicted from GARCH(1,1).

Option life (days)	10	30	50	100	500
Increase in volatility (%)	0.97	0.92	0.87	0.77	0.33

Table 23.4 shows the effect of a volatility change on options of varying maturities for the S&P 500 data considered above. We assume as before that $V(0) = 0.0003$, so that $\sigma(0) = \sqrt{252} \times \sqrt{0.0003} = 27.50\%$. The table considers a 100-basis-point change in the instantaneous volatility from 27.50% per year to 28.50% per year. This means that $\Delta\sigma(0) = 0.01$, or 1%.

Many financial institutions use analyses such as this when determining the exposure of their books to volatility changes. Rather than consider an across-the-board increase of 1% in implied volatilities when calculating vega, they relate the size of the volatility increase that is considered to the maturity of the option. Based on Table 23.4, a 0.97% volatility increase would be considered for a 10-day option, a 0.92% increase for a 30-day option, a 0.87% increase for a 50-day option, and so on.

23.7 CORRELATIONS

The discussion so far has centered on the estimation and forecasting of volatility. As explained in Chapter 22, correlations also play a key role in the calculation of VaR. In this section, we show how correlation estimates can be updated in a similar way to volatility estimates.

The correlation between two variables X and Y can be defined as

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

where σ_X and σ_Y are the standard deviations of X and Y and $\text{cov}(X, Y)$ is the covariance between X and Y. The covariance between X and Y is defined as

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

where μ_X and μ_Y are the means of X and Y, and E denotes the expected value. Although it is easier to develop intuition about the meaning of a correlation than it is for a covariance, it is covariances that are the fundamental variables of our analysis.[15]

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

$$x_i = \frac{X_i - X_{i-1}}{X_{i-1}}, \qquad y_i = \frac{Y_i - Y_{i-1}}{Y_{i-1}}$$

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

[15] An analogy here is that variance rates were the fundamental variables for the EWMA and GARCH procedures in the first part of this chapter, even though volatilities are easier to understand.

following:

$\sigma_{x,n}$: Daily volatility of variable X, estimated for day n

$\sigma_{y,n}$: Daily volatility of variable Y, estimated for day n

cov_n: Estimate of covariance between daily changes in X and Y, calculated on day n.

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

$$\frac{\text{cov}_n}{\sigma_{x,n}\,\sigma_{y,n}}$$

Using equal weighting and assuming that the means of x_i and y_i are zero, equation (23.3) shows that the variance rates of X and Y can be estimated from the most recent m observations as

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

A similar estimate for the covariance between X and Y is

$$\text{cov}_n = \frac{1}{m}\sum_{i=1}^{m} x_{n-i}\,y_{n-i} \tag{23.16}$$

One alternative for updating covariances is an EWMA model similar to equation (23.7). The formula for updating the covariance estimate is then

$$\text{cov}_n = \lambda\,\text{cov}_{n-1} + (1-\lambda)x_{n-1}\,y_{n-1}$$

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

Example 23.3

Suppose that $\lambda = 0.95$ and that the estimate of the correlation between two variables X and Y on day $n-1$ is 0.6. Suppose further 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 of the covariance between the 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:

$$\sigma_{x,n}^2 = 0.95 \times 0.01^2 + 0.05 \times 0.005^2 = 0.00009625$$
$$\sigma_{y,n}^2 = 0.95 \times 0.02^2 + 0.05 \times 0.025^2 = 0.00041125$$
$$\text{cov}_n = 0.95 \times 0.00012 + 0.05 \times 0.005 \times 0.025 = 0.00012025$$

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

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

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

$$\text{cov}_n = \omega + \alpha x_{n-1} y_{n-1} + \beta \, \text{cov}_{n-1}$$

and the long-term average covariance is $\omega/(1-\alpha-\beta)$. Formulas similar to those in equations (23.13) and (23.14) can be developed for forecasting future covariances and calculating the average covariance during the life of an option.[16]

Consistency Condition for Covariances

Once all the variances and covariances have been calculated, a variance–covariance matrix can be constructed. As explained in Section 22.4, when $i \neq j$, the (i, j)th element of this matrix shows the covariance between variable i and variable j. When $i = j$, 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

$$\boldsymbol{w}^{\mathsf{T}} \Omega \boldsymbol{w} \geqslant 0 \tag{23.17}$$

for all $N \times 1$ vectors $\boldsymbol{w}$, where $\boldsymbol{w}^{\mathsf{T}}$ is the transpose of $\boldsymbol{w}$. A matrix that satisfies this property is known as *positive-semidefinite*.

To understand why the condition in equation (23.17) must hold, suppose that $\boldsymbol{w}^{\mathsf{T}}$ is $[w_1, w_2, \dots, w_n]$. The expression $\boldsymbol{w}^{\mathsf{T}} \Omega \boldsymbol{w}$ is the variance of $w_1 x_1 + w_2 x_2 + \cdots + w_n x_n$, where x_i is the value of variable i. As such, it cannot be negative.

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 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 second variables. This seems strange. When $\boldsymbol{w}$ is set equal to $(1, 1, -1)$, the condition in equation (23.17) is not satisfied, proving that the matrix is not positive-semidefinite.[17]

---

[16] The ideas in this chapter can be extended to multivariate GARCH models, where an entire variance–covariance matrix is updated in a consistent way. For a discussion of alternative approaches, see R. Engle and J. Mezrich, "GARCH for Groups," *Risk*, August 1996: 36–40.

[17] It can be shown that the condition for a 3×3 matrix of correlations to be internally consistent is

$$\rho_{12}^2 + \rho_{13}^2 + \rho_{23}^2 - 2\rho_{12}\rho_{13}\rho_{23} \leqslant 1$$

where ρ_{ij} is the coefficient of correlation between variables i and j.

23.8 APPLICATION OF EWMA TO FOUR-INDEX EXAMPLE

We now return to the example considered in Section 22.2. This involved a portfolio on September 25, 2008, consisting of a \$4 million investment in the Dow Jones Industrial Average, a \$3 million investment in the FTSE 100, a \$1 million investment in the CAC 40, and a \$2 million investment in the Nikkei 225. Daily returns were collected over 500 days ending on September 25, 2008. Data and all calculations presented here can be found at: `www.rotman.utoronto.ca/~hull/OFOD/VaRExample`.

The correlation matrix that would be calculated on September 25, 2008, by giving equal weight to the last 500 returns is shown in Table 23.5. The FTSE 100 and CAC 40 are very highly correlated. The Dow Jones Industrial Average is moderately highly correlated with both the FTSE 100 and the CAC 40. The correlation of the Nikkei 225 with other indices is less high.

The covariance matrix for the equal-weight case is shown in Table 23.6. From equation (22.3), this matrix gives the variance of the portfolio losses (\$000s) as 8,761.833. The standard deviation is the square root of this, or 93.60. The one-day 99% VaR in \$000s is therefore $2.33 \times 93.60 = 217.757$. This is \$217,757, which compares with \$253,385, calculated using the historical simulation approach in Section 22.2.

Instead of calculating variances and covariances by giving equal weight to all observed returns, we now use the exponentially weighted moving average method with $\lambda = 0.94$. This gives the variance–covariance matrix in Table 23.7.[18] From equation (22.3), the

Table 23.5 Correlation matrix on September 25, 2008, calculated by giving equal weight to the last 500 daily returns: variable 1 is DJIA; variable 2 is FTSE 100; variable 3 is CAC 40; variable 4 is Nikkei 225.

$$\begin{bmatrix} 1 & 0.489 & 0.496 & -0.062 \\ 0.489 & 1 & 0.918 & 0.201 \\ 0.496 & 0.918 & 1 & 0.211 \\ -0.062 & 0.201 & 0.211 & 1 \end{bmatrix}$$

Table 23.6 Covariance matrix on September 25, 2008, calculated by giving equal weight to the last 500 daily returns: variable 1 is DJIA; variable 2 is FTSE 100; variable 3 is CAC 40; variable 4 is Nikkei 225.

$$\begin{bmatrix} 0.0001227 & 0.0000768 & 0.0000767 & -0.0000095 \\ 0.0000768 & 0.0002010 & 0.0001817 & 0.0000394 \\ 0.0000767 & 0.0001817 & 0.0001950 & 0.0000407 \\ -0.0000095 & 0.0000394 & 0.0000407 & 0.0001909 \end{bmatrix}$$

[18] In the EWMA calculations, the variance was initially set equal to the population variance. This is an alternative to setting it equal to the first squared return as in Table 23.1. The two approaches give similar final variances, and the final variance is all we are interested in.

Table 23.7 Covariance matrix on September 25, 2008, calculated using the EWMA method with $\lambda = 0.94$: variable 1 is DJIA; variable 2 is FTSE 100; variable 3 is CAC 40; variable 4 is Nikkei 225.

$$\begin{bmatrix} 0.0004801 & 0.0004303 & 0.0004257 & -0.0000396 \\ 0.0004303 & 0.0010314 & 0.0009630 & 0.0002095 \\ 0.0004257 & 0.0009630 & 0.0009535 & 0.0001681 \\ -0.0000396 & 0.0002095 & 0.0001681 & 0.0002541 \end{bmatrix}$$

variance of portfolio losses ($000s) is 40,995.765. The standard deviation is the square root of this, or 202.474. The one-day 99% VaR is therefore

$$2.33 \times 202.474 = 471.025$$

This is $471,025, over twice as high as the value given when returns are equally weighted. Tables 23.8 and 23.9 show the reasons. The standard deviation of a portfolio consisting of long positions in securities increases with the standard deviations of security returns and also with the correlations between security returns. Table 23.8 shows that the estimated daily standard deviations are much higher when EWMA is used than when data are equally weighted. This is because volatilities were much higher during the period immediately preceding September 25, 2008, than during the rest of the 500 days covered by the data. Comparing Table 23.9 with Table 23.5, we see that correlations had also increased.[19]

Table 23.8 Volatilities (% per day) using equal weighting and EWMA.

	DJIA	*FTSE 100*	*CAC 40*	*Nikkei 225*
Equal weighting:	1.11	1.42	1.40	1.38
EWMA:	2.19	3.21	3.09	1.59

Table 23.9 Correlation matrix on September 25, 2008, calculated using the EWMA method: variable 1 is DJIA; variable 2 is FTSE 100; variable 3 is CAC 40; variable 4 is Nikkei 225.

$$\begin{bmatrix} 1 & 0.611 & 0.629 & -0.113 \\ 0.611 & 1 & 0.971 & 0.409 \\ 0.629 & 0.971 & 1 & 0.342 \\ -0.113 & 0.409 & 0.342 & 1 \end{bmatrix}$$

[19] This is an example of the phenomenon that correlations tend to increase in adverse market conditions.

SUMMARY

Most popular option pricing models, such as Black–Scholes–Merton, assume that the volatility of the underlying asset is constant. This assumption is far from perfect. In practice, the volatility of an asset, like the asset's price, is a stochastic variable. Unlike the asset price, it is not directly observable. This chapter has discussed procedures for attempting to keep track of the current level of volatility.

We define u_i as the percentage change in a market variable between the end of day $i-1$ and the end of day i. The variance rate of the market variable (that is, the square of its volatility) is calculated as a weighted average of the u_i^2. The key feature of the procedures that have been discussed here is that they do not give equal weight to the observations on the u_i^2. The more recent an observation, the greater the weight assigned to it. In the EWMA and the GARCH(1,1) models, the weights assigned to observations decrease exponentially as the observations become older. The GARCH(1,1) model differs from the EWMA model in that some weight is also assigned to the long-run average variance rate. It has a structure that enables forecasts of the future level of variance rate to be produced relatively easily.

Maximum likelihood methods are usually used to estimate parameters from historical data in the EWMA, GARCH(1,1), and similar models. These methods involve using an iterative procedure to determine the parameter values that maximize the chance or likelihood that the historical data will occur. Once its parameters have been determined, a GARCH(1,1) model can be judged by how well it removes autocorrelation from the u_i^2.

For every model that is developed to track variances, there is a corresponding model that can be developed to track covariances. The procedures described here can therefore be used to update the complete variance–covariance matrix used in value at risk calculations.

FURTHER READING

Bollerslev, T. "Generalized Autoregressive Conditional Heteroscedasticity," *Journal of Econometrics*, 31 (1986): 307–27.

Cumby, R., S. Figlewski, and J. Hasbrook. "Forecasting Volatilities and Correlations with EGARCH Models," *Journal of Derivatives*, 1, 2 (Winter 1993): 51–63.

Engle, R. F. "Autoregressive Conditional Heteroscedasticity with Estimates of the Variance of UK Inflation," *Econometrica* 50 (1982): 987–1008.

Engle R. F., and J. Mezrich. "Grappling with GARCH," *Risk*, September 1995: 112–117.

Engle, R. F., and J. Mezrich, "GARCH for Groups," *Risk*, August 1996: 36–40.

Engle, R. F., and V. Ng, "Measuring and Testing the Impact of News on Volatility," *Journal of Finance*, 48 (1993): 1749–78.

Noh, J., R. F. Engle, and A. Kane. "Forecasting Volatility and Option Prices of the S&P 500 Index," *Journal of Derivatives*, 2 (1994): 17–30.

Practice Questions (Answers in Solutions Manual)

23.1. Explain the exponentially weighted moving average (EWMA) model for estimating volatility from historical data.

23.2. What is the difference between the exponentially weighted moving average model and the GARCH(1,1) model for updating volatilities?

23.3. The most recent estimate of the daily volatility of an asset is 1.5% and the price of the asset at the close of trading yesterday was \$30.00. The parameter λ in the EWMA model is 0.94. Suppose that the price of the asset at the close of trading today is \$30.50. How will this cause the volatility to be updated by the EWMA model?

23.4. A company uses an EWMA model for forecasting volatility. It decides to change the parameter λ from 0.95 to 0.85. Explain the likely impact on the forecasts.

23.5. The volatility of a certain market variable is 30% per annum. Calculate a 99% confidence interval for the size of the percentage daily change in the variable.

23.6. A company uses the GARCH(1,1) model for updating volatility. The three parameters are ω, α, and β. Describe the impact of making a small increase in each of the parameters while keeping the others fixed.

23.7. The most recent estimate of the daily volatility of the US dollar/sterling exchange rate is 0.6% and the exchange rate at 4 p.m. yesterday was 1.5000. The parameter λ in the EWMA model is 0.9. Suppose that the exchange rate at 4 p.m. today proves to be 1.4950. How would the estimate of the daily volatility be updated?

23.8. Assume that S&P 500 at close of trading yesterday was 1,040 and the daily volatility of the index was estimated as 1% per day at that time. The parameters in a GARCH(1,1) model are $\omega = 0.000002$, $\alpha = 0.06$, and $\beta = 0.92$. If the level of the index at close of trading today is 1,060, what is the new volatility estimate?

23.9. Suppose that the daily volatilities of asset A and asset B, calculated at the close of trading yesterday, are 1.6% and 2.5%, respectively. The prices of the assets at close of trading yesterday were \$20 and \$40 and the estimate of the coefficient of correlation between the returns on the two assets was 0.25. The parameter λ used in the EWMA model is 0.95.
(a) Calculate the current estimate of the covariance between the assets.
(b) On the assumption that the prices of the assets at close of trading today are \$20.5 and \$40.5, update the correlation estimate.

23.10. The parameters of a GARCH(1,1) model are estimated as $\omega = 0.000004$, $\alpha = 0.05$, and $\beta = 0.92$. What is the long-run average volatility and what is the equation describing the way that the variance rate reverts to its long-run average? If the current volatility is 20% per year, what is the expected volatility in 20 days?

23.11. Suppose that the current daily volatilities of asset X and asset Y are 1.0% and 1.2%, respectively. The prices of the assets at close of trading yesterday were \$30 and \$50 and the estimate of the coefficient of correlation between the returns on the two assets made at this time was 0.50. Correlations and volatilities are updated using a GARCH(1,1) model. The estimates of the model's parameters are $\alpha = 0.04$ and $\beta = 0.94$. For the correlation $\omega = 0.000001$, and for the volatilities $\omega = 0.000003$. If the prices of the two assets at close of trading today are \$31 and \$51, how is the correlation estimate updated?

23.12. Suppose that the daily volatility of the FTSE 100 stock index (measured in pounds sterling) is 1.8% and the daily volatility of the dollar/sterling exchange rate is 0.9%. Suppose further that the correlation between the FTSE 100 and the dollar/sterling exchange rate is 0.4. What is the volatility of the FTSE 100 when it is translated to US dollars? Assume that the dollar/sterling exchange rate is expressed as the number of

US dollars per pound sterling. (*Hint*: When $Z = XY$, the percentage daily change in Z is approximately equal to the percentage daily change in X plus the percentage daily change in Y.)

23.13. Suppose that in Problem 23.12 the correlation between the S&P 500 Index (measured in dollars) and the FTSE 100 Index (measured in sterling) is 0.7, the correlation between the S&P 500 Index (measured in dollars) and the dollar/sterling exchange rate is 0.3, and the daily volatility of the S&P 500 index is 1.6%. What is the correlation between the S&P 500 index (measured in dollars) and the FTSE 100 index when it is translated to dollars? (*Hint*: For three variables X, Y, and Z, the covariance between $X + Y$ and Z equals the covariance between X and Z plus the covariance between Y and Z.)

23.14. Show that the GARCH (1,1) model $\sigma_n^2 = \omega + \alpha u_{n-1}^2 + \beta\sigma_{n-1}^2$ in equation (23.9) is equivalent to the stochastic volatility model $dV = a(V_L - V)\,dt + \xi V\,dz$, where time is measured in days, V is the square of the volatility of the asset price, and

$$a = 1 - \alpha - \beta, \qquad V_L = \frac{\omega}{1 - \alpha - \beta}, \qquad \xi = \alpha\sqrt{2}$$

What is the stochastic volatility model when time is measured in years? (*Hint*: The variable u_{n-1} is the return on the asset price in time Δt. It can be assumed to be normally distributed with mean zero and standard deviation σ_{n-1}. It follows from the moments of the normal distribution that the mean and variance of u_{n-1}^2 are σ_{n-1}^2 and $2\sigma_{n-1}^4$, respectively.)

23.15. At the end of Section 23.8, the VaR for the four-index example was calculated using the model-building approach. How does the VaR calculated change if the investment is \$2.5 million in each index? Carry out calculations when (a) volatilities and correlations are estimated using the equally weighted model and (b) when they are estimated using the EWMA model with $\lambda = 0.94$. Use the spreadsheets on the author's website.

23.16. What is the effect of changing λ from 0.94 to 0.97 in the EWMA calculations in the four-index example at the end of Section 23.8. Use the spreadsheets on the author's website.

Further Questions

23.17. Suppose that the price of gold at close of trading yesterday was \$600 and its volatility was estimated as 1.3% per day. The price at the close of trading today is \$596. Update the volatility estimate using
(a) The EWMA model with $\lambda = 0.94$
(b) The GARCH(1,1) model with $\omega = 0.000002$, $\alpha = 0.04$, and $\beta = 0.94$.

23.18. Suppose that in Problem 23.17 the price of silver at the close of trading yesterday was \$16, its volatility was estimated as 1.5% per day, and its correlation with gold was estimated as 0.8. The price of silver at the close of trading today is unchanged at \$16. Update the volatility of silver and the correlation between silver and gold using the two models in Problem 23.17. In practice, is the ω parameter likely to be the same for gold and silver?

23.19. An Excel spreadsheet containing over 900 days of daily data on a number of different exchange rates and stock indices can be downloaded from the author's website:

`www.rotman.utoronto.ca/~hull/data.`

Choose one exchange rate and one stock index. Estimate the value of λ in the EWMA model that minimizes the value of $\sum_i (v_i - \beta_i)^2$, where v_i is the variance forecast made at the end of day $i-1$ and β_i is the variance calculated from data between day i and day $i+25$. Use the Solver tool in Excel. Set the variance forecast at the end of the first day equal to the square of the return on that day to start the EWMA calculations.

23.20. Suppose that the parameters in a GARCH (1,1) model are $\alpha = 0.03$, $\beta = 0.95$, and $\omega = 0.000002$.

(a) What is the long-run average volatility?

(b) If the current volatility is 1.5% per day, what is your estimate of the volatility in 20, 40, and 60 days?

(c) What volatility should be used to price 20-, 40-, and 60-day options?

(d) Suppose that there is an event that increases the current volatility by 0.5% to 2% per day. Estimate the effect on the volatility in 20, 40, and 60 days.

(e) Estimate by how much the event increases the volatilities used to price 20-, 40-, and 60-day options?

23.21. The calculations for the four-index example at the end of Section 23.8 assume that the investments in the DJIA, FTSE 100, CAC 40, and Nikkei 225 are \$4 million, \$3 million, \$1 million, and \$2 million, respectively. How does the VaR calculated change if the investments are \$3 million, \$3 million, \$1 million, and \$3 million, respectively? Carry out calculations when (a) volatilities and correlations are estimated using the equally weighted model and (b) when they are estimated using the EWMA model. What is the effect of changing λ from 0.94 to 0.90 in the EWMA calculations? Use the spreadsheets on the author's website.

23.22. Estimate parameters for EWMA and GARCH(1, 1) from data on the euro–USD exchange rate between July 27, 2005, and July 27, 2010. This data can be found on the author's website:

`www.rotman.utoronto.ca/~hull/data.`