Financial Econometrics
Lecture 3: ARCH & GARCH models

Richard G. Pierse

1 Introduction

One observed feature of many financial time series, such as returns on stocks or exchange rate differences, is that they have a non-constant variance. Such series typically exhibit periods of high volatility interspersed with periods of low volatility. The *ARCH* model of Engle (1982) and its *GARCH* generalisation by Bollerslev (1986) are simple models to capture such observed behaviour.

2 The *ARCH(q)* Model

Consider the model

$$y_t = \mathbf{x}_t'\beta + u_t \quad , \quad t = 1, \cdots, T \tag{2.1}$$

where $\mathbf{x}_t$ is a $k \times 1$ vector of regressors at time period t, and u_t is a disturbance process. The disturbance u_t has a constant mean

$$\mathrm{E}(u_t) = 0$$

but u_t^2 is time varying and follows an autoregressive process of order q

$$u_t^2 = \alpha_0 + \alpha_1 u_{t-1}^2 + \alpha_2 u_{t-2}^2 + \cdots + \alpha_q u_{t-q}^2 + \varepsilon_t \tag{2.2}$$

where ε_t is a disturbance with mean zero and unit variance, assumed to be distributed independently of u_t.

In order to ensure that $u_t^2 > 0$, we impose the restrictions that $\alpha_0 > 0$ and $\alpha_i \geq 0$, $i = 1, \cdots, q$. In addition, for the autoregressive process to be *weakly stationary* we require that the roots of

$$\alpha(L) = \alpha_1 L + \alpha_2 L^2 + \cdots + \alpha_q L^q$$

all lie *outside the unit circle*. This implies the condition that

$$\sum_{i=1}^{q} \alpha_i < 1.$$

The *conditional* variance of u_t, defined by $\mathrm{E}(u_t^2|u_{t-1}, u_{t-2}, \cdots)$, is given by

$$\mathrm{E}(u_t^2|u_{t-1}, u_{t-2}, \cdots) = \alpha_0 + \alpha_1 u_{t-1}^2 + \alpha_2 u_{t-2}^2 + \cdots + \alpha_q u_{t-q}^2 = \sigma_t^2. \qquad (2.3)$$

The *unconditional* variance, $\mathrm{E}(u_t^2)$, is constant over time and is given by

$$\mathrm{E}(u_t^2) = \frac{\alpha_0}{1 - \sum_{i=1}^{q} \alpha_i} = \sigma^2. \qquad (2.4)$$

This model, first suggested by Engle (1982), exhibits *AutoRegressive Conditional Heteroscedasticity* and is therefore known as the *ARCH(q)* model.

2.1 Leptokurtosis and *ARCH* models

Kurtosis is the scaled fourth moment of a probability distribution and is defined (for a random variable x_t with mean μ) by

$$\kappa(x_t) = \frac{\mathrm{E}(x_t - \mu)^4}{[\mathrm{E}(x_t - \mu)^2]^2} = \frac{\mu_4}{\mu_2^2}.$$

Kurtosis measures the fatness of the tails of a distribution, which is the probability of 'outliers'. For the normal distribution, $\kappa(u_t) = 3$. Distributions with fatter tails than the normal, like the t-distribution, have $\kappa(u_t) > 3$ and are said to be *leptokurtic*. Distributions with thinner tails than the normal have $\kappa(u_t) < 3$ and are said to be *platykurtic*. In empirical investigation of financial data series, many series show evidence of leptokurtosis. *ARCH* model disturbances have the property of being leptokurtic. For the case of the *ARCH(1)* model, the kurtosis of the *unconditional* distribution of u_t is given by

$$\kappa(u_t) = 3\frac{(1 - \alpha_1^2)}{(1 - 3\alpha_1^2)} > 3$$

for $\alpha_1 > 0$ and $3\alpha_1^2 < 1$ (otherwise the kurtosis is not finite). Thus *ARCH* models exhibit fatter tails than the normal distribution although this may not be enough fully to account for the kurtosis observed in real financial data series.

Kurtosis of Stock returns 6 January 1986 to 31 December 1997.
Source: Franses and van Dijk (2000)

Stock Market	*Daily*	*Weekly*
Amsterdam	19.795	11.929
Frankfurt	15.066	8.093
Hong Kong	119.241	18.258
London	27.408	15.548
New York	99.680	11.257
Paris	10.560	9.167
Singapore	28.146	23.509
Tokyo	14.798	4.897

2.2 Estimation of the *ARCH* model

The parameters of the *ARCH(q)* model can be estimated by the *maximum likelihood* method. Firstly, note that the joint log-likelihood function of all T observations of any model can always be written as the sum:

$$\log L(\mathbf{y};\theta) = \log L(y_1) + \sum_{t=2}^{T} \log L(y_t|y_{t-1}) \tag{2.5}$$

where $L(y_t|y_{t-1})$ is the *conditional density* of y_t given y_{t-1} and $L(y_1)$ is the *marginal density* of the first observation, y_1, which in practice is often dropped.

In the case of the *ARCH(q)* model, the conditional density of y_t given y_{t-1} has mean $\mathbf{x}_t'\beta$ and variance given by the conditional variance of u_t, which is σ_t^2. If we assume a normal density for $y_t|y_{t-1}$, then we can write

$$\begin{aligned} L(y_t|y_{t-1}) &\sim N(\mathbf{x}_t'\beta,\sigma_t^2) \\ &= -\frac{1}{2}\log 2\pi - \frac{1}{2}\log\sigma_t^2 - \frac{1}{2\sigma_t^2}(y_t - \mathbf{x}_t'\beta)^2. \end{aligned}$$

and the joint log-likelihood function is defined by

$$\log L(\mathbf{y};\beta,\alpha) = \log L(y_1) - \frac{T-1}{2}\log 2\pi - \frac{1}{2}\sum_{t=2}^{T}\log\sigma_t^2 - \frac{1}{2}\sum_{t=2}^{T}\frac{(y_t - \mathbf{x}_t'\beta)^2}{\sigma_t^2} \tag{2.6}$$

where

$$\sigma_t^2 = \alpha_0 + \sum_{i=1}^{q}\alpha_i(y_{t-i} - \mathbf{x}_{t-i}'\beta)^2.$$

Maximising this log-likelihood function with respect to the unknown parameters, β, α_0, α_i, $i = 1,\cdots,q$, defines the maximum likelihood estimators, $\widetilde{\beta}$, $\widetilde{\alpha}_0$, $\widetilde{\alpha}_i$. No explicit expression for these estimators is possible but the estimators can be found by numerical maximisation of (2.6). Note that the first q observations of the sample must be dropped in order to define σ_t^2. In general, some restrictions on the likelihood function may also need to be imposed to ensure that the estimators satisfy the conditions that $\widetilde{\alpha}_0 > 0$ and $\widetilde{\alpha}_i \geq 0$, $i = 1, \cdots, q$.

It is also possible to choose a non-normal conditional density function for $y_t|y_{t-1}$. One alternative choice which is sometimes used is the *t-distribution*. This has fatter tails than the normal distribution, with kurtosis $\kappa = 3+6/(\lambda-4)$ where λ is the degrees of freedom parameter. Use of the t-distribution is consistent with the excess kurtosis often found in financial data. Maximum likelihood estimation with the t-distribution proceeds as before, using (2.5) to define the joint density function for all observations and maximising this with respect to the unknown parameters.

3 The *GARCH(p,q)* model

Bollerslev (1986) proposed a *generalised ARCH* or *GARCH* model, in which the conditional variance of u_t, σ_t^2, depends not only on lagged u_t^2 but also on lags of σ_t^2 itself. The *GARCH(p,q)* model is defined by

$$\sigma_t^2 = \alpha_0 + \alpha_1 u_{t-1}^2 + \alpha_2 u_{t-2}^2 + \cdots + \alpha_q u_{t-q}^2 + \beta_1 \sigma_{t-1}^2 + \beta_2 \sigma_{t-2}^2 + \cdots + \beta_p \sigma_{t-p}^2 \quad (3.1)$$

or

$$\beta(L)\sigma_t^2 = \alpha_0 + \alpha(L)u_t^2$$

where

$$\beta(L) = 1 - \beta_1 L - \beta_2 L^2 - \cdots - \beta_p L^p.$$

This is a form of *ARMA* process for u_t^2 as can be seen by adding u_t^2 to each side of (3.1) and rewriting the equation as

$$\begin{aligned} u_t^2 &= \alpha_0 + (\alpha_1 + \beta_1)u_{t-1}^2 + (\alpha_2 + \beta_2)u_{t-2}^2 + \cdots + (\alpha_m + \beta_m)u_{t-m}^2 \\ &\quad + (u_t^2 - \sigma_t^2) - \beta_1(u_{t-1}^2 - \sigma_{t-1}^2) - \beta_2(u_{t-2}^2 - \sigma_{t-2}^2) - \cdots - \beta_p(u_{t-p}^2 - \sigma_{t-p}^2) \end{aligned}$$

where $m = \max(p, q)$ and $\alpha_j = 0$, $j > q$, and $\beta_j = 0$, $j > p$. Defining $\varepsilon_t = u_t^2 - \sigma_t^2$, this is

$$\begin{aligned} u_t^2 &= \alpha_0 + (\alpha_1 + \beta_1)u_{t-1}^2 + (\alpha_2 + \beta_2)u_{t-2}^2 + \cdots + (\alpha_m + \beta_m)u_{t-m}^2 \\ &\quad + \varepsilon_t - \beta_1 \varepsilon_{t-1} - \beta_2 \varepsilon_{t-2} - \cdots - \beta_p \varepsilon_{t-p} \end{aligned}$$

which is an *ARMA(m,p)* process for u_t^2. Thus the *GARCH(p,q)* model is an *ARMA(m,p)* process and *GARCH(0,q)* coincides with the pure *AR* process *ARCH(q)*. The advantage of a *GARCH* process over a pure *ARCH* process is *parsimony.* A *GARCH* model can capture complicated patterns of time variability in the conditional variance using fewer parameters than an *ARCH* model. In empirical applications, the *GARCH(1,1)* model has often proved sufficiently general to capture the heteroscedasticity in financial time series.

For weak stationarity of σ_t^2, we require that

$$\sum_{i=1}^{q} \alpha_i + \sum_{j=1}^{p} \beta_j < 1.$$

In that case the unconditional variance σ^2 is given by

$$\begin{aligned} \sigma^2 &= \frac{\alpha_0}{1 - \sum_{i=1}^{q} \alpha_i - \sum_{j=1}^{p} \beta_j} \\ &= \frac{\alpha_0}{\beta(1) - \alpha(1)}. \end{aligned}$$

To ensure that σ_t^2 and σ^2 are always positive, a *sufficient* condition is that $\alpha_0 \geq 0$ and that $\alpha_i \geq 0$ and $\beta_j \geq 0$, $\forall i, j$.

3.1 Testing for *(G)ARCH* effects

When *ARCH* effects are suspected, we can consider testing the null hypothesis of *homoscedasticity*

$$\sigma_t^2 = \sigma^2 = \alpha_0$$

against the alternative hypothesis of an *ARCH(q)* model:

$$\sigma_t^2 = \alpha_0 + \alpha_1 u_{t-1}^2 + \alpha_2 u_{t-2}^2 + \cdots + \alpha_q u_{t-q}^2.$$

On the null hypothesis, *OLS* is an appropriate estimator and so a Lagrange Multiplier diagnostic test can be constructed, based on the residuals, $\widehat{u}_t$, from the *OLS* regression of equation (2.1). The test is based on the R^2 from the following regression of the squared *OLS* residuals:

$$\widehat{u}_t^2 = \widehat{\alpha}_0 + \widehat{\alpha}_1 \widehat{u}_{t-1}^2 + \cdots + \widehat{\alpha}_q \widehat{u}_{t-q}^2$$

and Engle (1982) showed that, on the null hypothesis,

$$TR^2 \sim_a \chi_q^2.$$

Testing the null of *homoscedasticity* against the more general alternative of a *GARCH(p,q)* model is less straightforward and, in fact, no test exists for the general case. However, tests for some special cases are available. In particular, a test of *homoscedasticity* against the *GARCH(p,0)* model does exist and coincides with the *ARCH* test just considered, for the alternative hypothesis *ARCH(p)*.

References

[1] Bollerslev, T. (1986), 'Generalised autoregressive conditional heteroskedasticity', *Journal of Econometrics*, 32, 307–327.

[2] Engle, R.F. (1982), 'Autoregressive conditional heteroskedasticity with estimates of the variance of UK inflation', *Econometrica*, 50, 987–1008.

[3] Engle, R.F. (1995) (ed.) *ARCH Selected Readings*, Oxford University Press, Oxford, UK.

[4] Franses, P.H. and D. van Dijk (2000), *Non-linear time series models in empirical finance*, Cambridge University Press, Cambridge, UK.