

Time series: Variance modelling

Bernt Arne Ødegaard

24 November 2021

Contents	
1 Motivation	1
1.1 Variance clustering	1
1.2 Relation to heteroskedasticity	2
1.3 Second motivation: implied volatility	3
1.4 Uses of ARCH processes	4
1.5 Testing for ARCH	4
1.6 Estimation of ARCH type processes	5
1.7 Extensions to ARCH	7
2 Summary	7
3 Formal definitions	8
4 Literature	9

Introduction

The simplest assumption for time series is that variance is constant. Unfortunately that assumption is often violated in actual data.

In this lecture we look at the implications of time varying conditional variance, or ARCH (Autoregressive Conditional Heteroscedasticity).

1 Motivation

Two motivations, volatility clustering, and implied volatility

1.1 Variance clustering


Exercise 1.

Collect data for the S&P 500. Use data between 1990 and 2000.

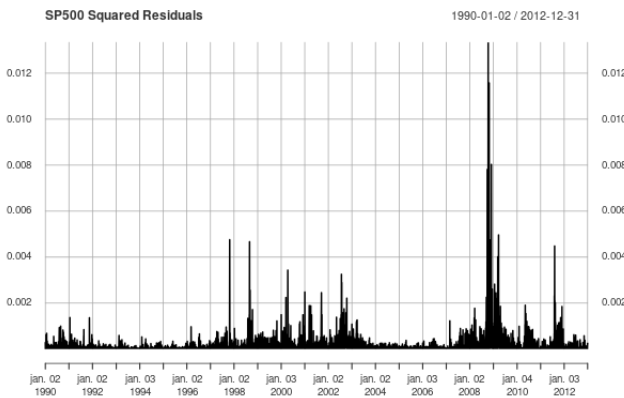
1. Calculate daily returns $R_{m,t}$. Plot the time series of returns.
2. Calculate the squared daily returns $R_{m,t}^2$. Plot the time series of squared returns returns.
3. Are there signs of “volatility clustering”?

Solution to Exercise 1.

1. Plotting daily returns



2. Plotting squared daily returns.



3. Volatility is the standard deviations of returns. Under a constant mean this is estimated as (the square root of) the squared return.

Clustering is when similar volatility patterns are “close” in time, sometimes low volatility, other times high volatility. From the pictures it certainly looks like there is some volatility clustering.

The example of S&P 500 daily returns and their squares is a clear illustration of “variance clustering”

1.2 Relation to heteroskedasticity

Relative to what we have seen before this is heteroskedasticity in the errors, when considering the time series.

Simply put, due to the time series nature of the data, the heteroskedasticity is related to time.

If we could identify the structure of the variance-covariance matrix Ω , we would use it in a GLS estimation

$$\hat{\mathbf{b}}^{gls} = (\mathbf{X}'\Omega^{-1}\mathbf{X})^{-1}\mathbf{X}'\Omega^{-1}\mathbf{y}$$

The challenge is to find a way of modelling the dependence between time and volatility.

One obvious way to start is to model the time t variance σ_t^2 as an autoregressive processes, using past values of σ_t^2 to predict the next. Suppose we have a time series estimation

$$Y_t = \beta_0 + \beta_1 Y_{t-1} + \gamma_1 X_{t-1} + u_t$$

Two well known examples: $ARCH(p)$ (AutoRegressive Conditional Heteroskedasticity of order p) (Engle, 1982)

$$\sigma_t^2 = \sigma_0 + \alpha_1 u_{t-1}^2 + \alpha_2 u_{t-2}^2 + \dots + \alpha_p u_{t-p}^2$$

and $GARCH(p, q)$ (Generalized AutoRegressive Conditional Heteroskedasticity of order (p, q) .) (Bollerslev, 1986)

$$\sigma_t^2 = \sigma_0 + \alpha_1 u_{t-1}^2 + \alpha_2 u_{t-2}^2 + \dots + \alpha_p u_{t-p}^2 + \phi_1 \sigma_{t-1}^2 + \phi_2 \sigma_{t-2}^2 + \dots + \phi_q \sigma_{t-q}^2$$

Note that an ARCH is analogous to a distributed lag model, and an GARCH is analogous to an ADL model.

This is the econometric motivation for working with ARCH/GARCH processes, to alleviate problems of heteroskedasticity in estimation. Here the (G)ARCH estimation is just nuisance parameters that help us improving precision of estimates.

However, predictability of variances (and covariances) is something that is of interest in itself, in particular in finance applications, where so much relies on these measures of variability.

1.3 Second motivation: implied volatility

Let us look at another motivation for why (G)ARCH is important for finance, and look at options pricing.

Exercise 2.

Consider pricing of options. The price of any option is a function of the following five inputs:

- Current price of underlying.
- Exercise price of the option.
- Current risk free interest rate.
- Time to maturity of the option.
- Volatility of the underlying.

$$\text{OptPrice}_t = f(S_t, K, r_t, (T - t), \sigma_t)$$

All of these except the volatility of the underlying are observables.


Given an observed market price of the option, this equation can be reversed, by asking what volatility would have given the observed option market price.

This is called the “implied volatility”.

Calculate the implied volatility of options on the Norwegian OBX stock market index in the period July 2001 to November 2001, and plot the time series.

Does the implied volatility change over time?

Solution to Exercise 2.



The implied volatility is after the fact, what one ideally want to look at is ways of predicting the volatility, and use it as inputs to option pricing.

1.4 Uses of ARCH processes

The previous points out two the two most important uses of ARCH type processes.

- Correcting for ARCH effects in errors to improve estimation (ARCH as nuisance parameters)
- Using ARCH models to forecast variance.

1.5 Testing for ARCH

ARCH type of effects involves variance as a function of time. Testing for ARCH is then mainly a matter of running time series models on the estimated errors, and testing whether we find predictability.

Exercise 3.

ARCH is time varying conditional variance. We want to test for its presence in the returns r_m of a stock market index.

The null hypothesis of no ARCH effects involve a constant mean and variance

$$r_{mt} \sim (E[r_m], \sigma_m^2)$$

Estimating the mean and variance, testing for ARCH effects can be done using the residuals

$$u_t = r_{mt} - \widehat{E[r_m]}$$

and running the regression

$$\hat{u}_t^2 = b_0 + \sum_{i=1}^p b_i \hat{u}_{t-i}^2 + \text{residual},$$

where p is the number of lags in the ARCH/GARCH formulation.

Under a null of no ARCH effects all the coefficients b_i for $i > 0$ are jointly zero. This is the basis for testing for ARCH effects. Under the null T times the R^2 of the regression has a $\chi^2(p)$ distribution. (T is the number of observations). Alternatively, the usual F test can be applied.

Perform the test on the S&P 500 index using daily data for the period 1990–2011. Use $p = 2$.

Solution to Exercise 3.

```
> library(zoo)
> SP500 <- read.zoo("../data/usa/us_equity_indices/sp500_daily.csv",
  format="%Y-%m-%d", sep="," ,header=TRUE)
> sp500 <- SP500[,6]
> # the post 1990 period
> sp500 <- sp500[10055:length(sp500),]
> dsp500 <- diff(log(sp500))
> u <- dsp500 - mean(dsp500)
> u2 <- u^2
> data <- merge(u2, lag(u2, -1), lag(u2, -2), all=FALSE)
> u2 <- data[,1]
> u2lag1 <- data[,2]
> u2lag2 <- data[,3]
> arch <- lm(u2~u2lag1+u2lag2)
> summary(arch)
```

Call:

```
lm(formula = u2 ~ u2lag1 + u2lag2)
```

Residuals:

Min	1Q	Median	3Q	Max
-0.0032956	-0.0000951	-0.0000700	-0.0000029	0.0098092

```

Coefficients:
              Estimate Std. Error t value Pr(>|t|)
(Intercept) 7.372e-05  6.017e-06  12.25  <2e-16 ***
u2lag1      1.396e-01  1.263e-02  11.05  <2e-16 ***
u2lag2      3.348e-01  1.263e-02  26.50  <2e-16 ***
---
Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1

Residual standard error: 0.0004171 on 5564 degrees of freedom
Multiple R-squared:  0.1512, Adjusted R-squared:  0.1509
F-statistic: 495.5 on 2 and 5564 DF,  p-value: < 2.2e-16

```

1.6 Estimation of ARCH type processes

Note that both u_t and σ_t are unobservable. Estimation of ARCH type models is therefore subject to the same problems as estimation of MA parameters.

For now, let me just state that ARCH type effects can be estimated using Maximum Likelihood and/or GMM, Usually we just leave it to the computer.

Exercise 4.

Consider daily returns on the S&P 500. Testing shows signs of ARCH effects. Using data for 1995–2006, estimate five different specifications: ARCH(1), ARCH(2), GARCH(1,1), GARCH(2,2).

What are the different parameters?

Solution to Exercise 4.

Estimation of ARCH(1)

```

Call:
garch(x = as.matrix(dsp500), order = c(0, 1))

Model:
GARCH(0,1)

Residuals:
      Min       1Q   Median       3Q      Max
-9.04915 -0.44568  0.04755  0.51239  9.16175

Coefficient(s):
      Estimate Std. Error t value Pr(>|t|)
a0 1.009e-04  1.263e-06  79.91  <2e-16 ***
a1 3.008e-01  1.396e-02  21.54  <2e-16 ***
---
Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1

Diagnostic Tests:
Jarque Bera Test

data: Residuals
X-squared = 14046.71, df = 2, p-value < 2.2e-16

Box-Ljung test

data: Squared.Residuals
X-squared = 3.1029, df = 1, p-value = 0.07815

```

Estimation of ARCH(2)

Model:

GARCH(0,2)

Residuals:

Min	1Q	Median	3Q	Max
-6.51078	-0.47933	0.05438	0.55421	6.67325

Coefficient(s):

	Estimate	Std. Error	t value	Pr(> t)
a0	6.631e-05	1.388e-06	47.76	<2e-16 ***
a1	1.844e-01	1.345e-02	13.71	<2e-16 ***
a2	3.486e-01	1.473e-02	23.66	<2e-16 ***

Signif. codes: 0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1

Diagnostic Tests:

Jarque Bera Test

data: Residuals

X-squared = 2170.128, df = 2, p-value < 2.2e-16

Box-Ljung test

data: Squared.Residuals

X-squared = 3.2871, df = 1, p-value = 0.06983

Estimation of GARCH(1,1)

Model:

GARCH(1,1)

Residuals:

Min	1Q	Median	3Q	Max
-6.85962	-0.52642	0.06148	0.62802	4.23953

Coefficient(s):

	Estimate	Std. Error	t value	Pr(> t)
a0	9.432e-07	1.243e-07	7.588	3.24e-14 ***
a1	7.300e-02	4.474e-03	16.315	< 2e-16 ***
b1	9.202e-01	4.980e-03	184.789	< 2e-16 ***

Signif. codes: 0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1

Diagnostic Tests:

Jarque Bera Test

data: Residuals

X-squared = 956.2433, df = 2, p-value < 2.2e-16

Box-Ljung test

```
data: Squared.Residuals
X-squared = 3.3017, df = 1, p-value = 0.06921
```

Garch (2,2)

```
> summary(res)
```

```
Call:
garch(x = as.matrix(dsp500), order = c(2, 2))
```

```
Model:
GARCH(2,2)
```

```
Residuals:
      Min       1Q   Median       3Q      Max
-6.74479 -0.52896  0.06152  0.63680  4.29440
```

```
Coefficient(s):
      Estimate Std. Error t value Pr(>|t|)
a0 1.796e-06      NA      NA      NA
a1 2.943e-02      NA      NA      NA
a2 1.011e-01      NA      NA      NA
b1 2.833e-01      NA      NA      NA
b2 5.728e-01      NA      NA      NA
```

```
Diagnostic Tests:
Jarque Bera Test
```

```
data: Residuals
X-squared = 835.0642, df = 2, p-value < 2.2e-16
```

Box-Ljung test

```
data: Squared.Residuals
X-squared = 0.4195, df = 1, p-value = 0.5172
```

1.7 Extensions to ARCH

There is a whole alphabet soup of extensions to the basic ARCH/GARCH models. Most of them are motivated by a better fit to observed features of data series

For example, the exponential GARCH models the log of the variance instead of the variance, which guarantees that the estimated variance is positive, and also a formulation that allows for asymmetry, negative shocks may have different effects from positive shocks.

Another example is a GARCH-in-mean, where the GARCH process is allowed to affect the mean of the underlying variable.

2 Summary

Modelling volatility

Empirical observations:

- Volatility clustering.

- Time series of implied volatility.

ARCH(p)

$$\sigma_t^2 = \sigma_0 + \alpha_1 u_{t-1}^2 + \alpha_2 u_{t-2}^2 + \cdots + \alpha_p u_{t-p}^2$$

GARCH(p, q)

$$\sigma_t^2 = \sigma_0 + \alpha_1 u_{t-1}^2 + \alpha_2 u_{t-2}^2 + \cdots + \alpha_p u_{t-p}^2 + \phi_1 \sigma_{t-1}^2 + \phi_2 \sigma_{t-2}^2 + \cdots + \phi_p \sigma_{t-p}^2$$

Uses of ARCH processes.

- Correcting for ARCH effects in errors to improve estimation (ARCH as nuisance parameters)
- Using ARCH models to forecast variance.

Extensions to ARCH

3 Formal definitions

Generic relation

$$\epsilon_t = z_t \sigma_t$$

$$z_t iid$$

$$E[z_t] = 0$$

$$\text{var}(z_t) = 1$$

$$y_t = g(x_{t-1}; b) + \epsilon_t$$

Estimation by ML

$$\epsilon_t \sim f(\cdot)$$

log likelihood function

$$L(\theta) = \sum_{t=1}^T [\log f(\epsilon_t \sigma_t^{-1}) - \log \sigma_t]$$

Since $\epsilon_t = y_t - g(x_{t-1}; b)$, estimate by

$$\hat{b} = \arg \max_b \sum_{t=1}^T \log f((y_t - g(x_{t-1}; b)) \sigma_t^{-1}) - \log \sigma_t$$

Typical representations, univariate

ARCH(q):

$$\sigma_t^2 = \omega + \sum_{i=1}^q \alpha_i \epsilon_{t-i}^2$$

GARCH(p, q):

$$\sigma_t^2 = \omega + \sum_{i=1}^q \alpha_i \epsilon_{t-i}^2 + \sum_{i=1}^p \beta_i \sigma_{t-i}^2$$

EGARCH(p, q):

$$\log \sigma_t^2 = \omega + \sum_{i=1}^q \alpha_i (\tau z_{t-i} + \gamma [|z_{t-i}| - E[z_{t-i}]] + \sum_{i=1}^p \beta_i \log \sigma_{t-i}^2$$

Of course also vector representations.

$$\epsilon_t = z_t \Omega_t^{\frac{1}{2}}$$

where Ω is a covariance matrix
vector GARCH(p,q)

$$\text{vech}(\Omega_t) = W + \sum_{i=1}^q A_i \text{vech}(\epsilon_{t-i} \epsilon'_{t-i}) + \sum_{i=1}^p B_i \text{vech}(\Omega_{t-i})$$

W : $\frac{N(N+1)}{2} \times 1$ matrix

A,B : $\frac{N(N+1)}{2} \times \frac{N(N+1)}{2}$ matrix

4 Literature

For a survey of volatility modelling, including ARCH, see Andersen, Bollerslev, and Diebold (2010)

References

Torben G Andersen, Tim Bollerslev, and Francis X Diebold. Parametric and nonparametric volatility estimation. In Yacine Ait-Sahalia and Lars Peter Hansen, editors, *Handbook of Financial Econometrics*. North-Holland, 2010.

Tim Bollerslev. Generalized autoregressive conditional heteroskedasticity. *Journal of Business*, pages 307–327, 1986.

Robert F Engle. Autoregressive conditional heteroscedasticity with estimates of the variance of United Kingdom inflation. *Econometrica*, 50:987–1007, 1982.