Financial Econometrics and Volatility Models
Univariate GARCH Models

Eric Zivot
Updated: April 5, 2010

Introduction to ARCH and GARCH Models

ARCH (AutoRegressive Conditional Heteroskedasticity) models were proposed by Engle in 1982.

GARCH (Generalized ARCH) models proposed by Bollerslev in 1986.

Engle received the Nobel price in 2003. The GARCH model framework is considered as one of the most important contributions in empirical finance over the last 20 years.

Robert Engle, NYU

Nobel Prize citation: "for methods of analyzing economic time series with time-varying volatility (ARCH)"

Champion Pairs Skater Too!

Other Acronyms for ARCH

Anything Can Really Happen

Another Risk Can't Hurt

All Reality Comes Here

Applied Research Can Help

Another Rather Crazy Hypothesis

Almost Right Conjected Heuristic

Almost Real Crummy Hogwash

Applied Research on Changing Histories

A Real Chinese Hero

A Rice Cake History

A Real Clever Heuristic

Another Really Cute Hunch

Ahh, Robert...Call Home!

A Really Cool Hat

And Robert Can Hit

All Risks Compensate Highly

The ARCH Family

Bolerslev (2008) identified over 150 different ARCH models. Here are some of the most common:

- GJR-GARCH
- TARCH
- STARCH
- AARCH
- NARCH
- MARCH
- SWARCH
- SNPARCH
- APARCH
- TAYLOR-SCHWERT
- FIGARCH
- FIEGARCH
- Component
- Asymmetric Component
- SQGARCH
- CESGARCH
- Student t
- GED
- SPARCH

Stylized Facts of Daily Asset Returns

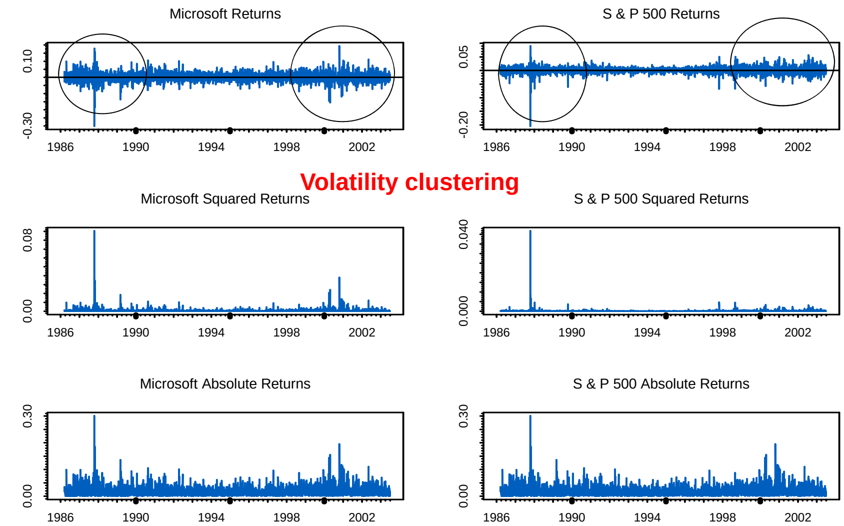

Sample Autocorrelations of Daily Returns

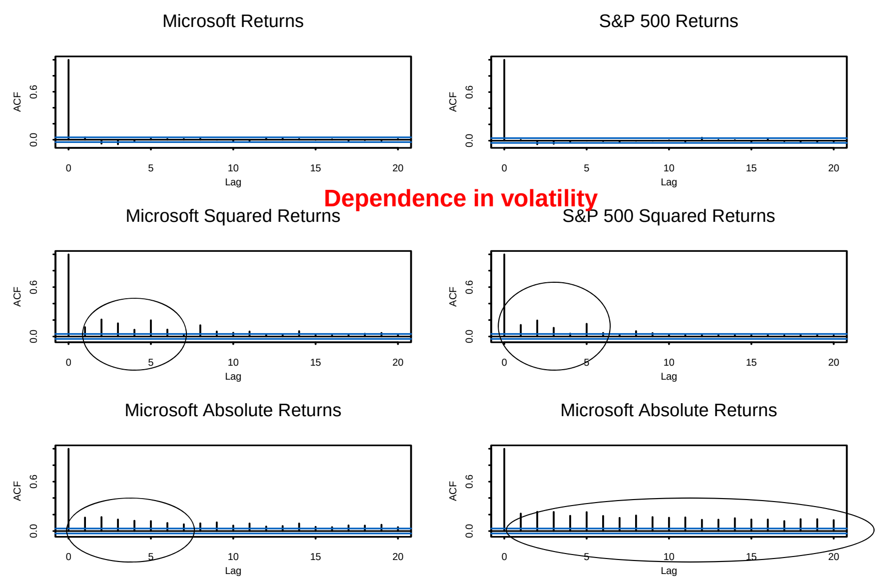

Stylized Facts for Monthly Asset Returns

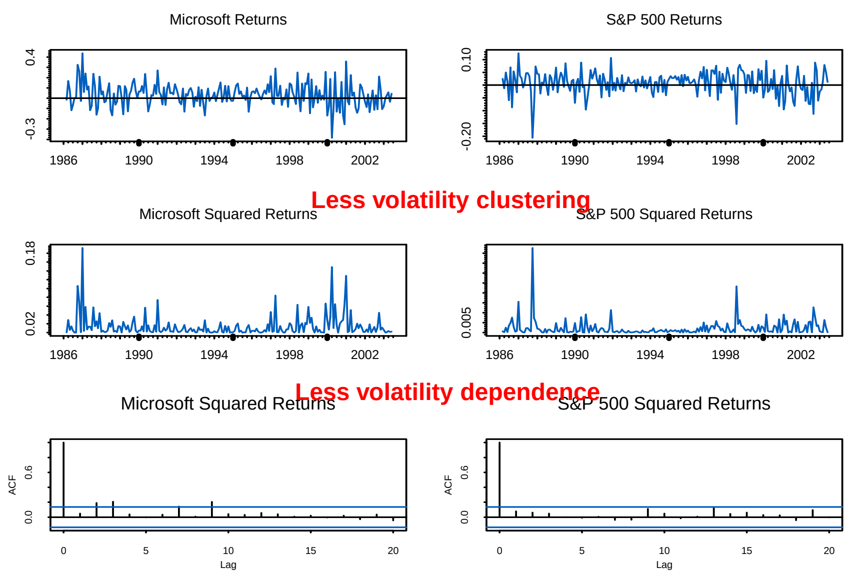

Summary statistics

```
> summaryStats(msft.ts)
> summaryStats(sp500.ts)
> normalTest(msft.ts, method="jb")
> normalTest(sp500.ts, method="jb")
```

Highly non-normal data

Asset	Mean	Med	Min	Max	Std. Dev	Skew	Kurt	JB
				Daily Returns				
MSFT	0.0016	0.0000	-0.3012	0.1957	0.0253	-0.2457	11.66	13693
S&P 500	0.0004	0.0005	-0.2047	0.0909	0.0113	-1.486	32.59	160848
				Monthly Returns				
MSFT	0.0336	0.0336	-0.3861	0.4384	0.1145	0.1845	4.004	9.922
S&P 500	0.0082	0.0122	-0.2066	0.1250	0.0459	-0.8377	5.186	65.75

Notes: Sample period is 03/14/86 - 06/30/03 giving 4365 daily observations.

Simulating an ARCH(1) Model

```
> sim.arch1 <-
+   simulate.garch(model=list(a.value=1,arch=0.8),
+                   n=500, rseed=124)
> names(sim.arch1)
[1] "et"      "sigma.t"
```

c = 0 (by default), a.value = a_0, arch = a_1, $e_t = \varepsilon_t$ and sigma.t = σ_t (conditional standard veviation).

Simulating an ARCH(1) Model

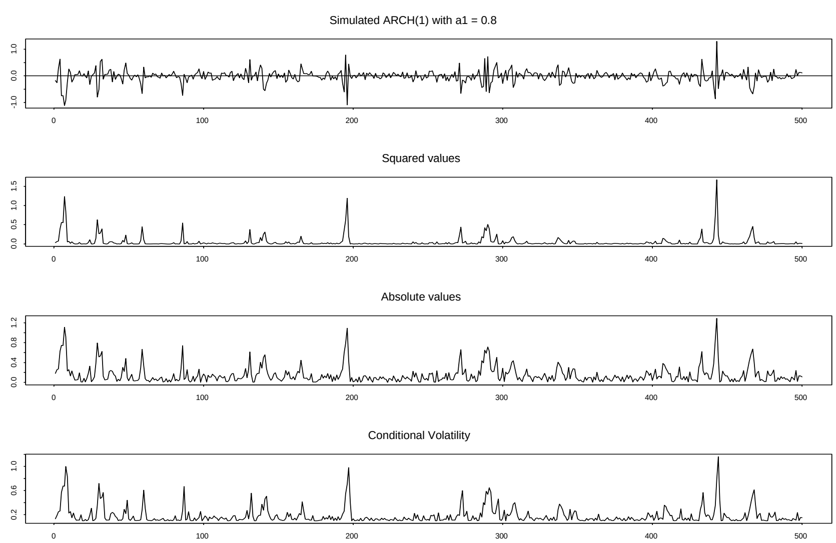

Sample Autocorrelations from ARCH(1)

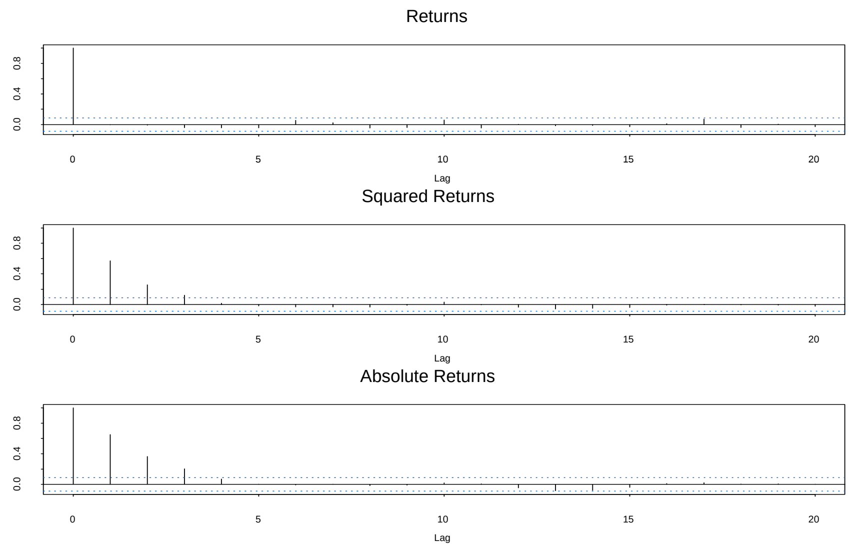

Distribution of Simulated ARCH(1) Data

Sample Quantiles:

min	1Q	median	3Q	max
-1.11	-0.1068	-0.006157	0.0915	1.292

Sample Moments:

mean	std	skewness	kurtosis
-0.0168	0.229	-0.3925	8.732

Number of Observations: 500

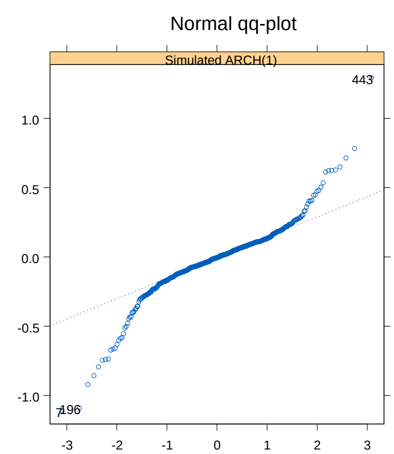

Simulated GARCH(1,1) Data

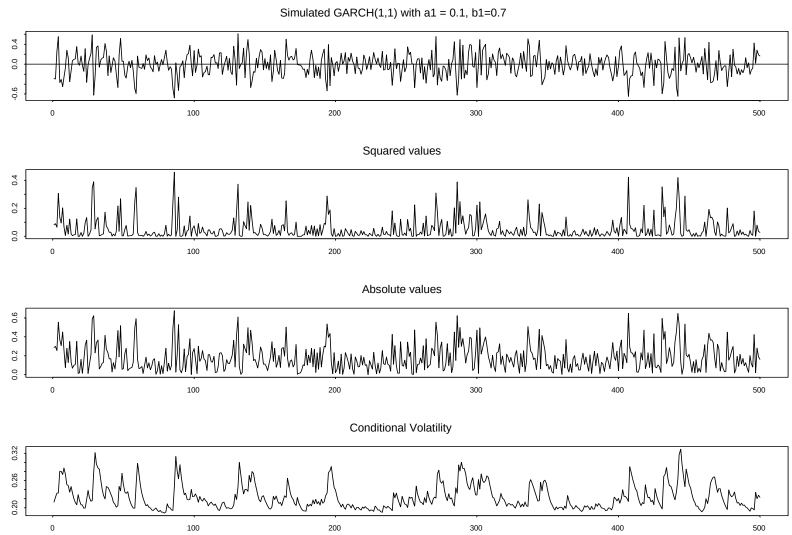

Sample Autocorrelations from GARCH(1,1)

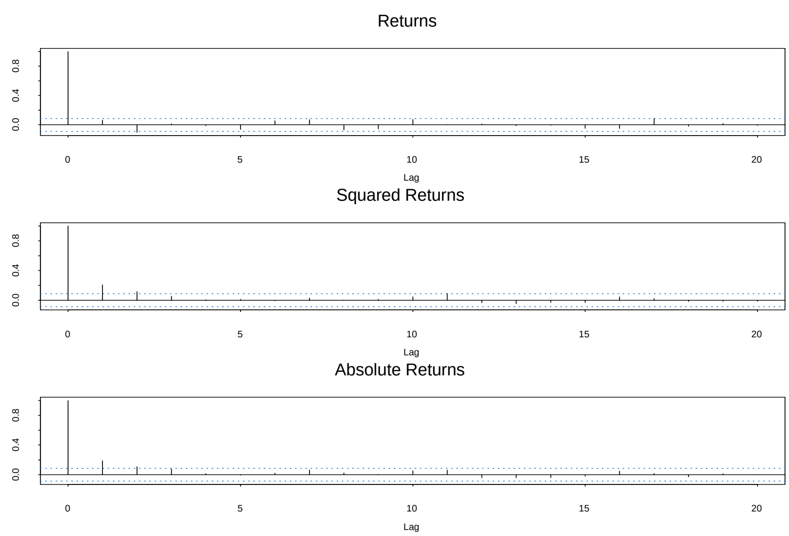

Distribution of Simulated GARCH(1,1) Data

```
Sample Quantiles:
     min      1Q    median     3Q    max
 -0.6763 -0.1587 -0.01016 0.1342 0.6103

Sample Moments:
     mean    std skewness kurtosis
 -0.01029 0.2296 -0.05683    3.102

Number of Observations:  500
```

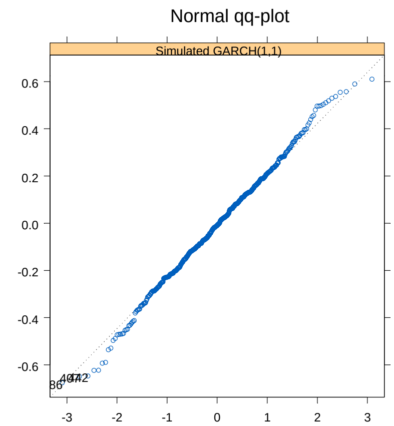

Testing for ARCH effects in daily MSFT returns

Modified Box-Ljung Statistic

```
> autocorTest(seriesMerge(msft.ts^2,
+          abs(msft.ts)), lag.n = 5)

Test for Autocorrelation: Ljung-Box

Null Hypothesis: no autocorrelation

Test Statistics:
            MSFT.1   MSFT.2
Test Stat 562.1016 463.8208
  p.value   0.0000   0.0000

Dist. under Null: chi-square with 5
degrees of freedom
   Total Observ.: 4365
```

`lag.n = 5` specifies 5 lags in modified Q stat

Easily reject H0: no autocorrelation

Testing for ARCH effects in daily MSFT returns

Engle's LM Test

```
> archTest(msft.ts, lag.n = 5)

Test for ARCH Effects: LM Test

Null Hypothesis: no ARCH effects

Test Statistics:
            MSFT
Test Stat 377.9377
  p.value   0.0000

Dist. under Null: chi-square with 5
degrees of freedom
   Total Observ.: 4365
```

`lag.n = 5` specifies 5 lags for LM test

Easily reject H0: no ARCH effects

Note: Less evidence for ARCH effects in monthly returns

Estimate GARCH(1,1) for MSFT

```
> msft.g11 = garch(msft.ts ~ 1,  ~ garch(1, 1))
Iteration   0  Step Size =  1.00000  Likelihood =  2.32406
Iteration   0  Step Size =  2.00000  Likelihood =  2.32445
Iteration   0  Step Size =  4.00000  Likelihood =  2.23425
Iteration   1  Step Size =  1.00000  Likelihood =  2.32489
Iteration   1  Step Size =  2.00000  Likelihood =  2.32497
Iteration   1  Step Size =  4.00000  Likelihood =  2.32402
Iteration   2  Step Size =  1.00000  Likelihood =  2.32506
Iteration   2  Step Size =  2.00000  Likelihood =  2.32500

Convergence R-Square = 4.710438e-006 is less than tolerance =
0.0001
Convergence reached.
```

Conditional mean formula: `msft ~ 1`

Conditional variance formula: `~ garch(1,1)`

Garch Object

```
> class(msft.g11)
[1] "garch"
```

print, summary, plot, predict, simulate

coef, fitted, residuals, sigma.t, vcov

```
> names(msft.g11)
 [1] "residuals"   "sigma.t"     "df.residual" "coef"   "model"
 [6] "cond.dist"   "likelihood"  "opt.index"   "cov"    "prediction
[11] "call"        "asymp.sd"    "series"
```

```
> msft.g11$coef
                 COEF
        C 0.00187904290
        A 0.00002803065
  ARCH(1) 0.09044383310
 GARCH(1) 0.86581806057
```

```
> coef(msft.g11)
        C 0.00187904290
        A 0.00002803065
  ARCH(1) 0.09044383310
 GARCH(1) 0.86581806057
```

```
> msft.g11$asymp.sd
[1] 0.0254471
```

$$\bar{\sigma} = \left(a_0 / (1 - a_1 - b_1)\right)^{1/2}$$

GARCH methods

```
> summary(msft.g11)

Call:
garch(formula.mean = msft.ts ~ 1, formula.var =  ~ garch(1, 1))

Conditional Distribution:  gaussian
--------------------------------------------------------------

Estimated Coefficients:
--------------------------------------------------------------
              Value  Std.Error t value   Pr(>|t|)
       C 0.00187904 3.256e-004   5.771 8.430e-009
       A 0.00002803 3.419e-006   8.199 4.441e-016
 ARCH(1) 0.09044383 5.924e-003  15.267 0.000e+000
GARCH(1) 0.86581806 1.016e-002  85.187 0.000e+000
```

Note: $a_1 + b_1 = 0.09 + 0.87 = 0.96$

Half life of shock to volatility: ln(0.5)/ln(0.96)=15.5 days

Examining a garch object

```
> names(msft.g11$model)
 [1] "c.which"   "c.value"     "MA"           "AR"         "arch"
 [6] "garch"     "a.which"     "a.value"      "info" "power.value"

> msft.g11$model

-- Constants in mean ---
    value which
 0.001879     1

-- Constants in variance--
      value which
 0.00002803     1

----- ARCH --------------
        value which
lag 1 0.09044     1

----- GARCH -------------
       value which
lag 1 0.8658     1
```

```
> msft.g11$model$garch
$order:
[1] 1

$which:
[1] T

$value:
[1] 0.8658181
```


Revising a garch model

```
  # revise model starting values
  > g11.model = msft.g11$model
  > g11.model$garch$value=0.9
  > g11.model$arch$value=0.1

 # reestimate with new values
 msft.g11.new = garch(series = msft.ts, model = g11.model)
 Iteration   0  Step Size =  1.00000  Likelihood =  2.31871
 Iteration   3  Step Size =  8.00000  Likelihood =  2.32465
 …
 Convergence R-Square = 9.360924e-006 is less than tolerance
 = 0.0001

> cbind(coef(msft.g11), coef(msft.g11.new))
                  [,1]           [,2]
       C 0.00190253165 0.00195889793
       A 0.00002755037 0.00002730452
 ARCH(1) 0.08994106397 0.08967032459
GARCH(1) 0.86751369929 0.86806428699
```

```
> cbind(msft.g11$likelihood,
msft.g11.new$likelihood)
         [,1]    [,2]
[1,] 10148.91 10148.9
```

The basic GARCH model assumes normal errors.

If the errors are not normal but the Gaussian likelihood is used then the resulting estimator is called the QMLE

Under general conditions, the QMLE is still consistent and asymptotically normal but the variance must be adjusted

```
> summary(msft.g11, method="QMLE")
> vcov(msft.g11, method="QMLE")
```


General Conditional Mean Specification

$$E_{t-1}[y_t] = c + \sum_{j=1}^{r} \phi_j y_{t-j} + \sum_{k=1}^{s} \theta_k \varepsilon_{t-k} + \sum_{l=0}^{L} \beta_l' z_{t-l}$$

General Conditional Variance Specification

$$\sigma_t^2 = a_0 + \sum_{i=1}^{p} a_i \varepsilon_{t-i}^2 + \sum_{j=1}^{q} b_j \sigma_{t-j}^2 + \sum_{l=0}^{L} \delta_k' z_{t-k}$$

```
args(garch)
function(formula.mean =  ~ arma(0, 0),
         formula.var =  ~ garch(0, 0),
         series.start = 1, x = NULL, x.start = 1, xlag = 0,
         z = NULL, z.start = 1, …
```

Note: Exogenous variables should be directly specified in formula.mean and formula.var

GARCH Model Diagnostics

- test if a coefficient is equal to zero
  - Standard t-statistics and p-value
- test if distribution of residuals is normal
  - `normalTest()`
  - Jarque-Bera test & Shapiro-Wilk test
- test if squared residuals are correlated
  - `autocorTest()`
  - Box-Pierce & modified Box-Pierce (or Ljung-Box)
- test if any ARCH effects remain in residuals
  - `archTest()`

Model diagnostics are computed and printed from `summary()`

Model diagnostics are reported graphically from `plot()`

Garch model diagnostics: summary()

```
> summary(msft.g11)
…
Normality Test:
--------------------------------------------------------------
 Jarque-Bera P-value Shapiro-Wilk    P-value
        1752       0       0.9782 2.367e-025

Ljung-Box test for standardized residuals:
--------------------------------------------------------------
 Statistic P-value Chi^2-d.f.
     17.62  0.1278         12

Ljung-Box test for squared standardized residuals:
--------------------------------------------------------------
 Statistic P-value Chi^2-d.f.
     4.792  0.9646         12

Lagrange multiplier test:
--------------------------------------------------------------
    TR^2 P-value F-stat P-value
 4.768  0.9653 0.4339  0.9918
```

Standardized residuals are non-normal

No autocorrelation or ARCH effects in residuals

Garch model diagnostics: plot()

```
> plot(msft.g11)

Make a plot selection (or 0 to exit):

1: plot: All
2: plot: Series and Conditional SD
3: plot: Series with 2 Conditional SD Superimposed
4: plot: ACF of the Observations
5: plot: ACF of Squared Observations
6: plot: Cross Corr. between Squared Series and Series
7: plot: Residuals
8: plot: Conditional Standard Deviations
9: plot: Standardized Residuals
10: plot: ACF of Standardized Residuals
11: plot: ACF of Squared Standardized Residuals
12: plot: Cross Corr. between Squared Std.Res and Std.Res
13: plot: QQ-Plot of Standardized Residuals
Selection:
```


Garch model diagnostics: conditional SD

```
> plot(msft.g11, which.plot = 1)
```

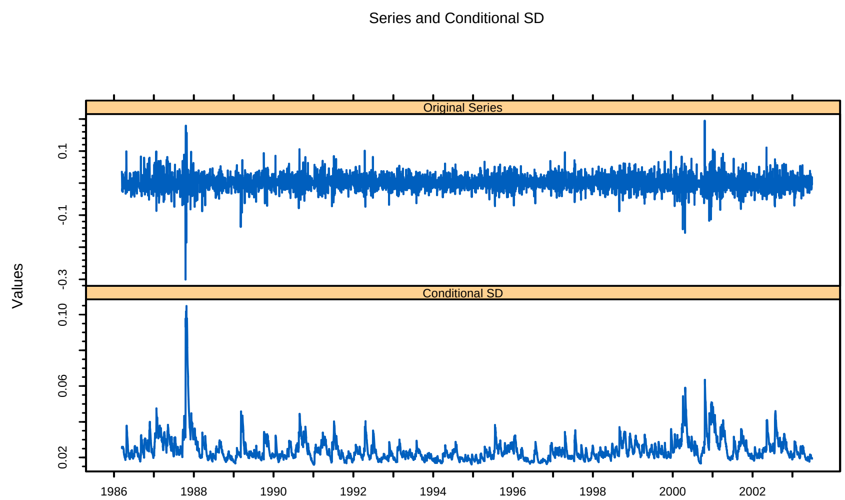

Garch model diagnostics: residuals

```
> plot(msft.g11, which.plot = 10) > plot(msft.g11, which.plot = 12
```

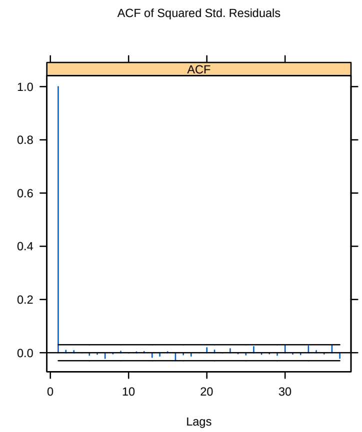

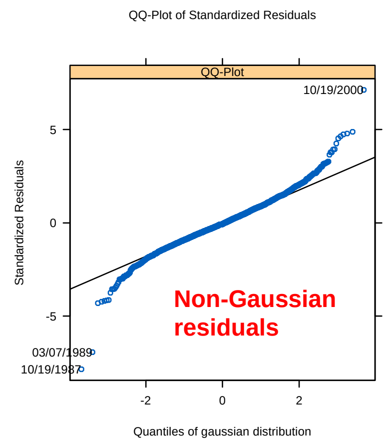

GARCH Model Selection

```
> msft.a1 = garch(msft.ts~1, ~garch(1,0))
> msft.a2 = garch(msft.ts~1, ~garch(2,0))
> msft.a3 = garch(msft.ts~1, ~garch(3,0))
> msft.a4 = garch(msft.ts~1, ~garch(4,0))
> msft.a5 = garch(msft.ts~1, ~garch(5,0))
> msft.g21 = garch(msft.ts~1, ~garch(2,1))
> msft.g22 = garch(msft.ts~1, ~garch(2,2))
> msft.g12 = garch(msft.ts~1, ~garch(1,2))

> compare.mgarch(msft.a1, msft.a2, msft.a3, msft.a4, msft.a5,
      msft.g11, msft.g12, msft.g21, msft.g22)
           msft.a1 msft.a2 msft.a3 msft.a4 msft.a5
       AIC  -19977  -20086  -20175  -20196  -20211
       BIC  -19958  -20060  -20143  -20158  -20166
Likelihood    9992   10047   10092   10104   10113

           msft.g11 msft.g12 msft.g21 msft.g22
       AIC   -20290   -20290   -20292   -20288
       BIC   -20264   -20258   -20260   -20249
Likelihood    10149    10150    10151    10150
```

Compare ARCH(5) to GARCH(1,1)

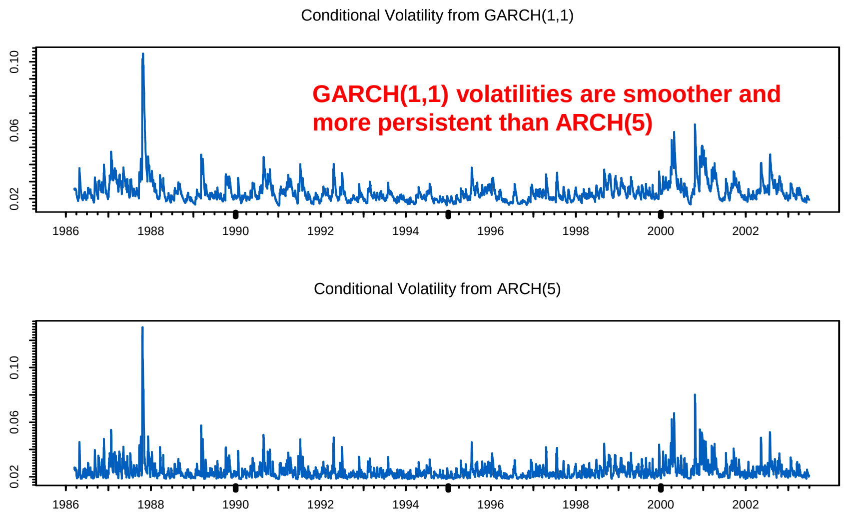

Rolling Analysis of GARCH Models

```
roll.garch <- function(data)
{
  my.mod = garch(data~1, ~garch(1,1),trace=F)
  return(list(a0 = coef(my.mod)[2],
              a1 = coef(my.mod)[3],
              b1 = coef(my.mod)[4],
              asymp.sd = my.mod$asymp.sd))
}
```

Rolling analysis is performed with FinMetrics function **`roll`**

```
# rolling GARCH(1,1) estimation for MSFT
# 2 years of daily data, rolled ahead by 1 month
> roll.fit.msft <- roll(FUN="roll.garch", data=msft.ts,
                        width=500, incr=20)
```

Rolling Analysis of GARCH(1,1)

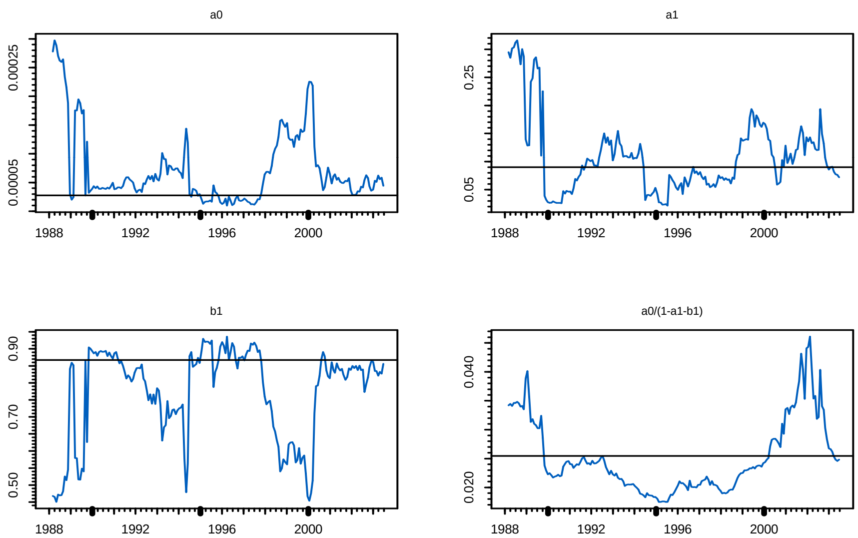

Rolling Analysis of GARCH(1,1)

Persistence parameter: $a_1 + b_1$

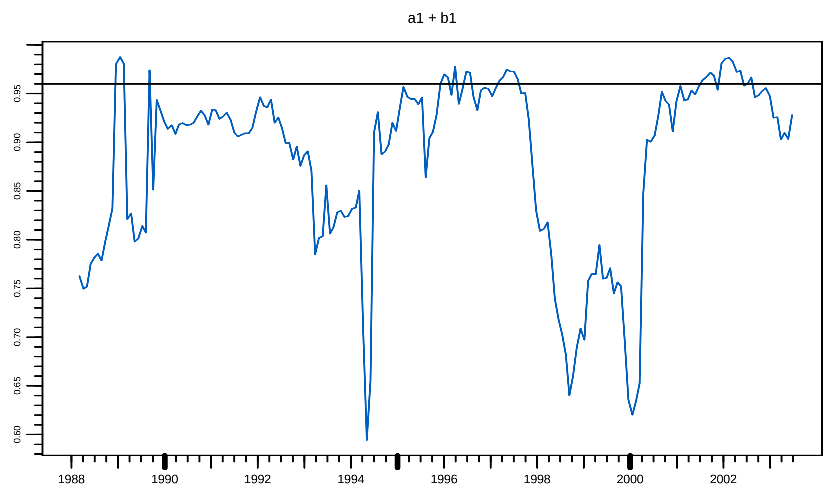

Garch predict() method

```
> msft.g11.pred = predict(msft.g11, n.predict = 250)
> class(msft.g11.pred)
[1] "predict.garch"
> names(msft.g11.pred)
[1] "series.pred" "sigma.pred"  "asymp.sd"
> plot(msft.g11.pred)
```

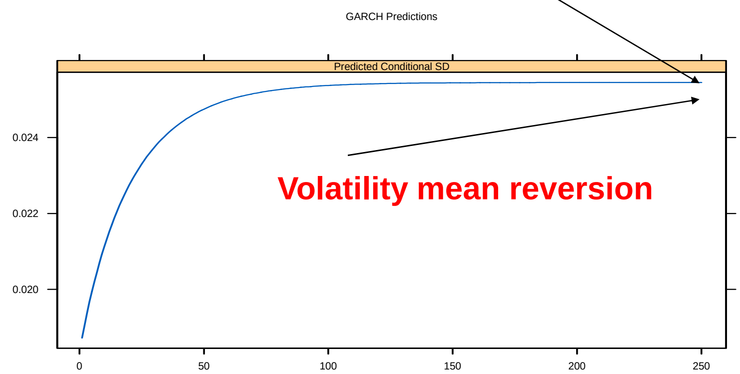

Mean reversion in volatility forecast

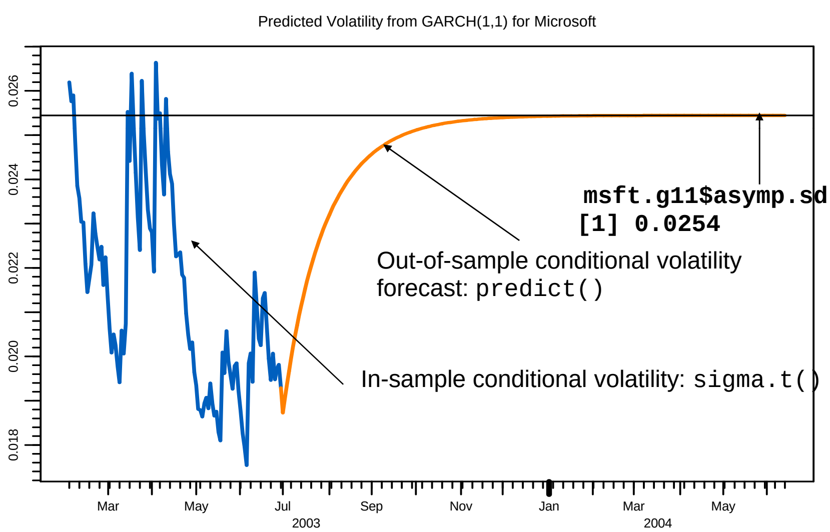

Volatility Term Structure

```
> term.dat = 1:200
> msft.volTermStruc = cumsum(predict(msft.g11,
                        n.predict=200)$sigma.pred^2)/term.dat
> tsplot(msft.volTermStruc)
```

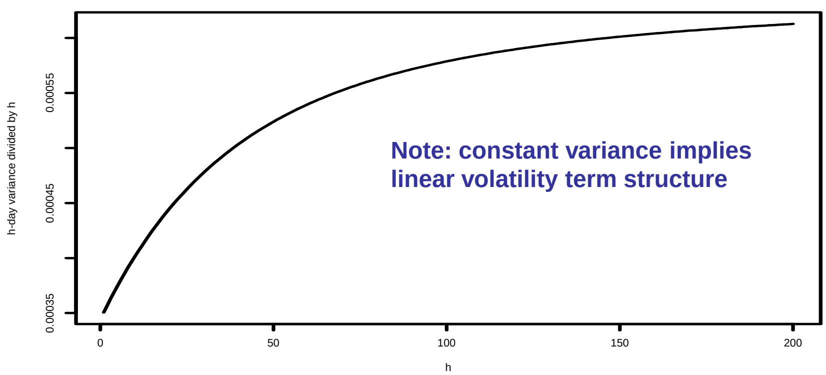

Testing for Asymmetry

```
> SignBias.fit = OLS(ehat2~tslag(S.minus), data=tmp.ts)
> NegativeSizeBias.fit = OLS(ehat2~tslag(S.minus*ehat),
                             data=tmp.ts)
> PositiveSizeBias.fit = OLS(ehat2~tslag(S.plus*ehat),
                             data=tmp.ts)
> All.fit =
+  OLS(ehat2~S.minus+tslag(S.minus*ehat)+tslag(S.plus*ehat),
   data=tmp.ts)
```

Asset	$\text{corr}(r_t^2, r_{t-1})$	Sign Bias	Negative Size Bias	Positive Size Bias
Microsoft	-0.0315	-0.4417	-6.816	3.174
		(0.6587)	(0.000)	(0.001)
S&P 500	-0.098	2.457	-11.185	1.356
		(0.014)	(0.000)	(0.175)

Notes: p-values are in parentheses.

Clear evidence for leverage effects

Exponential GARCH Model (EGARCH)

$$h_t = \alpha_0 + \sum_{i=1}^{p} \alpha_i \frac{|\varepsilon_{t-i}| + \gamma_i \varepsilon_{t-i}}{\sigma_{t-i}} + \sum_{j=1}^{q} \beta_j h_{t-j}$$

$$h_t = log(\sigma_t^2), \sigma_t^2 = \exp(h_t)$$

- good news effect $\varepsilon_{\tau-\iota} : (1+\gamma_\iota)\varepsilon_{\tau-\iota}$
- bad news effect $\varepsilon_{\tau-\iota} : (1-\gamma_\iota)|\varepsilon_{\tau-\iota}|$
- usually, $\gamma_i < 0$

Note: Conditional variance is always positive due to exponential transformation

```
garch(hp.s~1, ~egarch(1,1), leverage=T)
```


Threshold GARCH Model (TGARCH)

$$\sigma_t^2 = a_0 + \sum_{i=1}^{p} a_i \varepsilon_{t-i}^2 + \sum_{i=1}^{p} \gamma_i S_{t-i} \varepsilon_{t-i}^2 + \sum_{j=1}^{q} b_j \sigma_{t-j}^2$$

$$S_t = \begin{cases} 1 & if \quad \varepsilon_t < 0 \\ 0 & if \quad \varepsilon_t \geq 0 \end{cases}$$

- good news effect $\varepsilon_{t-i} : a_i \varepsilon_{t-i}^2$
- bad news effect $\varepsilon_{t-i} : (a_i + \gamma_i)\varepsilon_{t-i}^2$
- usually, $\gamma_i > 0$

```
garch(hp.s~1, ~tgarch(1,1))
```

Power GARCH Model (PGARCH)

$$\sigma_t^d = a_0 + \sum_{i=1}^{p} a_i(|\varepsilon_{t-i}| + \gamma_i \varepsilon_{t-i})^d + \sum_{j=1}^{q} b_j \sigma_{t-j}^d$$

d is a positive exponent (can be fixed or estimated from data).

- γ_i denotes the coefficient of the leverage effects.
- d can be fixed (~garch(p,q,d)) or estimated by ML method (~garch(p,q))

```
garch(hp.s~1, ~pgarch(1,1,1), leverage=T)
```


Estimation of Asymmetric GARCH models

```
msft.egarch = garch(msft.ts~1, ~egarch(1,1), leverage=T)
msft.tgarch = garch(msft.ts~1, ~tgarch(1,1))
msft.pgarch.2 = garch(msft.ts~1, ~garch(1,1), leverage=T)
msft.pgarch.1 = garch(msft.ts~1, ~pgarch(1,1,1),leverage=T)
```

Model	a_0	a_1	b_1	γ_1	BIC
			Microsoft		
EGARCH	-0.7273	0.2144	0.9247	-0.2417	-20265
	[0.4064]	[0.0594]	[0.0489]	[0.0758]	
TGARCH	$3.01e^{-5}$	0.0564	0.8581	0.0771	-20291
	$[1.02e^{-5}]$	[0.0141]	[0.0342]	[0.0306]	
PGARCH 2	$2.87e^{-5}$	0.0853	0.8672	-0.2164	-20290
	$[9.27e^{-6}]$	[0.0206]	[0.0313]	[0.0579]	
PGARCH 1	0.0010	0.0921	0.8876	-0.2397	-20268
	[0.0006]	[0.0236]	[0.0401]	[0.0813]	

News Impact curves of GARCH processes

Functional relationship between condtional variance at time t and error term at time t-1, holding constant information dated t-2 and earlier and with all lagged conditional variances evaluated at unconditional variance

$$
\begin{array}{l|l}
GARCH(1,1) & \begin{array}{l} \sigma_t^2 = A + a_1(|\varepsilon_{t-1}| + \gamma_1 \varepsilon_{t-1})^2 \\ A = a_0 + b_1\bar{\sigma}^2 \end{array} \\
TGARCH(1,1) & \begin{array}{l} \sigma_t^2 = A + (a_1 + \gamma_1 S_{t-1})\varepsilon_{t-1}^2 \\ A = a_0 + b_1\bar{\sigma}^2 \end{array} \\
PGARCH(1,1,1) & \begin{array}{l} \sigma_t^2 = A + 2\sqrt{A}a_1(|\varepsilon_{t-1}| + \gamma_1\varepsilon_{t-1}) + a_1^2(|\varepsilon_{t-1}| + \gamma_1\varepsilon_{t-1})^2 \\ A = (a_0 + b_1\bar{\sigma})^2 \end{array} \\
EGARCH(1,1) & \begin{array}{l} \sigma_t^2 = A\exp\{a_1(|\varepsilon_{t-1}| + \gamma_1\varepsilon_{t-1})/\bar{\sigma}\} \\ A = \bar{\sigma}^{2b_1}\exp\{a_0\} \end{array}
\end{array}
$$

Estimated News Impact Curves

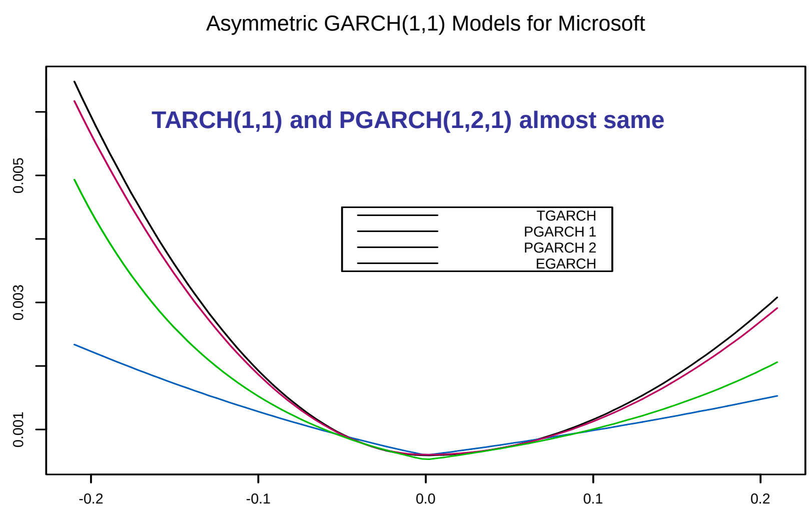

Estimation of Non-Gaussian GARCH(1,1) Models

```
msft.g11t = garch(msft.ts~1, ~garch(1,1), cond.dist="t“)
msft.tgarch.t = garch(msft.ts~1, ~tgarch(1,1), cond.dist="t")
```

Model	a_0	a_1	b_1	γ_1	v	BIC
			Microsoft			
GARCH	$3.39e^{-5}$ $[1.52e^{-5}]$	0.0939 [0.0241]	0.8506 [0.0468]		6.856 [0.7121	-20504
TGARCH	$3.44e^{-5}$ $[1.20e^{-5}]$	0.0613 [0.0143]	0.8454 [0.0380]	0.0769 [0.0241]	7.070 [0.7023]	-20511

Normal GARCH BIC = -20264

Normal TARCH BIC = -20291

QQ-plot of Standardized Residuals from TGARCH with t-errors

Nice and linear QQ-plot

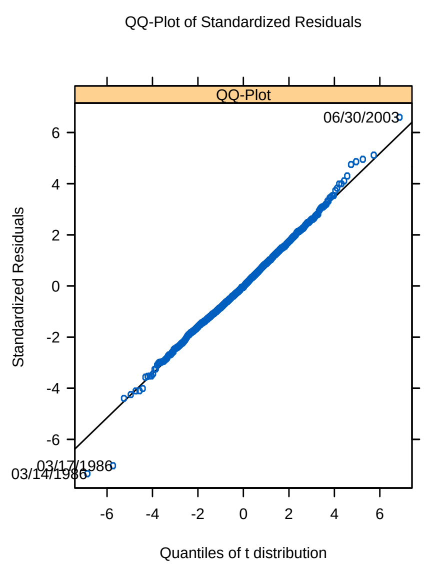

Comparison of GARCH Model Forecasts

Volatility Forecasts for Microsoft from Competing GARCH Models