Volatility

- Many economic series, and most financial series, display conditional volatility
  - The conditional variance changes over time
  - There are periods of high volatility
    - When large changes frequently occur
  - And periods of low volatility
    - When large changes are less frequent

Weekly Stock Prices
Levels and Returns

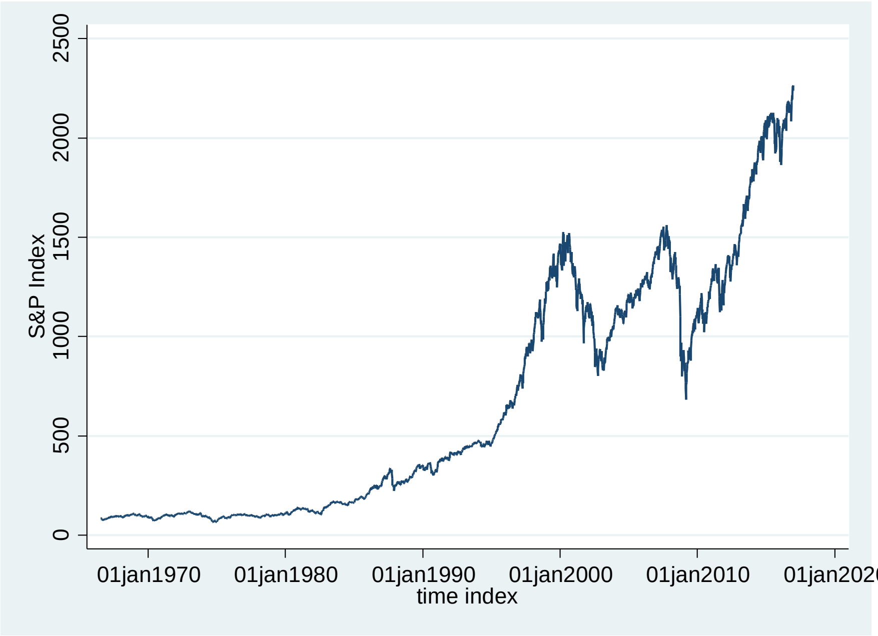

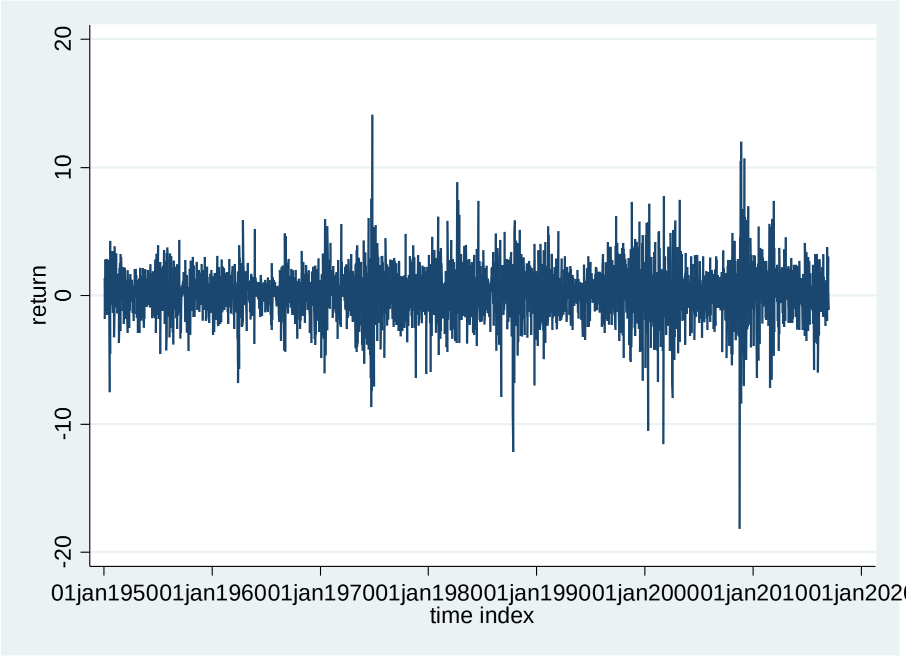

Conditional Mean

- The conditional mean of y is

$$E\left(y_t \mid \Omega_{t-1}\right)$$

- The regression error is mean zero and unforecastable

$$E\left(e_t \mid \Omega_{t-1}\right) = 0$$

Conditional Variance

- The conditional variance of y is

$$\operatorname{var}(y_t \mid \Omega_{t-1}) = E\left(\left(y_t - E(y_t \mid \Omega_{t-1})\right)^2 \mid \Omega_{t-1}\right)$$
$$= E\left(e_t^2 \mid \Omega_{t-1}\right)$$

- The squared regression error can be forecastable

Forecastable Conditional Variance

- If the squared error is forecastable, then the conditional variance is time-varying and correlated.
  - The magnitude of changes is predictable
  - The sign is not predictable

Stock returns are unpredictable

```
. reg return L(1/4).return, r

Linear regression                               Number of obs     =      3,491
                                                F(4, 3486)        =       1.26
                                                Prob > F          =     0.2846
                                                R-squared         =     0.0032
                                                Root MSE          =     2.0601
```

return	Coef.	Robust Std. Err.	t	P>|t|	[95% Conf.	Interval]
return						
L1.	-.0313688	.028884	-1.09	0.278	-.088	.0252623
L2.	.0374993	.0253879	1.48	0.140	-.0122772	.0872759
L3.	-.0077905	.0264173	-0.29	0.768	-.0595854	.0440044
L4.	-.0279404	.0249751	-1.12	0.263	-.0769077	.0210268
_cons	.1655475	.039166	4.23	0.000	.0887568	.2423381

```
. testparm L(1/4).return

 ( 1)  L.return = 0
 ( 2)  L2.return = 0
 ( 3)  L3.return = 0
 ( 4)  L4.return = 0

       F(  4,  3486) =    1.26
            Prob > F =    0.2846
```

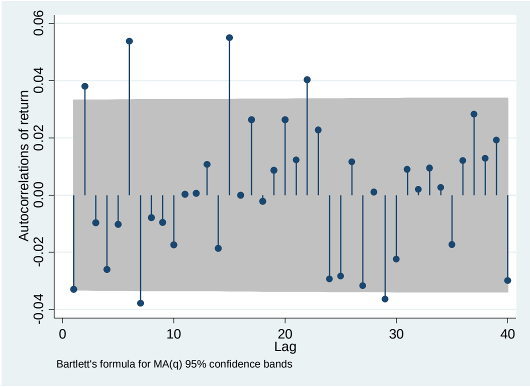

Squared Returns are predictable

```
. gen y = (return-0.1629)^2

. reg y L(1/4).y, r

Linear regression                               Number of obs     =      3,491
                                                F(4, 3486)        =       8.82
                                                Prob > F          =     0.0000
                                                R-squared         =     0.1155
                                                Root MSE          =      10.55

------------------------------------------------------------------------------
             |               Robust
           y |      Coef.   Std. Err.      t    P>|t|     [95% Conf. Interval]
-------------+----------------------------------------------------------------
           y |
         L1. |   .2140044   .0983976     2.17   0.030     .0210817    .4069272
         L2. |   .0712205   .0309978     2.30   0.022      .010445    .1319961
         L3. |     .14427   .0436719     3.30   0.001      .058645     .229895
         L4. |   .0687664   .0535968     1.28   0.200    -.0363178    .1738506
             |
       _cons |   2.132529   .3273144     6.52   0.000     1.490782    2.774276
------------------------------------------------------------------------------

. testparm L(1/4).y

 ( 1)  L.y = 0
 ( 2)  L2.y = 0
 ( 3)  L3.y = 0
 ( 4)  L4.y = 0

       F(  4,  3486) =    8.82
            Prob > F =    0.0000
```

Squared Returns

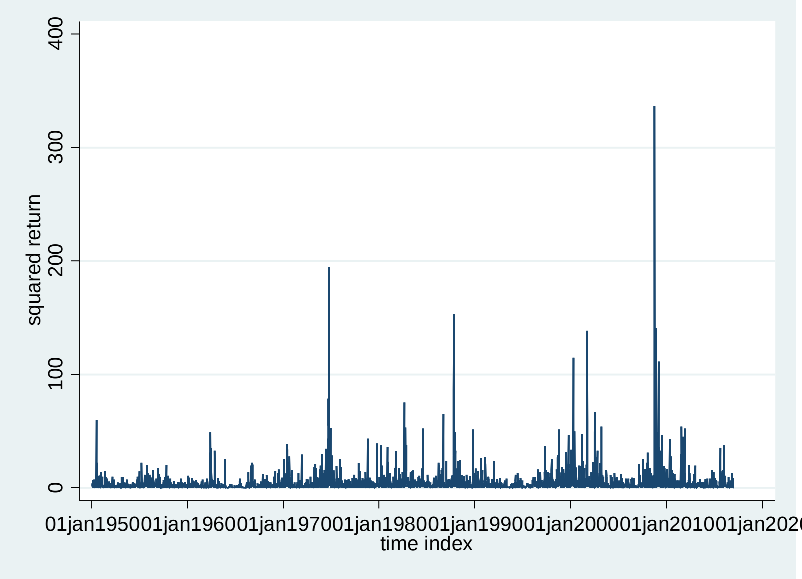

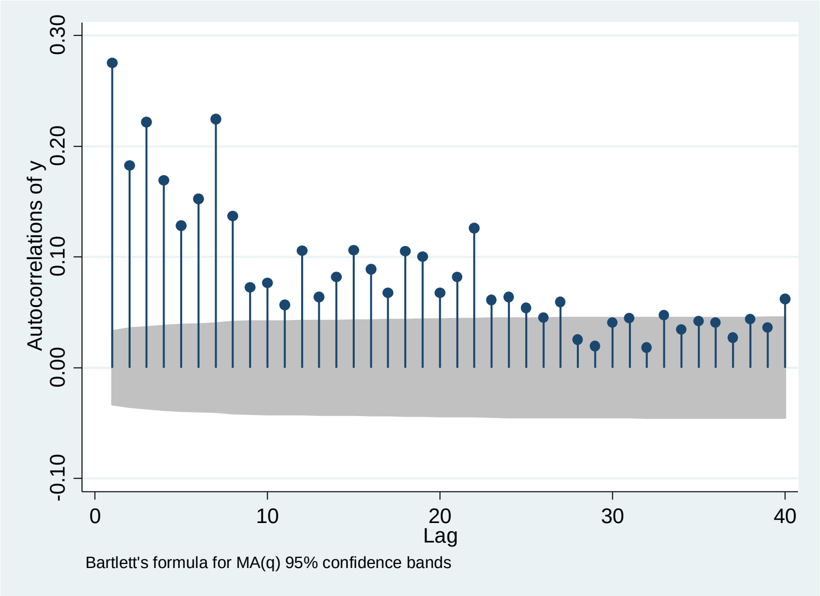

ARCH

- Robert Engle (1982) proposed a model for the conditional variance
  - AutoRegressive Conditional Heteroskedasticity
  - “ARCH” now describes volatility models
- Nobel Prize 2003

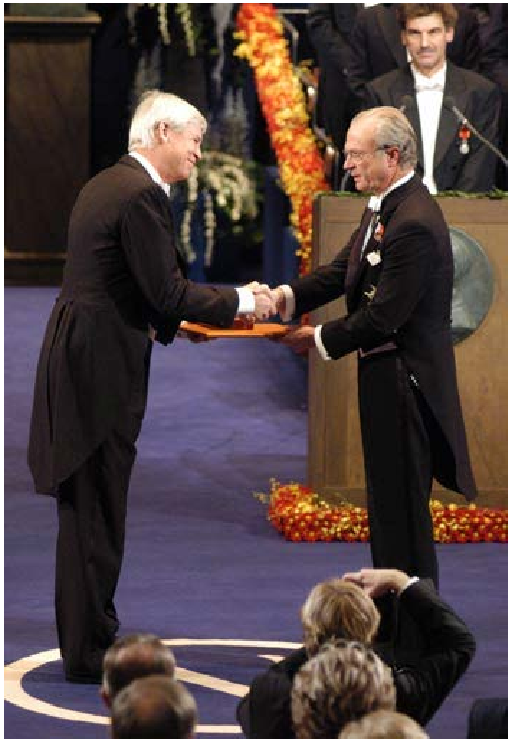

ARCH(1) Model

$$y_t = \mu + e_t$$

$$\sigma_t^2 = \mathrm{var}(e_t \mid \Omega_{t-1}) = \omega + \alpha e_{t-1}^2$$

$$\omega > 0$$

$$\alpha \geq 0$$

- $\alpha>0$ means that the conditional variance is high when the lagged squared error is high
- Large errors (either sign) today mean high expected errors (in magnitude) tomorrow.
- Small magnitude errors forecast next period small magnitude errors.

Unconditional variance

- A property of expectations is that expected (average) conditional expectations are unconditional expectations.
- So the average conditional variance is the average variance – the variance of the regression error.

$$\sigma^2 = E\left(\sigma_t^2\right) = \omega + \alpha E\left(e_{t-1}^2\right) = \omega + \alpha\sigma^2$$

- Solving for the variance: $\sigma^2 = \dfrac{\omega}{1-\alpha}$

- Rewriting, this implies

$$\omega = \sigma^2 (1 - \alpha)$$

- Substituting into ARCH(1) equation

$$\sigma_t^2 = (1 - \alpha)\sigma^2 + \alpha e_{t-1}^2$$

or

$$\sigma_t^2 = \sigma^2 + \alpha\left(e_{t-1}^2 - \sigma^2\right)$$

- This shows that the conditional variance is a combination of the unconditional variance, and the deviation of the squared error from its average value.

ARCH(1) as AR(1) in squares

- The model

$$\text{var}\left(e_t \mid \Omega_{t-1}\right) = E\left(e_t^2 \mid \Omega_{t-1}\right) = \omega + \alpha e_{t-1}^2$$

implies the regression

$$e_t^2 = \omega + \alpha e_{t-1}^2 + u_t$$

where u is white noise

- Thus e-squared is an AR(1)

Estimation

- **.arch return, arch(1)**

```
Sample: 13jan1950 - 30dec2016                   Number of obs   =      3,495
Distribution: Gaussian                          Wald chi2(.)    =          .
Log likelihood = -7302.702                      Prob > chi2     =          .
```

return	Coef.	OPG Std. Err.	z	P>|z|	[95% Conf.	Interval]
return						
_cons	.2150146	.0300858	7.15	0.000	.1560476	.2739817
ARCH						
arch						
L1.	.2794291	.0200765	13.92	0.000	.2400799	.3187783
_cons	2.986409	.0622812	47.95	0.000	2.86434	3.108478

Variance Forecast

- Given the parameter estimates, the estimated conditional variance for period t is

$$\hat{\sigma}_t^2 = \hat{\omega} + \hat{\alpha}\hat{e}_{t-1}^2 = \hat{\omega} + \hat{\alpha}\left(y_{t-1} - \hat{\mu}\right)^2$$

- The forecasted out-of-sample variance is

$$\hat{\sigma}_{n+1}^2 = \hat{\omega} + \hat{\alpha}\left(y_n - \hat{\mu}\right)^2$$

Forecast Interval for the mean

- You can use the estimated conditional standard deviation to obtain forecast intervals for the mean

$$\hat{y}_{n+1|n} \pm Z_{\alpha/2}\hat{\sigma}_{n+1}$$

- These forecast intervals will vary in width depending on the estimated conditional variance.
  - Wider in periods of high volatility
  - More narrow in periods of low volatility

ARCH(p) model

- Allow p lags of squared errors

$$y_t = \mu + e_t$$
$$\sigma_t^2 = \omega + \alpha_1 e_{t-1}^2 + \alpha_2 e_{t-2}^2 + \cdots + \alpha_p e_{t-p}^2$$

- Similar to AR(p) in squares
- Estimation: ARCH(8)
  - **.arch return, arch(1/8)**
  - ARCH model with lags 1 through 8

ARCH(8) Estimates

- **.arch return, arch(1/8)**

```
Sample: 13jan1950 - 30dec2016                   Number of obs     =      3,495
Distribution: Gaussian                          Wald chi2(.)      =          .
Log likelihood = -7117.767                      Prob > chi2       =          .
```

return	Coef.	OPG Std. Err.	z	P>|z|	[95% Conf.	Interval]
return						
_cons	.2281217	.0273503	8.34	0.000	.1745161	.2817272
ARCH						
arch						
L1.	.1670606	.0153543	10.88	0.000	.1369668	.1971545
L2.	.1137801	.0200296	5.68	0.000	.0745229	.1530374
L3.	.1533685	.0206159	7.44	0.000	.112962	.193775
L4.	.0981257	.0185268	5.30	0.000	.0618138	.1344377
L5.	.0243815	.016271	1.50	0.134	-.0075092	.0562721
L6.	.0812137	.0182713	4.44	0.000	.0454026	.1170249
L7.	.0404796	.0137457	2.94	0.003	.0135385	.0674206
L8.	.0631067	.0157473	4.01	0.000	.0322426	.0939708
_cons	1.1884	.0949953	12.51	0.000	1.002212	1.374587

ARCH needs many lags

- Notice that we included 8 lags, and all appeared significant.
- This is commonly observed in estimated ARCH models
  - The conditional variance appears to be a function of many lagged past squares

GARCH Model

- Tim Bollerslev (1986)
  - A student of Engle
  - Professor at Duke University

proposed the GARCH model to simplify this problem

$$\sigma_t^2 = \omega + \beta\sigma_{t-1}^2 + \alpha e_{t-1}^2$$

$$\beta > 0$$

$$\omega > 0$$

$$\alpha \geq 0$$

GARCH(1,1)

- This makes the variance a function of all past lags:

$$\sigma_t^2 = \omega + \beta\sigma_{t-1}^2 + \alpha e_{t-1}^2$$

$$= \sum_{j=0}^{\infty} \beta^j \left(\omega + \alpha e_{t-1-j}^2\right)$$

- It is also smoother than an ARCH model with a small number of lags

GARCH(p,q)

- p lags of squared error
- q lags of conditional variance

$$\sigma_t^2 = \omega + \beta_1 \sigma_{t-1}^2 + \cdots + \beta_q \sigma_{t-q}^2 + \alpha_1 e_{t-1}^2 + \cdots + \alpha_p e_{t-p}^2$$

- GARCH(1,1):
  - **.arch r, arch(1) garch(1)**
- GARCH(3,2):
  - **.arch r, arch(1/3) garch(1/2)**

GARCH(1,1)

```
Sample: 13jan1950 - 30dec2016                     Number of obs   =      3,495
Distribution: Gaussian                            Wald chi2(.)    =          .
Log likelihood = -7111.227                        Prob > chi2     =          .
```

return	Coef.	OPG Std. Err.	z	P>|z|	[95% Conf.	Interval]
return						
_cons	.2168081	.0277768	7.81	0.000	.1623666	.2712497
ARCH						
arch L1.	.131982	.0098944	13.34	0.000	.1125893	.1513748
garch L1.	.8386383	.0124113	67.57	0.000	.8143125	.8629641
_cons	.1379501	.0234193	5.89	0.000	.092049	.1838511

- Common GARCH features
  - Lagged variance (garch) has large coefficient
  - Sum of two coefficients very close to (but less than) one

GARCH(2,2) for Stock Returns

```
Sample: 13jan1950 - 30dec2016                   Number of obs   =      3,495
Distribution: Gaussian                          Wald chi2(.)    =          .
Log likelihood = -7110.359                      Prob > chi2     =          .
```

return	Coef.	OPG Std. Err.	z	P>|z|	[95% Conf.	Interval]
return						
_cons	.2154139	.0278286	7.74	0.000	.1608707	.269957
ARCH						
arch						
L1.	.1489542	.0143911	10.35	0.000	.1207481	.1771603
L2.	.0202212	.0719248	0.28	0.779	-.1207487	.1611912
garch						
L1.	.4584953	.525055	0.87	0.383	-.5705937	1.487584
L2.	.3349021	.4410889	0.76	0.448	-.5296162	1.19942
_cons	.1735723	.0774155	2.24	0.025	.0218406	.325304

GARCH(1,1)

- The GARCH(1,1) often fits well, and is a useful benchmark.
  - Daily, weekly, or monthly asset returns, exchange rates, or interest rates

Extensions

- There are many extensions of the basic GARCH model, developed to handle a variety of situations
  - Asymmetric Response
  - Garch-in-mean
  - Explanatory variables in variance
  - Non-normal errors

Asymetric GARCH

- Threshold GARCH

$$\sigma_t^2 = \omega + \beta\sigma_{t-1}^2 + \alpha e_{t-1}^2 + \gamma e_{t-1}^2 1\left(e_{t-1} > 0\right)$$

- The last tern is dummy variable for positive lagged errors
- This model specifies that the ARCH effect depends on whether the error was positive or negative
  - If the error is negative, the effect is α
  - If the error is positive, the full effect is α+γ

TARCH estimation

- **.arch return, arch(1) tarch(1) garch(1)**
- Negative errors have coefficient of 0.21
- Positive errors have coefficient of 0.04
- Negative returns increase volatility much more than positive returns

```
Sample: 13jan1950 - 30dec2016                    Number of obs   =      3,495
Distribution: Gaussian                           Wald chi2(.)    =          .
Log likelihood = -7071.684                       Prob > chi2     =          .
```

return	Coef.	OPG Std. Err.	z	P>|z|	[95% Conf.	Interval]
return						
_cons	.1633676	.0292604	5.58	0.000	.1060183	.2207169
ARCH						
arch L1.	.2128052	.0178016	11.95	0.000	.1779147	.2476957
tarch L1.	-.1735807	.017781	-9.76	0.000	-.2084309	-.1387306
garch L1.	.8279936	.0141225	58.63	0.000	.800314	.8556732
_cons	.1842874	.0239465	7.70	0.000	.1373531	.2312217

Leverage Effect

- This model describes what is called the "leverage effect"
  - A negative shock to equity increases the ratio debt/equity of investors
  - This increases the *leverage* of their portfolios
  - This increases risk, and the conditional variance
  - Negative shocks have stronger effect on variance than positive shocks

GARCH-in-mean

- If investors are risk averse, risky assets will earn higher returns (a risk premium) in market equilibrium
- If assets have varying volatility (risk), their expected return will vary with this volatility
  - Expected return should be positively correlated with volatility

GARCH-M model

$$y = \beta_1 + \beta_1 \sigma_{t-1}^2 + e_t$$

$$\sigma_t^2 = \omega + \beta \sigma_{t-1}^2 + \alpha e_{t-1}^2$$

- **.arch return, arch(1) garch(1) archm**

GARCH-M for Stock Returns

- Positive effect

```
Sample: 13jan1950 - 30dec2016                    Number of obs   =      3,495
Distribution: Gaussian                           Wald chi2(1)    =       6.64
Log likelihood = -7107.479                       Prob > chi2     =     0.0100
```

return	Coef.	OPG Std. Err.	z	P>|z|	[95% Conf.	Interval]
return						
_cons	.1136797	.04841	2.35	0.019	.0187979	.2085616
ARCHM						
sigma2	.0351553	.0136456	2.58	0.010	.0084103	.0619003
ARCH						
arch L1.	.1321214	.0100517	13.14	0.000	.1124205	.1518223
garch L1.	.838641	.0125663	66.74	0.000	.8140116	.8632705
_cons	.1372439	.0236595	5.80	0.000	.0908722	.1836156

TARCH and GARCH-M

- **.arch return, arch(1) tarch(1) garch(1) archm**
- archm effect reduced, appears insignificant

```
Sample: 13jan1950 - 30dec2016                   Number of obs   =      3,495
Distribution: Gaussian                          Wald chi2(1)    =       1.11
Log likelihood = -7071.057                      Prob > chi2     =     0.2911
```

return	Coef.	OPG Std. Err.	z	P>|z|	[95% Conf.	Interval]
return						
_cons	.124786	.0472928	2.64	0.008	.0320939	.2174781
ARCHM						
sigma2	.013834	.0131047	1.06	0.291	-.0118508	.0395188
ARCH						
arch L1.	.2106083	.0176877	11.91	0.000	.1759411	.2452755
tarch L1.	-.1696995	.0176998	-9.59	0.000	-.2043905	-.1350086
garch L1.	.8248105	.0147597	55.88	0.000	.795882	.8537389
_cons	.1921108	.0255139	7.53	0.000	.1421045	.2421171

Estimated standard deviation

- Estimated TARCH model
- **.predict v, variance**
- **.gen s=sqrt(v)**
- Average volatility=1.94

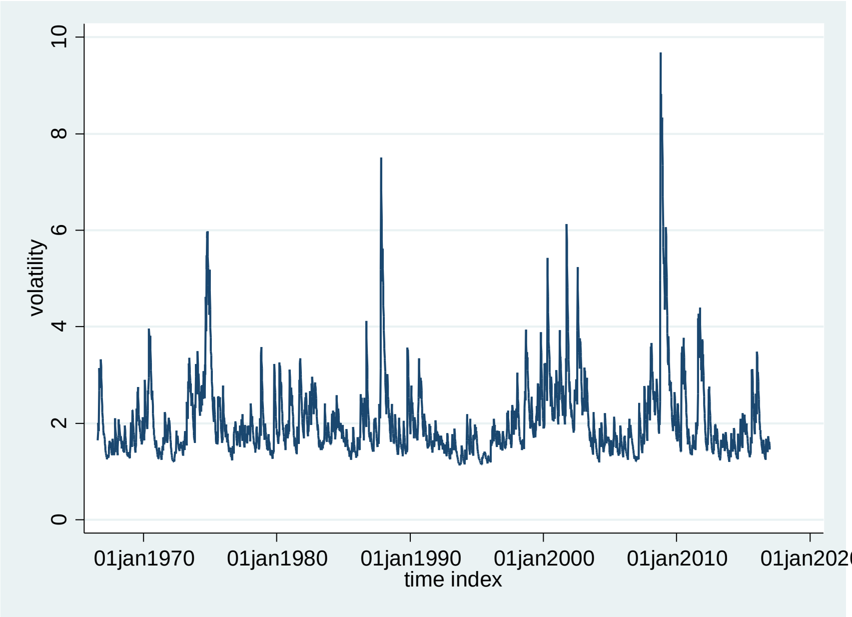

S&P, returns, and estimated volatilty 2006-2016