Financial Econometrics
Shanghai Jiao Tong University
Summer School – 2026
Empirical Exercise 16.2
With Solutions

Fatih Kansoy

Reference: Stock, J.H. and Watson, M.W., Introduction to Econometrics (4th Edition, Global Edition)

Data: ***us_macro_quarterly.csv****. Real GDP is the column* ***GDPC1****; compute* $GDPGR_t = 400\,[\ln(GDP_t) - \ln(GDP_{t-1})]$*. Sample: 1960:Q1–2012:Q4.*

Empirical Exercise 16.2. The data file `us_macro_quarterly.csv` contains quarterly data on US real GDP. Real GDP is the column `GDPC1`. Construct the quarterly growth rate of GDP at an annual rate,

$$GDPGR_t = 400\left[\ln(GDP_t) - \ln(GDP_{t-1})\right],$$

where GDP_t denotes real GDP in quarter t. The factor of 400 converts the quarterly log change into an annualised percentage growth rate.

(a) Using data on $\texttt{GDPGR}_t$ from 1960:Q1 to 2012:Q4, estimate an AR(2) model for the growth rate of GDP with GARCH(1,1) errors. Report the estimated conditional mean equation and the estimated conditional variance equation.

(b) Plot the residuals from the AR(2) model together with $\pm\,\hat{\sigma}_t$ bands, where $\hat{\sigma}_t$ is the conditional standard deviation implied by the estimated GARCH(1,1) model, in the manner of Figure 16.4 in the text.

(c) Some macroeconomists have claimed that there was a sharp drop in the variability of the growth rate of GDP around 1983, which they call the *Great Moderation*. Is this Great Moderation evident in the plot that you formed in part (b)?

Solution:

Data Overview

The dataset `us_macro_quarterly.csv` contains quarterly US macroeconomic series. We use the real GDP column `GDPC1` and construct the annualised quarterly growth rate

$$GDPGR_t = 400\left[\ln(GDP_t) - \ln(GDP_{t-1})\right].$$

The estimation sample runs from 1960:Q1 to 2012:Q4 (212 quarterly observations of `GDPGR`). Over this period the growth rate averages roughly 3% per year, with a standard deviation of about 3.4 percentage points, and it shows clear episodes of clustered volatility, that is, calm and turbulent periods that arrive in runs rather than at random.

Object	**Symbol**	**Meaning**
Level series	`GDPC1`	Real GDP (chained dollars)
Growth rate	$GDPGR_t$	$400\,[\ln(GDP_t) - \ln(GDP_{t-1})]$, annualised %
Conditional mean	AR(2)	$GDPGR_t$ on its first two lags
Conditional variance	GARCH(1,1)	time-varying σ_t^2 of the error

Key Notation and Model

The conditional mean is an AR(2) process,

$$GDPGR_t = \beta_0 + \beta_1\, GDPGR_{t-1} + \beta_2\, GDPGR_{t-2} + u_t,$$

where the error $u_t = \sigma_t\,\nu_t$ with ν_t independent of past information and having mean zero and unit variance. The conditional variance follows a GARCH(1,1) process,

$$\sigma_t^2 = \alpha_0 + \alpha_1\, u_{t-1}^2 + \phi_1\, \sigma_{t-1}^2.$$

The term $\alpha_1 u_{t-1}^2$ (the ARCH term) lets a large shock last quarter raise this quarter's variance; the term $\phi_1 \sigma_{t-1}^2$ (the GARCH term) lets the variance carry over smoothly from one quarter to the next. The model is estimated by maximum likelihood.

(a) AR(2) Model with GARCH(1,1) Errors

Estimating the model by maximum likelihood over 1960:Q1–2012:Q4 gives the following results (standard errors in parentheses).

Conditional Mean Equation

$$\boxed{\widehat{GDPGR}_t = 1.55 + 0.304\, GDPGR_{t-1} + 0.200\, GDPGR_{t-2}}$$

$$(0.33) \quad (0.07) \qquad\qquad (0.07)$$

Conditional Variance Equation

$$\boxed{\hat{\sigma}_t^2 = 0.41 + 0.22\, u_{t-1}^2 + 0.76\, \hat{\sigma}_{t-1}^2}$$

$$(0.22) \quad (0.05) \qquad (0.05)$$

Equation	**Parameter**	**Estimate**	**Std. Error**
Mean (AR(2))	β_0 (constant)	1.55	0.33
	β_1 (lag 1)	0.304	0.07
	β_2 (lag 2)	0.200	0.07
Variance (GARCH(1,1))	α_0 (constant)	0.41	0.22
	α_1 (ARCH, u_{t-1}^2)	0.22	0.05
	ϕ_1 (GARCH, σ_{t-1}^2)	0.76	0.05

Interpretation:

- **Conditional mean.** Both autoregressive coefficients are positive and individually significant (each estimate is several times its standard error). Growth has positive persistence: a quarter of above-average growth tends to be followed by further above-average growth. The long-run (unconditional) mean implied by the AR(2) is $1.55/(1 - 0.304 - 0.200) \approx 3.1\%$ per year, which matches the sample average.

- **Conditional variance.** Both GARCH coefficients are strongly significant. The ARCH coefficient $\hat{\alpha}_1 = 0.22$ shows that a large surprise in growth raises next quarter's variance; the GARCH coefficient $\hat{\phi}_1 = 0.76$ shows that variance is highly persistent and decays only slowly.

- **Volatility persistence.** The sum $\hat{\alpha}_1 + \hat{\phi}_1 = 0.22 + 0.76 = 0.98$ is close to (but below) one. The conditional variance is stationary, but shocks to volatility die out only gradually, so high-volatility and low-volatility episodes persist for many quarters. This is exactly the volatility clustering that the GARCH model is designed to capture.

Note: estimates may differ slightly across software packages depending on the optimiser and the treatment of starting values.

(b) Residuals with $\pm\,\hat{\sigma}_t$ Bands

[See Figure: AR(2) residuals $\hat{u}_t$ with $\pm\,\hat{\sigma}_t$ GARCH(1,1) bands, 1960:Q1–2012:Q4]

The plot shows the AR(2) residuals $\hat{u}_t$ (the part of growth not explained by its own two lags) against time, with the time-varying conditional-standard-deviation bands $\pm\,\hat{\sigma}_t$ drawn around zero. The bands are wide when the estimated conditional standard deviation is large and narrow when it is small, so their width is a direct read-out of the model's volatility forecast.

What to look for in the figure:

- In the 1960s and especially the 1970s and very early 1980s, the residuals are large and scattered and the $\pm\,\hat{\sigma}_t$ bands are wide. This period contains the oil-price shocks and the deep 1973–75 and 1981–82 recessions.
- From the mid-1980s onward the residuals are noticeably smaller and the bands are much narrower, indicating a sustained fall in the conditional variance of GDP growth.
- Volatility returns briefly at the very end of the sample, around the 2007–09 financial crisis and Great Recession, when the bands widen again.

(c) Is the Great Moderation Evident?

Yes. The GARCH conditional-standard-deviation bands narrow considerably in the early 1980s, providing clear evidence of a decrease in the volatility of GDP growth. The width of the $\pm\,\hat{\sigma}_t$ bands drops sharply around 1983 and stays low for the next two decades, which is precisely the pattern that the term *Great Moderation* describes.

Key finding: The estimated GARCH(1,1) bands are wide through the 1960s and 1970s and then narrow markedly in the early 1980s. This drop in the conditional standard deviation of $GDPGR_t$ is the Great Moderation, the sustained fall in US macroeconomic volatility from roughly 1983 onward.

Reading the economics from the bands:

- The narrowing is not a one-off quiet quarter but a persistent regime of lower volatility, consistent with the high estimated persistence ($\hat{\alpha}_1 + \hat{\phi}_1 = 0.98$) found in part (a): once volatility falls, the GARCH dynamics keep it low.
- The episode is visible directly in the residual plot, without any need to split the sample by hand. The time-varying variance model lets the data reveal the change in volatility on its own.

- The brief widening of the bands near the end of the sample (the 2007–09 crisis) is a useful reminder that the Great Moderation describes a long calm interval, not a permanent end to macroeconomic volatility.

Summary: Key Takeaways

Question	Finding	Interpretation
AR(2) mean	$1.55 + 0.30\,L_1 + 0.20\,L_2$	Growth is positively persistent
GARCH variance	$0.41 + 0.22\,u_{t-1}^2 + 0.76\,\sigma_{t-1}^2$	Volatility clusters and persists
$\hat{\alpha}_1 + \hat{\phi}_1$	0.98	Stationary but very persistent
$\pm\,\hat{\sigma}_t$ bands	Wide pre-1983, narrow after	Volatility falls sharply
Great Moderation?	**Yes**	Bands narrow in the early 1980s

The Big Picture

This exercise illustrates several core ideas about modelling time-varying volatility:

1. **Two equations, one model.** An AR(2)–GARCH(1,1) model describes both the conditional *mean* (how growth depends on its own recent past) and the conditional *variance* (how risk changes over time). Estimating them jointly by maximum likelihood gives sharper inference than treating the variance as constant.

2. **Volatility clustering.** The large GARCH coefficient ($\hat{\phi}_1 = 0.76$) and the near-unit sum $\hat{\alpha}_1 + \hat{\phi}_1 = 0.98$ capture the empirical fact that calm and turbulent periods come in runs. This is a defining feature of most macroeconomic and financial series.

3. **Volatility as a story.** Plotting $\pm\,\hat{\sigma}_t$ bands turns the estimated variance into something you can read off a chart. The bands make the Great Moderation visible: a clear, sustained narrowing from the early 1980s, with a temporary return of volatility during the 2007–09 crisis.

4. **Structural change without a hand-picked break.** Rather than splitting the sample at an arbitrary date, the GARCH model lets the data signal the change in volatility itself, which is a more honest way to detect a shift such as the Great Moderation.

Calculations performed in Stata (`arch` command) and cross-checked in Python (`arch` package).