Financial Econometrics
Shanghai Jiao Tong University
Summer School – 2026
Empirical Exercise 14.1
With Solutions

Fatih Kansoy

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

Data: `us_macro_quarterly.csv` *(1955:Q1–2017:Q4). S&W's price index* `PCEP` *is the column* `PCECTPI`*. Use the sample period 1963:Q1–2012:Q4 (earlier data may serve as initial values for lags).*

Empirical Exercise 14.1. The data file `us_macro_quarterly.csv` contains quarterly data on several macroeconomic series for the United States over 1955:Q1–2017:Q4. The variable `PCEP` (stored in the column `PCECTPI`) is the price index for personal consumption expenditures from the U.S. National Income and Product Accounts. In this exercise you will construct forecasting models for the rate of inflation, based on `PCEP`. For this analysis, use the sample period 1963:Q1–2012:Q4 (where data before 1963 may be used, as necessary, as initial values for lags in regressions).

(a) (i) Compute the inflation rate, $\mathit{Infl}_t = 400\,[\ln(PCEP_t) - \ln(PCEP_{t-1})]$. What are the units of *Infl*? (Is *Infl* measured in dollars, percentage points, percentage per quarter, percentage per year, or something else? Explain.)

(ii) Plot the value of *Infl* from 1963:Q1 through 2012:Q4. Based on the plot, do you think that *Infl* has a stochastic trend? Explain.

(b) (i) Compute the first four autocorrelations of $\Delta\mathit{Infl}$.

(ii) Plot the value of $\Delta\mathit{Infl}$ from 1963:Q1 through 2012:Q4. The plot should look "choppy" or "jagged". Explain why this behaviour is consistent with the first autocorrelation that you computed in part (i).

(c) (i) Run an OLS regression of $\Delta Infl_t$ on $\Delta Infl_{t-1}$. Does knowing the change in inflation this quarter help predict the change in inflation next quarter? Explain.

(ii) Estimate an AR(2) model for $\Delta Infl$. Is the AR(2) model better than an AR(1) model? Explain.

(iii) Estimate an AR(p) model for $p = 0, \ldots, 8$. What lag length is chosen by BIC? What lag length is chosen by AIC?

(iv) Use the AR(2) model to predict the change in inflation from 2012:Q4 to 2013:Q1, that is, predict the value of $\Delta Infl_{2013:Q1}$.

(v) Use the AR(2) model to predict the level of the inflation rate in 2013:Q1, that is, $Infl_{2013:Q1}$.

(d) (i) Use the ADF test for the regression in Equation (14.31) with two lags of $\Delta Infl$ to test for a stochastic trend in $Infl$.

(ii) Is the ADF test based on Equation (14.31) preferred to the test based on Equation (14.32) for testing for a stochastic trend in $Infl$? Explain.

(iii) In (i) you used two lags of $\Delta Infl$. Should you use more lags? Fewer lags? Explain.

(iv) Based on the test you carried out in (i), does the AR model for $Infl$ contain a unit root? Explain carefully. (Hint: does the failure to reject a null hypothesis mean that the null hypothesis is true?)

(e) Use the QLR test with 15% trimming to test the stability of the coefficients in the AR(2) model for $\Delta Infl$. Is the AR(2) model stable? Explain.

(f) (i) Using the AR(2) model for $\Delta Infl$ with a sample period that begins in 1963:Q1, compute pseudo out-of-sample forecasts for the change in inflation beginning in 2003:Q1 and going through 2012:Q4 (that is, compute $\Delta Infl_{2003:Q1\,|\,2002:Q4}$, $\Delta Infl_{2003:Q2\,|\,2003:Q1}, \ldots, \Delta Infl_{2012:Q4\,|\,2012:Q3}$).

(ii) Are the pseudo out-of-sample forecasts biased? That is, do the forecast errors have a nonzero mean?

(iii) How large is the RMSFE of the pseudo out-of-sample forecasts? Is this consistent with the AR(2) model for $\Delta Infl$ estimated over the 1963:Q1–2002:Q4 sample period?

(iv) There is a large outlier in 2008:Q4. Why did inflation fall so much in 2008:Q4? (Hint: collect some data on oil prices. What happened to oil prices during 2008?)

Solution:

Data Overview

The analysis uses quarterly U.S. macroeconomic data from `us_macro_quarterly.csv`, which spans 1955:Q1–2017:Q4. The relevant series is the personal consumption expenditures price index `PCEP`, stored in the column `PCECTPI`. The estimation sample is 1963:Q1–2012:Q4, with observations before 1963 used only to supply initial values for the lagged regressors.

The quarterly inflation rate, expressed at an annual rate, is constructed as

$$\boxed{\mathit{Infl}_t = 400\,[\ln(PCEP_t) - \ln(PCEP_{t-1})]}$$

and its quarter-to-quarter change is $\Delta \mathit{Infl}_t = \mathit{Infl}_t - \mathit{Infl}_{t-1}$. The factor of 400 converts the fractional change of the price index (multiply by 100 for percentage points) into an annual rate (multiply by 4 because the data are quarterly).

Note: The analysis distinguishes the *level* of inflation Infl, which turns out to be highly persistent, from its *change* $\Delta \mathit{Infl}$, which is the stationary series that the autoregressions are fitted to.

(a) Constructing and Inspecting the Inflation Rate

(i) Units of *Infl*

Infl is measured in **percentage points at an annual rate**. The difference of logarithms $\ln(PCEP_t) - \ln(PCEP_{t-1})$ is the (approximate) fractional change in the price index over one quarter. Multiplying by 100 turns the fraction into percentage points, and the further factor of 4 (so 400 in total) annualises the quarterly rate, giving the percentage increase that would occur over a full year if prices kept rising at the current quarterly pace.

(ii) Plot of *Infl*, 1963:Q1–2012:Q4

[See Figure: Inflation rate Infl, 1963:Q1–2012:Q4]

Inflation rose over the 20-year period 1960–1980, then declined over the following decade, and has been reasonably stable since then. Because the series shows these long, slow swings and does not revert quickly to a fixed mean, it appears to have a **stochastic trend**.

(b) Autocorrelations and Behaviour of $\Delta \boldsymbol{Infl}$

(i) First four autocorrelations of $\Delta \boldsymbol{Infl}$

Lag j	1	2	3	4
$\hat{\rho}_j$	-0.25	-0.20	0.14	-0.08

(ii) Plot of $\Delta \boldsymbol{Infl}$ and the link to $\hat{\rho}_1$

[See Figure: Change in inflation $\Delta Infl$, 1963:Q1–2012:Q4]

The change in inflation is **slightly negatively serially correlated** (the first autocorrelation is $\hat{\rho}_1 = -0.25$), so that values above the mean tend to be followed by values below the mean. This is exactly what produces the "choppy" or "jagged" appearance of the plot: a quarter in which inflation rises by more than average is, on average, followed by a quarter in which it falls back, so the series repeatedly reverses direction from one quarter to the next.

(c) Autoregressive Models for $\Delta \boldsymbol{Infl}$

(i) AR(1) model

$$\boxed{\widehat{\Delta Infl}_t = 0.007 - 0.247\,\Delta Infl_{t-1}, \qquad R^2 = 0.06}$$

(standard errors: 0.0965 on the intercept, 0.068 on the slope).

The coefficient on lagged $\Delta Infl$ is statistically significant (the slope is -0.247 with a standard error of 0.068, so the t-statistic is roughly -3.6). Knowing the change in inflation this quarter therefore **does help predict** the change in inflation next quarter. The negative coefficient again reflects the partial reversal already seen in part (b).

(ii) AR(2) model

$$\boxed{\widehat{\Delta Infl}_t = 0.0061 - 0.318\,\Delta Infl_{t-1} - 0.284\,\Delta Infl_{t-2}, \qquad R^2 = 0.13}$$

(standard errors: 0.0926, 0.064, 0.076 respectively).

The estimated coefficient on $\Delta Infl_{t-2}$ is statistically significant (-0.284 with standard error 0.076), so the **AR(2) model is preferred to the AR(1) model**. Note also that the R^2 rises from 0.06 in the AR(1) model to 0.13 in the AR(2) model, confirming the second lag adds predictive content.

(iii) Lag-length selection by AIC and BIC

p	BIC(p)	AIC(p)
0	0.8049	0.7895
1	0.7669	0.7360
2	**0.7077**	**0.6614**
3	0.7322	0.6705
4	0.7437	0.6665
5	0.7533	0.6608
6	0.7689	0.6609
7	0.7931	0.6697
8	0.8163	0.6774

Both criteria are minimised at $p = 2$, so **AIC and BIC both select two lags**. (BIC has a clear minimum at $p = 2$; AIC dips slightly again at $p = 5$–6 but the smallest value remains at $p = 2$.)

(iv) Predicted change in inflation, 2013:Q1

Using the AR(2) model, the predicted change in inflation from 2012:Q4 to 2013:Q1 is

$$\boxed{\widehat{\Delta Infl}_{2013:Q1} = -0.72 \text{ percentage points.}}$$

(v) Predicted level of inflation, 2013:Q1

Adding the predicted change to the level of inflation in 2012:Q4 gives the forecast level

$$\boxed{\widehat{Infl}_{2013:Q1} = 1.96 \text{ percentage points (annual rate).}}$$

(d) Testing for a Stochastic Trend (Unit Root)

(i) ADF test with two lags

The augmented Dickey–Fuller (ADF) t-statistic for the regression with an intercept and two lags of $\Delta Infl$ is

$$\boxed{\text{ADF } t = -2.74.}$$

The relevant critical values for the intercept-only specification are -2.57 at the 10% level and -2.86 at the 5% level. Since -2.74 is more negative than -2.57 but not more negative than -2.86, the **unit-root null hypothesis can be rejected at the 10% level but not at the 5% level**.

ADF t-statistic	10% critical value	5% critical value	Decision
-2.74	-2.57	-2.86	Reject at 10%, not at 5%

(ii) Which specification?

The specification **with an intercept but no time trend is appropriate**. The inflation rate does not exhibit a deterministic linear trend (it rises and then falls back over the sample), so the ADF regression should not include a deterministic time trend. The intercept-only specification of Equation (14.31) is therefore preferred to the trend specification of Equation (14.32).

(iii) Number of lags

Two lags is a reasonable choice. In part (c) both AIC and BIC selected $p = 2$ for $\Delta Infl$, and the same information criteria here **also choose two lags**. There is no need to add more lags, and using fewer would risk leaving serial correlation in the residuals, which would distort the ADF test.

(iv) Does the AR model contain a unit root?

We cannot conclude that the model contains a unit root. The test *fails to reject* the unit-root null at the 5% level, but failure to reject a null hypothesis does **not** mean the null is true. Inflation is highly persistent: its largest autoregressive root is close to 1.0, and given the imprecision of the estimate the null that the root equals 1.0 cannot be rejected at 5%, which is consistent with a true value close to (but not necessarily exactly equal to) 1.0. Ideally one would report a confidence interval for the largest root, although constructing such intervals when the root is near 1.0 is delicate and goes beyond the methods covered here.

(e) QLR Test for Stability of the AR(2) Model

The Quandt likelihood ratio (QLR) statistic, computed with 15% trimming, is

$$\boxed{\text{QLR} = 5.12, \text{ achieved in } 2009 : Q1.}$$

The 5% critical value (for two restrictions tested over the trimmed range) is 4.71. Since $5.12 > 4.71$, the **null hypothesis of coefficient stability is rejected at the 5% level**, so the AR(2) model shows evidence of instability.

However, the break date 2009:Q1 falls immediately after the very large change in inflation in 2008:Q4. This timing suggests that the apparent instability may be driven by that single outlier rather than by a genuine, persistent change in the relationship.

Absent that one extreme observation the model may well be stable, and it would be worth re-running the test to investigate this conjecture.

(f) Pseudo Out-of-Sample Forecasting

(i) Constructing the POOS forecasts

[See Figure: Actual $\Delta Infl$ (black) and pseudo out-of-sample forecasts (blue), 2003:Q1–2012:Q4]

The AR(2) model is re-estimated recursively, starting from a sample that begins in 1963:Q1, and one-quarter-ahead forecasts $\Delta Infl_{t\,|\,t-1}$ are produced for every quarter from 2003:Q1 through 2012:Q4. Each forecast uses only data available up to the previous quarter, mimicking the position of a real-time forecaster.

(ii) Are the forecasts biased?

	Mean forecast error	Standard error
POOS forecast errors	0.00	0.23

The sample mean of the forecast errors is 0.00 (to two decimal places) with a standard error of 0.23. The mean is far from significant relative to its standard error, so there is **no significant bias** in the pseudo out-of-sample forecasts.

(iii) RMSFE

The root mean squared forecast error of the pseudo out-of-sample forecasts is

$$\boxed{\text{RMSFE} = 1.80.}$$

The in-sample regression estimated over 1963:Q1–2002:Q4 has a standard error of 1.20, so the out-of-sample RMSFE (1.80) is **larger than the in-sample regression would lead one to expect**. As part (iv) explains, much of this gap is due to the outlier in 2008:Q4, which produced a single forecast error of about -9.7%. Excluding that one observation reduces the RMSFE to 1.29, which is much closer to the in-sample SER of 1.20.

	RMSFE	Comparison
All POOS quarters	1.80	Above in-sample SER (1.20)
Excluding 2008:Q4	1.29	Close to in-sample SER (1.20)

(iv) The 2008:Q4 outlier

Inflation fell sharply in 2008:Q4 because **oil prices plummeted** at the end of 2008. During the global financial crisis crude oil prices collapsed from their mid-2008 peak (around \$140 per barrel) to well below \$50 per barrel within a few months. Because energy enters the consumption basket directly, this rapid fall in oil prices pulled headline PCE inflation down abruptly, producing the large negative forecast error that the AR(2) model, which uses only past inflation, could not anticipate.

Summary: Key Takeaways

Part	Finding	Interpretation
(a) Units / trend	Percentage points, annual rate	*Infl* has a stochastic trend
(b) $\hat{\rho}_1$ of $\Delta Infl$	-0.25	Mild negative serial correlation
(c) AR(1) slope	-0.247 (significant)	Lag helps predict $\Delta Infl$
(c) AR(2) & criteria	$R^2 = 0.13$; AIC, BIC pick $p = 2$	AR(2) preferred
(c) Forecasts	$\Delta Infl = -0.72$; $Infl = 1.96$	2013:Q1 predictions
(d) ADF t	-2.74	Reject unit root at 10%, not 5%
(e) QLR	$5.12 > 4.71$ at 2009:Q1	Instability, likely the 2008:Q4 outlier
(f) RMSFE	1.80 (1.29 without 2008:Q4)	Outlier-driven; no bias

The Big Picture

This exercise illustrates several core ideas in time series forecasting:

1. **Levels versus changes.** The level of inflation is highly persistent (a near unit root), so forecasting is carried out on the stationary first difference $\Delta Infl$ rather than on the level itself.

2. **Information criteria.** AIC and BIC both select an AR(2), trading off in-sample fit against parsimony. Here they agree, which gives extra confidence in the chosen lag length.

3. **Unit-root inference is asymmetric.** Failing to reject the unit-root null does not prove that a unit root is present. The data are simply consistent with a largest root close to 1.0.

4. **Stability and outliers.** The QLR test rejects stability, but the break sits right after the 2008:Q4 oil-price collapse, a reminder that a single extreme observation can masquerade as a structural break.

5. **Out-of-sample honesty.** Pseudo out-of-sample forecasting gives a realistic picture of forecast accuracy. The forecasts are unbiased, but the RMSFE (1.80) exceeds the in-sample SER (1.20) almost entirely because of the 2008:Q4 oil shock; removing it brings the two measures back into line.

Calculations performed using Stata. See accompanying `E14_1_solution.do` *file.*