Financial Econometrics
Shanghai Jiao Tong University
Summer School – 2026
Empirical Exercise 16.1
With Solutions

Fatih Kansoy

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

Data: `us_macro_quarterly.csv`*. Inflation is built from the PCE price index* `PCEP` = `PCECTPI` *as* $\mathit{Infl}_t = 400\,[\ln(PCEP_t) - \ln(PCEP_{t-1})]$*. Sample period 1963:Q1–2012:Q4.*

Empirical Exercise 16.1. This exercise is an extension of Empirical Exercise 14.1. Using the data file `us_macro_quarterly.csv`, compute inflation from the price index for personal consumption expenditures as

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

where `PCEP` = `PCECTPI`. For all regressions use the sample period 1963:Q1–2012:Q4 (data before 1963 may be used as initial values for the lags in the regressions).

(a) Using the data on inflation through 2012:Q4 and an estimated AR(2) model:

(i) Forecast $\Delta \mathit{Infl}_{2013:\mathrm{Q1}}$, the change in inflation from 2012:Q4 to 2013:Q1.

(ii) Forecast $\Delta \mathit{Infl}_{2013:\mathrm{Q2}}$, the change in inflation from 2013:Q1 to 2013:Q2. (Use an iterated forecast.)

(iii) Forecast $\mathit{Infl}_{2013:\mathrm{Q2}} - \mathit{Infl}_{2012:\mathrm{Q4}}$, the change in inflation from 2012:Q4 to 2013:Q2.

(iv) Forecast $\mathit{Infl}_{2013:\mathrm{Q2}}$, the level of inflation in 2013:Q2.

(b) Repeat (a) using the direct forecasting method.

(c) In Empirical Exercise 14.1 you carried out an ADF test for a unit root in the autoregression for `Infl`. Now carry out the unit-root test using the DF-GLS test. Are the conclusions based on the DF-GLS test the same as those you reached using the ADF test? Explain.

Solution:

Data Overview

The series of interest is quarterly United States inflation, built from the personal consumption expenditures price index `PCEP` (the variable `PCECTPI` in `us_macro_quarterly.csv`). The quarterly inflation rate, expressed in percentage points at an annual rate, is

$$\mathit{Infl}_t = 400\left[\ln(PCEP_t) - \ln(PCEP_{t-1})\right].$$

The factor 400 converts the fractional quarterly log change into a percentage (the factor 100) at an annual rate (the further factor 4). All regressions use 1963:Q1–2012:Q4; the few quarters before 1963 supply initial values for the lagged regressors.

Quantity	**Symbol**	**Role**
PCE price index	`PCEP` (`PCECTPI`)	raw input
Inflation (annualised)	Infl_t	analysed series
Change in inflation	$\Delta \mathit{Infl}_t = \mathit{Infl}_t - \mathit{Infl}_{t-1}$	forecast target

Note: This exercise extends Empirical Exercise 14.1. There you treated inflation as plausibly having a random-walk stochastic trend, that is, you treated `Infl` as I(1) on the basis of an ADF test. The change in inflation $\Delta \mathit{Infl}_t$ is therefore modelled as stationary, and the AR(2) for $\Delta \mathit{Infl}_t$ is the one-period-ahead model used below.

Key Notation and Formulas

The one-period-ahead model is an AR(2) for the change in inflation,

$$\Delta \mathit{Infl}_t = \beta_0 + \beta_1\, \Delta \mathit{Infl}_{t-1} + \beta_2\, \Delta \mathit{Infl}_{t-2} + u_t.$$

Iterated forecast. The one-step forecast is computed first, then fed back in as data to obtain the two-step forecast:

$$\widehat{\Delta \mathit{Infl}}_{T+1|T} = \hat\beta_0 + \hat\beta_1\, \Delta \mathit{Infl}_T + \hat\beta_2\, \Delta \mathit{Infl}_{T-1},$$

$$\widehat{\Delta \mathit{Infl}}_{T+2|T} = \hat\beta_0 + \hat\beta_1\, \widehat{\Delta \mathit{Infl}}_{T+1|T} + \hat\beta_2\, \Delta \mathit{Infl}_T.$$

Direct forecast. The two-quarter horizon is obtained from a single regression in which all predictors are lagged by the extra horizon,

$$\Delta \mathit{Infl}_t = \delta_0 + \delta_1 \, \Delta \mathit{Infl}_{t-2} + \delta_2 \, \Delta \mathit{Infl}_{t-3} + u_t,$$

estimated with Newey–West (HAC) standard errors, because overlapping multiperiod errors are serially correlated.

Recovering the level. Because $\Delta \mathit{Infl}_t = \mathit{Infl}_t - \mathit{Infl}_{t-1}$, forecasts of the level are obtained by cumulating the forecast changes onto the last observed level:

$$\widehat{\mathit{Infl}}_{2013:\mathrm{Q2}|2012:\mathrm{Q4}} = \mathit{Infl}_{2012:\mathrm{Q4}} + \big(\widehat{\Delta \mathit{Infl}}_{2013:\mathrm{Q1}} + \widehat{\Delta \mathit{Infl}}_{2013:\mathrm{Q2}}\big).$$

(a) Iterated AR(2) forecasts

The AR(2) for $\Delta \mathit{Infl}_t$ is estimated over 1963:Q1–2012:Q4 and used to forecast forward from the end of the sample.

(i) One-quarter-ahead change, $\Delta \boldsymbol{Infl}_{\mathbf{2013:Q1}}$

The one-step forecast of the change in inflation from 2012:Q4 to 2013:Q1 is obtained directly from the fitted AR(2) using the last two observed changes:

$$\boxed{\widehat{\Delta \mathit{Infl}}_{2013:\mathrm{Q1}|2012:\mathrm{Q4}} = 0.721}$$

Inflation is forecast to rise by about 0.72 percentage points (annual rate) over the first quarter of 2013.

(ii) Two-quarter-ahead change, $\Delta \boldsymbol{Infl}_{\mathbf{2013:Q2}}$ (iterated)

The one-step forecast from (i) is now treated as data and substituted back into the AR(2) to forecast the change from 2013:Q1 to 2013:Q2:

$$\boxed{\widehat{\Delta \mathit{Infl}}_{2013:\mathrm{Q2}|2012:\mathrm{Q4}} = 0.098}$$

The further change is small and positive, so the iterated path settles down quickly, which is typical of a stationary AR.

(iii) Two-quarter cumulative change, $\boldsymbol{Infl}_{\mathbf{2013:Q2}} - \boldsymbol{Infl}_{\mathbf{2012:Q4}}$

The change in inflation from 2012:Q4 to 2013:Q2 is the sum of the two forecast quarterly changes:

$$0.721 + 0.098 = \boxed{0.819}.$$

(iv) Level of inflation, $\boldsymbol{Infl}_{\mathbf{2013:Q2}}$

Adding the cumulative change to the last observed level of inflation gives

$$\boxed{\widehat{\mathit{Infl}}_{2013:\mathrm{Q2}|2012:\mathrm{Q4}} = 1.87}$$

percentage points at an annual rate.

Forecast target	**Symbol**	**Iterated value**
Change to 2013:Q1	$\Delta \mathit{Infl}_{\mathrm{2013:Q1}}$	0.721
Change to 2013:Q2	$\Delta \mathit{Infl}_{\mathrm{2013:Q2}}$	0.098
Cumulative change Q4→Q2	$\mathit{Infl}_{\mathrm{2013:Q2}} - \mathit{Infl}_{\mathrm{2012:Q4}}$	0.819
Level in 2013:Q2	$\mathit{Infl}_{\mathrm{2013:Q2}}$	1.87

(b) Direct forecasts

The direct method replaces the iterated AR(2) by a single multiperiod regression for each horizon, with the predictors lagged by the forecast horizon and HAC standard errors. The one-quarter forecast coincides with the iterated AR(2), because at horizon one the direct and iterated models are identical.

(i) One-quarter-ahead change, $\Delta \boldsymbol{Infl}_{2013:\mathbf{Q}1}$

$$\boxed{\widehat{\Delta Infl}_{2013:\mathrm{Q}1|2012:\mathrm{Q}4} = 0.721}$$

(ii) Two-quarter-ahead change, $\Delta \boldsymbol{Infl}_{2013:\mathbf{Q}2}$ (direct)

Here the direct estimate differs slightly from the iterated one, because a separate two-quarter-ahead regression is fitted rather than iterating the one-step model:

$$\boxed{\widehat{\Delta Infl}_{2013:\mathrm{Q}2|2012:\mathrm{Q}4} = 0.090}$$

(iii) Two-quarter cumulative change, $\boldsymbol{Infl}_{2013:\mathbf{Q}2} - \boldsymbol{Infl}_{2012:\mathbf{Q}4}$

$$0.721 + 0.090 = \boxed{0.811}.$$

(iv) Level of inflation, $\boldsymbol{Infl}_{2013:\mathbf{Q}2}$

$$\boxed{\widehat{Infl}_{2013:\mathrm{Q}2|2012:\mathrm{Q}4} = 1.87}$$

Comparison: iterated versus direct

Forecast target	**Iterated (a)**	**Direct (b)**
Change to 2013:Q1	0.721	0.721
Change to 2013:Q2	0.098	0.090
Cumulative change Q4→Q2	0.819	0.811
Level in 2013:Q2	1.87	1.87

The two methods give the same one-quarter forecast and almost the same two-quarter forecast (0.098 against 0.090). The cumulative change differs only at the second decimal (0.819 against 0.811), and once rounded to the precision of the published inflation series the forecast level is the same, 1.87. This close agreement is what we expect when the one-step AR(2) is well specified: the theory in this chapter recommends the iterated method in such cases, because its coefficients are estimated more efficiently and its forecast paths are smoother across horizons.

(c) DF-GLS unit-root test versus the ADF test

Set-up. The DF-GLS test of Elliott, Rothenberg and Stock (1996) tests the null hypothesis that `Infl` has a unit autoregressive root, that is, that inflation is I(1), against the stationary alternative. It is computed in two steps. First, the deterministic terms (here an intercept) are removed by generalised least squares to produce a detrended series Infl_t^d. Second, a Dickey–Fuller regression with no intercept and no trend is run on the detrended series, and the t-statistic on the lagged level is the test statistic. The lag length is chosen, as in Empirical Exercise 14.1, by AIC or BIC.

Critical values. Because the deterministic terms are handled differently from the ADF test, the DF-GLS test uses its own critical values (Table 16.1). For the intercept-only version, the 10%, 5% and 1% critical values are −1.62, −1.95 and −2.58. The null of a unit root is rejected when the statistic is more negative than the critical value.

Deterministic terms	**10%**	**5%**	**1%**
Intercept only	−1.62	−1.95	−2.58
Intercept and time trend	−2.57	−2.89	−3.48

Conclusion. The DF-GLS test leads to the **same conclusion** as the ADF test carried out in Empirical Exercise 14.1: the null hypothesis of a unit root in `Infl` is not rejected at the usual significance levels. Inflation is therefore treated as I(1), so its stochastic trend is removed by first differencing and $\Delta \mathit{Infl}_t$ is the stationary series used for forecasting in parts (a) and (b).

Why the agreement is reassuring. The DF-GLS test has higher power than the ADF test, that is, it is better able to reject the unit-root null when the true root is large but less than one (for an AR(1) with coefficient 0.95 and $T = 200$, the rejection probabilities are about 31% for ADF against about 75% for DF-GLS). Because the more powerful DF-GLS test still fails to reject, the finding that inflation behaves like an I(1) series over this sample is not simply an artefact of the lower-power ADF test. Both tests point the same way, which strengthens the case for modelling inflation in differences.

Summary: Key Takeaways

Forecast target	Iterated	Direct
$\Delta \mathit{Infl}_{2013:\mathrm{Q1}}$	0.721	0.721
$\Delta \mathit{Infl}_{2013:\mathrm{Q2}}$	0.098	0.090
$\mathit{Infl}_{2013:\mathrm{Q2}} - \mathit{Infl}_{2012:\mathrm{Q4}}$	0.819	0.811
$\mathit{Infl}_{2013:\mathrm{Q2}}$ (level)	1.87	1.87
Unit-root test on `Infl`	DF-GLS does not reject; same as ADF (I(1))	

The Big Picture

This exercise illustrates several ideas from Chapter 16:

1. **Multiperiod forecasting.** A two-quarter forecast can be built by iterating a one-step AR(2) (feeding each forecast back as data) or by estimating a single direct regression at the two-quarter horizon. Both are legitimate; the iterated method is the default recommendation.

2. **Iterated versus direct agreement.** The two methods give the same one-quarter forecast and very close two-quarter forecasts (changes of 0.098 versus 0.090, cumulative 0.819 versus 0.811, level 1.87 either way). Close agreement is the signature of a well-specified one-step model, in which the iterated forecast is the more efficient choice.

3. **From changes to levels.** Because inflation is treated as I(1), the AR is fitted to the change $\Delta \mathit{Infl}_t$; forecasts of the level are recovered by cumulating the forecast changes onto the last observed level.

4. **Power and robustness of unit-root tests.** The DF-GLS test is more powerful than the ADF test, yet it reaches the same verdict here: it does not reject a unit root in inflation. Agreement between a low-power and a high-power test makes the I(1) treatment of inflation more credible.

Calculations performed in Python/Stata using `us_macro_quarterly.csv` over 1963:Q1–2012:Q4.