

# Time Series Regression and Forecasting

---

Juergen Meinecke

Practical Guidelines and Example

Practical Guidelines for Time Series Analysis

Suppose I give you a time series  $Y_t$  and ask you to produce forecasts, what should you do?

Let's try to tie up all loose ends and come up with some practical advice

Given the time series  $Y_t$ , here's what you should do to produce the best forecast...

- Graph it

Sometimes you can spot deterministic trends in graphs (e.g., GDP)

A graph will definitely not help you spot a stochastic trend

- Run an ADF unit root test on  $Y_t$

- If you reject the null hypothesis of a unit root:

Estimate the AR(p) model for  $Y_t$ :

$$Y_t = \beta_0 + \alpha t + \beta_1 Y_{t-1} + \beta_2 Y_{t-2} + \dots + \beta_p Y_{t-p} + u_t$$

(It's always safe to include a deterministic time trend)

You can easily estimate this in **Python**

Use the AR(p) model for  $Y_t$  to produce your forecasts

- If you cannot reject the null hypothesis of a unit root:  
Generate first differences  $\Delta Y_t$  of the data and run a unit root test on  $\Delta Y_t$

In 90% of practical situations, you will reject the null hypothesis of a unit root in  $\Delta Y_t$  (in other words, first differencing effectively removes the unit root)

This suggests that you should use an AR(p) model for  $\Delta Y_t$ :

$$\Delta Y_t = \beta_0 + \alpha t + \beta_1 \Delta Y_{t-1} + \beta_2 \Delta Y_{t-2} + \dots + \beta_p \Delta Y_{t-p} + u_t$$

Since  $\Delta Y_t$  is stationary, you can easily estimate this in **Python**

Use the AR(p) model for  $\Delta Y_t$  to produce your forecasts

# Time Series Regression and Forecasting

---

Juergen Meinecke

Practical Guidelines and Example

Example: CPI and Inflation



During the last two weeks I have used inflation as the running example to illustrate important concepts

Recall that we defined inflation to be

$$\text{infl}_t := 400 \cdot (\ln(\text{cpi}_t) - \ln(\text{cpi}_{t-1}))$$


(this is the *annualized* version of quarter-to-quarter inflation; but it would work just the same if we hadn't used the annualized version)

Our goal:

Come up with the best AR(p) model to forecast inflation

Let's go through the steps

## Graphing the data



Deterministic trend: not obvious

Stochastic trend: never obvious

Test for unit root

A few slides ago...

## Python Code

```
> from statsmodels.tsa.stattools import adfuller
> adfuller(df.infl.dropna(), maxlag = 4, regression='ct')
(-4.394580052261824,
 0.00223369411475301,
 4,
 303,
 '1%': -3.9889626832421174,    '5%': -3.4250798217949754,    '10%': -3.135626889644655,
 1316.468148600885)
```

Dickey Fuller critical value is -3.41

Reject the null hypothesis of unit root here

Therefore, our best time series model is this one:

$$\begin{aligned} \text{infl}_t = & \beta_0 + \alpha t + \beta_1 \text{infl}_{t-1} + \beta_2 \text{infl}_{t-2} \\ & + \beta_3 \text{infl}_{t-3} + \beta_4 \text{infl}_{t-4} + u_t \end{aligned}$$

Let's estimate it...

## Python Code (output edited)

```
> ar4trend = smf.ols('infl ~ l1infl + l2infl + l3infl + l4infl + trend',  
                    data=df, missing='drop').fit(use_t=False)  
> print(ar4trend.summary())
```

### OLS Regression Results

```
=====
```

|                   |               |                 |       |
|-------------------|---------------|-----------------|-------|
| Dep. Variable:    | infl          | R-squared:      | 0.560 |
| Model:            | OLS           | Adj. R-squared: | 0.552 |
| Method:           | Least Squares | F-statistic:    | 75.77 |
| No. Observations: | 304           |                 |       |
| Df Residuals:     | 298           |                 |       |
| Df Model:         | 5             |                 |       |
| Covariance Type:  | nonrobust     |                 |       |

```
=====
```

|           | coef      | std err | z      | P> z  | [0.025 | 0.975] |
|-----------|-----------|---------|--------|-------|--------|--------|
| Intercept | 0.7414    | 0.299   | 2.479  | 0.013 | 0.155  | 1.328  |
| l1infl    | 0.6183    | 0.057   | 10.826 | 0.000 | 0.506  | 0.730  |
| l2infl    | 0.0083    | 0.065   | 0.128  | 0.898 | -0.118 | 0.135  |
| l3infl    | 0.3126    | 0.065   | 4.836  | 0.000 | 0.186  | 0.439  |
| l4infl    | -0.1671   | 0.056   | -2.958 | 0.003 | -0.278 | -0.056 |
| trend     | 9.302e-05 | 0.001   | 0.067  | 0.946 | -0.003 | 0.003  |

```
=====
```

What should we do with these estimates now?

Reminder:

the time series `CPIAUCSL` ran from 1947:Q1 to 2024:Q1

That's 309 observations on `CPIAUCSL`

That's 308 observations on `infl`

Remember that our main goal is to produce forecasts

Let's produce an inflation forecast for 2024:Q2

It's very easy given our AR(4) estimates

Recall the observed data on inflation

## Python Code

```
> df.infl.tail(5)
quarter
2023Q1    3.684879
2023Q2    2.994927
2023Q3    3.370209
2023Q4    2.689546
2024Q1    3.735410
```

Obtain forecast by plugging into the estimated model:

$$\begin{aligned}\widehat{\text{infl}}_{2024:Q2|2024:Q1} &= 0.7414 + 0.00009302 \cdot 308 \\ &\quad + 0.6183 \cdot (3.7341) + 0.0083 \cdot (2.6895) \\ &\quad + 0.3126 \cdot (3.3702) - 0.1671 \cdot (2.9949) \\ &= 3.6542\end{aligned}$$

# Time Series Regression and Forecasting

---

Juergen Meinecke



Practical Guidelines and Example  
Pseudo Out-of Sample Forecasts

The CPIAUCSL time series runs from 1947:Q1 to 2024:Q1

To measure forecast performance, we could follow this little algorithm:

1. pretend like your data already ends in  $s=2005:Q1$
2. estimate an AR(4) model for  $\text{infl}_t$  for the time frame 1947:Q1 to  $s$
3. calculate a forecast for period  $s + 1$ :  $\widehat{\text{infl}}_{s+1|s}$
4. compare the forecast to the actual realization  $\text{infl}_{s+1}$ ; difference is the forecast error:  $\text{infl}_{s+1} - \widehat{\text{infl}}_{s+1|s}$
5. shift  $s$  up by one period and jump back to step 2

Repeat until you reach end of your data set ( $s \leq T$ ), with  $T := 2024:Q1$

This algorithm creates a shifting or rolling one-quarter pseudo out-of-sample forecast

So we obtain lots of one-quarter forecasts of inflation for the time period 2005:Q2 to 2024:Q1


This puts us in a position to compare our one-quarter forecasts to the *actual* inflation numbers

This is a bit complicated to code in **Python**, so I won't ask you to do it yourselves

Nevertheless, I have implemented this algorithm in **Python** and ran it over the inflation data


Here's the output in two pictures...

# Inflation: actual versus forecast



Forecast error

(vertical difference between two lines from previous slide)



Is this a good model or a bad model?

This simple AR(4) model already does a reasonable job  
And it took you only two and a half weeks of time series  
econometrics to be able to do this!

Where from here?

You could try different AR( $p$ ) models  
(What's the optimal 'choice' of  $p$ ?)

We could use an estimate of the so-called mean square forecast  
error (MSFE) to judge which model produces the best forecasts

You will learn much more on time series regression and forecasting  
in courses like EMET3007 or EMET3008

If you care more about the micro-econometric applications that we  
covered in weeks 1-9 of the semester, then EMET3004 or EMET3006  
are right for you