Department of Economics,
Management and Quantitative Methods
B-74-3-B Time Series Econometrics
Academic year 2019-2020

Computer Session 3 - Forecasting US Inflation

In this exercise, we forecast the US inflation at quarterly data, defined as the percent variation of GDP deflator over the quarter, at annual rate (i.e., four times the change over each quarter), and we compare our forecast with forecasts from the Survey of Professional Forecasters. Both observed and professionally forecast series are in the file 'gdp price deflator.wf1'. The series of Observed inflation is called 'true' in the worksheet.

We have data spanning 1968Q4-2015Q1, see picture

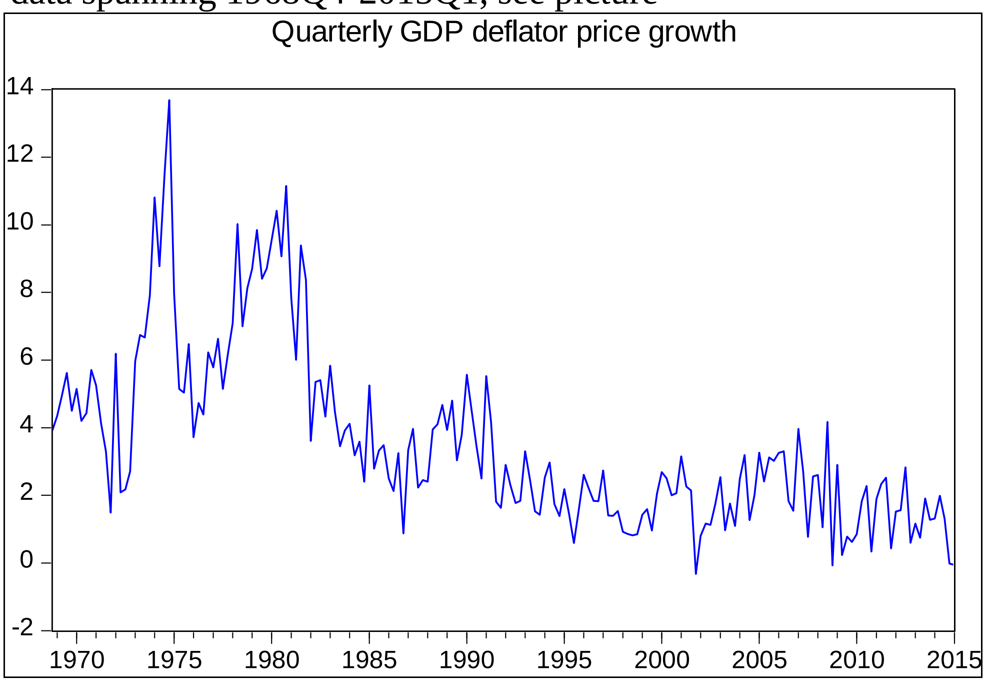

1) We discard observations spanning 1968-1984 as it is possible that inflation changed dynamics after 1985 (for example, in response to a change in monetary policy).
2) We estimated the model using the period 1985-2010
3) We check our forecasts over the period 2011-2014 (we also compare our forecasts against the forecasts from the Survey of Professional Forecasters).

1 Model Selection and Estimation

1.1 Set the sample to 1985-2010 (write '1985q1 2010q4' in the box *Sample*) (this is default in the original file).

1.2 Unit root test

Preliminary investigation with a Unit root test suggests the possibility of a unit root (test: ADF, Case 2, lags selected using the BIC).

Null Hypothesis: TRUE has a unit root
Exogenous: Constant
Lag Length: 3 (Automatic - based on SIC, maxlag=4)

		t-Statistic	Prob.*
Augmented Dickey-Fuller test statistic		-2.092440	0.2482
Test critical values:	1% level	-3.494378	
	5% level	-2.889474	
	10% level	-2.581741	

*MacKinnon (1996) one-sided p-values.

Augmented Dickey-Fuller Test Equation
Dependent Variable: D(TRUE)
Method: Least Squares
Sample: 1985Q1 2010Q4
Included observations: 104

Variable	Coefficient	Std. Error	t-Statistic	Prob.
TRUE(-1)	-0.229036	0.109459	-2.092440	0.0390
D(TRUE(-1))	-0.592641	0.124049	-4.777469	0.0000
D(TRUE(-2))	-0.453123	0.119946	-3.777713	0.0003
D(TRUE(-3))	-0.240537	0.097991	-2.454693	0.0158
C	0.488868	0.274206	1.782849	0.0777

R-squared	0.416638	Mean dependent var	-0.019863
Adjusted R-squared	0.393068	S.D. dependent var	1.295774
S.E. of regression	1.009484	Akaike info criterion	2.903639
Sum squared resid	100.8867	Schwarz criterion	3.030773
Log likelihood	-145.9892	Hannan-Quinn criter.	2.955144
F-statistic	17.67647	Durbin-Watson stat	1.933521
Prob(F-statistic)	0.000000		

What should we conclude from the unit root test in this case?
Strictly speaking, one cannot expect a unit root on inflation, as it means that the process is not bounded and does not revert to the mean: this clashes with our idea that there is “fair” level for inflation (i.e., the process tends to revert to a mean). The unit root may suggest that in fact the mean is subject to repeated breaks, that makes it look as a unit root, or that the autoregressive term is so strong to be very close to 1. A unit root may, in this case, give better forecasts (it is better to set rho=1 even if it is incorrect than estimating it).

For the sake of the presentation, however, we will discuss both using the model with unit root and without. We begin by discussing the model without unit root.

1.3 Preliminary investigation of the correlogram (model in levels)

Sample: 1985Q1 2010Q4
Included observations: 104

Autocorrelation	Partial Correlation		AC	PAC	Q-Stat	Prob
.|*** |	.|*** |	1	0.416	0.416	18.562	0
.|*** |	.|** |	2	0.376	0.245	33.809	0
.|*** |	.|** |	3	0.421	0.259	53.181	0
.|*** |	.|* |	4	0.416	0.194	72.229	0
.|** |	.|. |	5	0.326	0.039	84.07	0
.|** |	*|. |	6	0.234	-0.076	90.229	0
.|* |	*|. |	7	0.164	-0.129	93.281	0
.|** |	.|* |	8	0.322	0.186	105.18	0
.|* |	*|. |	9	0.122	-0.12	106.92	0
.|. |	*|. |	10	0.025	-0.138	106.99	0
.|* |	.|* |	11	0.188	0.165	111.2	0
.|* |	.|* |	12	0.204	0.128	116.2	0

This seems consistent with an ARMA(1,1).

1.4 Model selection (Model in Level)
We select the model using the Bayes information criterion. We consider up to a ARMA(4,4). (Note: the formula for BIC in e-views is calculated as $-2(l/T)+kln(T)/T$, where l is the maximized log-likelihood and k is the number of parameters). (Note: I used option CLS for the estimation).

	Iid	MA(1)	MA(2)	MA(3)	MA(4)
iid	3.26226	3.161375	3.165379	3.164661	3.118317
AR(1)	3.108767	2.97428	3.007477	3.016271	3.054542
AR(2)	3.076391	3.014657	3.040015	3.080197	3.076566
AR(3)	3.045199	3.043961	3.081695	3.122922	3.118811
AR(4)	3.030773	3.06319	3.093588	3.117743	3.036129

The ARMA(1,1) is indeed selected.

1.5 Estimation and validation
Note: Estimation Method: ARMA Conditional Least Squares (BFGS / Marquardt steps)

(True is the name given to the observed inflation)

Dependent Variable: TRUE
Method: ARMA Conditional Least Squares (BFGS / Marquardt steps)
Sample: 1985Q1 2010Q4
Included observations: 104
Convergence achieved after 21 iterations
Coefficient covariance computed using outer product of gradients
MA Backcast: 1984Q4

Variable	Coefficient	Std. Error	t-Statistic	Prob.
C	1.991494	0.477151	4.173720	0.0001
AR(1)	0.927567	0.054944	16.88201	0.0000
MA(1)	-0.698502	0.106051	-6.586487	0.0000

R-squared	0.314286	Mean dependent var	2.292682
Adjusted R-squared	0.300708	S.D. dependent var	1.214939
S.E. of regression	1.015977	Akaike info criterion	2.898000
Sum squared resid	104.2531	Schwarz criterion	2.974280
Log likelihood	-147.6960	Hannan-Quinn criter.	2.928903
F-statistic	23.14591	Durbin-Watson stat	1.987197
Prob(F-statistic)	0.000000		

Inverted AR Roots	.93
Inverted MA Roots	.70

Portmanteau test on the residuals

Sample: 1985Q1 2010Q4

Included observations: 104

Q-statistic probabilities adjusted for 2 ARMA terms

<table>
<tr><th>Autocorrelation</th><th>Partial Correlation</th><th></th><th>AC</th><th>PAC</th><th>Q-Stat</th><th>Prob</th></tr>
<tr><td>.|. |</td><td>.|. |</td><td>1</td><td>-0.024</td><td>-0.024</td><td>0.0628</td><td></td></tr>
<tr><td>.|. |</td><td>.|. |</td><td>2</td><td>-0.054</td><td>-0.054</td><td>0.3749</td><td></td></tr>
<tr><td>.|* |</td><td>.|* |</td><td>3</td><td>0.085</td><td>0.083</td><td>1.1653</td><td>0.28</td></tr>
<tr><td>.|* |</td><td>.|* |</td><td>4</td><td>0.145</td><td>0.147</td><td>3.4755</td><td>0.176</td></tr>
</table>

Thus, the Portmanteau test on the residuals confirms that the ARMA(1,1) is acceptable, if the series is stationary.

1.6 Preliminary investigation of the correlogram (model in first differences)

D(TRUE)

Sample: 1985Q1 2010Q4

Included observations: 104

<table>
<tr><th>Autocorrelation</th><th>Partial Correlation</th><th></th><th>AC</th><th>PAC</th><th>Q-Stat</th><th>Prob</th></tr>
<tr><td>***|. |</td><td>***|. |</td><td>1</td><td>-0.454</td><td>-0.454</td><td>22.057</td><td>0</td></tr>
<tr><td>*|. |</td><td>***|. |</td><td>2</td><td>-0.082</td><td>-0.362</td><td>22.775</td><td>0</td></tr>
<tr><td>.|. |</td><td>**|. |</td><td>3</td><td>0.034</td><td>-0.263</td><td>22.899</td><td>0</td></tr>
<tr><td>.|. |</td><td>*|. |</td><td>4</td><td>0.071</td><td>-0.104</td><td>23.462</td><td>0</td></tr>
<tr><td>.|. |</td><td>.|. |</td><td>5</td><td>0.004</td><td>0.003</td><td>23.464</td><td>0</td></tr>
<tr><td>.|. |</td><td>.|. |</td><td>6</td><td>-0.026</td><td>0.035</td><td>23.543</td><td>0.001</td></tr>
<tr><td>*|. |</td><td>**|. |</td><td>7</td><td>-0.169</td><td>-0.224</td><td>26.788</td><td>0</td></tr>
<tr><td>.|** |</td><td>.|. |</td><td>8</td><td>0.279</td><td>0.072</td><td>35.733</td><td>0</td></tr>
<tr><td>.|. |</td><td>.|* |</td><td>9</td><td>-0.042</td><td>0.137</td><td>35.942</td><td>0</td></tr>
<tr><td>**|. |</td><td>*|. |</td><td>10</td><td>-0.263</td><td>-0.204</td><td>44.065</td><td>0</td></tr>
<tr><td>.|* |</td><td>*|. |</td><td>11</td><td>0.147</td><td>-0.128</td><td>46.623</td><td>0</td></tr>
<tr><td>.|* |</td><td>.|. |</td><td>12</td><td>0.103</td><td>0.024</td><td>47.891</td><td>0</td></tr>
</table>

From this correlogram, it is very easy to see that we have a MA(1) (there is only one hit on the AC, and several hits for the PAC). The parameter should be negative.

1.7 Estimation and Validation (model in first differences)
We should select the model using the information criterion. However, the investigation of the correlogram (and previous investigation for the model in level) strongly recommend a MA(1), so we skip the selection via information criterion stage.

Dependent Variable: D(TRUE)
Method: ARMA Conditional Least Squares (BFGS / Marquardt steps)
Sample: 1985Q1 2010Q4
Included observations: 104
Convergence achieved after 16 iterations
Coefficient covariance computed using outer product of gradients
MA Backcast: 1984Q4

Variable	Coefficient	Std. Error	t-Statistic	Prob.
C	-0.021719	0.025484	-0.852253	0.3961
MA(1)	-0.755535	0.065017	-11.62056	0.0000

R-squared	0.377804	Mean dependent var	-0.019863
Adjusted R-squared	0.371704	S.D. dependent var	1.295774
S.E. of regression	1.027097	Akaike info criterion	2.910394
Sum squared resid	107.6027	Schwarz criterion	2.961247
Log likelihood	-149.3405	Hannan-Quinn criter.	2.930996
F-statistic	61.93538	Durbin-Watson stat	1.958513
Prob(F-statistic)	0.000000		

Inverted MA Roots	.76

Sample: 1985Q1 2010Q4
Included observations: 104
Q-statistic probabilities adjusted for 1 ARMA term

Autocorrelation	Partial Correlation		AC	PAC	Q-Stat	Prob
.|. |	.|. |	1	-0.008	-0.008	0.0065	
.|. |	.|. |	2	-0.05	-0.05	0.2793	0.597
.|* |	.|* |	3	0.077	0.076	0.9186	0.632
.|* |	.|* |	4	0.129	0.128	2.7426	0.433

The MA(1) fits the data well.

2. Forecasting

We consider four different forecasts, always on the interval 2011Q1-2014Q4.

F0: the forecast from the Survey of Professional Forecasters
F1: the forecast from the ARMA(1,1) assuming stationarity
F2: the forecast from the MA(1) assuming a unit root
F3: a naïve forecast, in which inflation is forecasted as the last available observation

To make the forecast F1, *estimate equation* "*true c AR(1) MA(1)*" with *sample* 1985q1-2010q4, then select the button *Forecast* and, therein, set the sample to 2011q1 2014q4 and option "*static*" and call the series generated in this way F1.
To make the forecast F2, *estimate equation* "*D(true) c MA(1)*" with *sample* 1985q1-2010q4, then select the button *Forecast* and, therein, set the sample to 2011q1 2014q4 and option "*static*" and call the series generated in this way F2.
Finally, generate F3=true(-1).

We then compute the forecast errors e0, e1, e2, e3, corresponding to the errors for the four forecasts for the sample 2011q1 2014q4
For this purpose, set sample 2011q1 2014q4 selecting *Sample* in the *Workfile*, then generate the series i.e., for example, *e0=F0-true* over the sample 2011q1 2014q4

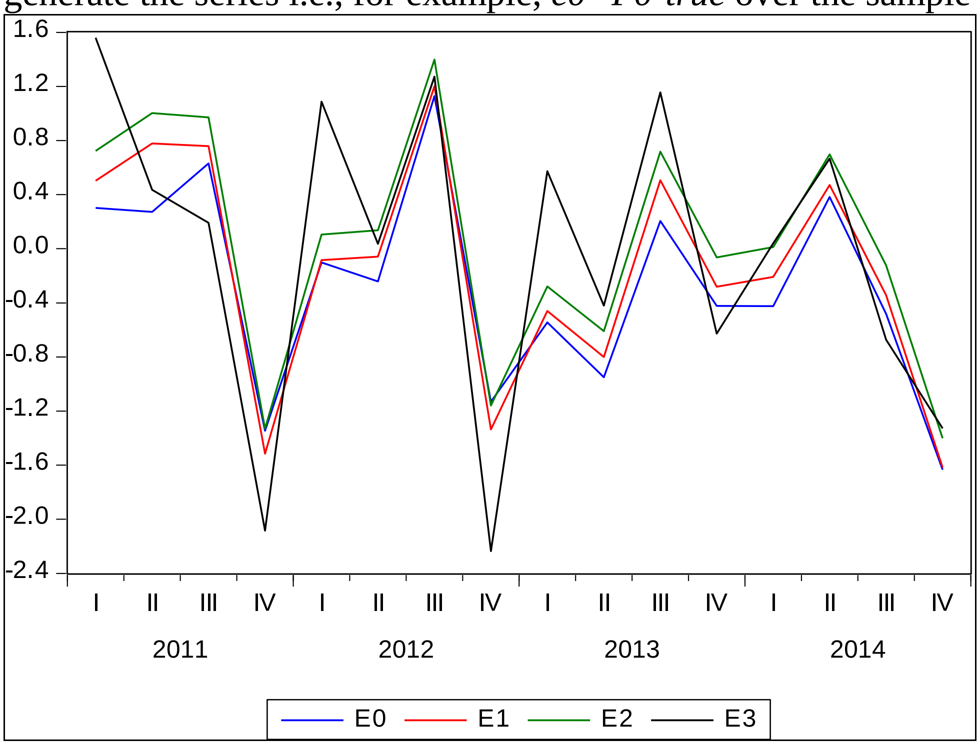

To appreciate more clearly which forecast is better, we compute squares e0^2, e1^2, e2^2, e3^2: we can look at these as measures of precision, as good forecasts have small (in absolute value) errors, and therefore the squares should be also low. We find that the averages of these series, over the sample 2011-2014, are

	E0SQ	E1SQ	E2SQ	E3SQ
Mean	0.60265	0.697612	0.687184	1.231669

We then see that the forecast from the SPF is the most precise one, followed by the one from the model in first difference. The naïve forecast is markedly less precise.

3. Comparing forecasts

Although the SPF forecast is better, the squared errors are quite close to the squared errors of the two ARMA / unit root and MA forecasts. Are the differences statistically significant?

We can compare the SPF and unit root and MA forecasts looking at the differences e0^2-e2^2 at the twelve points in time (2011Q1 to 2014Q4), and check if the difference is significant. We do this by running a regression of the difference e0^2-e2^2 on a constant.

Dependent Variable: E0SQ-E2SQ
Method: Least Squares
Sample: 2011Q1 2014Q4
Included observations: 16
HAC standard errors & covariance (Bartlett kernel, Newey-West fixed bandwidth = 3.0000)

Variable	Coefficient	Std. Error	t-Statistic	Prob.
C	-0.08453	0.120027	-0.70429	0.492

R-squared	0	Mean dependent var	-0.08453
Adjusted R-squared	0	S.D. dependent var	0.445643
S.E. of regression	0.445643	Akaike info criterion	1.281866
Sum squared resid	2.978969	Schwarz criterion	1.330152
Log likelihood	-9.25493	Hannan-Quinn criter.	1.284338
Durbin-Watson stat	1.363205		

Notice, here, the estimation of the variance of the regression estimate using the HAC estimate.

We can conclude that both the estimates are equally precise.

We might also compare our forecast against the naïve forecast: in this case we obtain:

Dependent Variable: E3SQ-E2SQ
Method: Least Squares
Sample: 2011Q1 2014Q4
Included observations: 16
HAC standard errors & covariance (Bartlett kernel, Newey-West fixed bandwidth = 3.0000)

Variable	Coefficient	Std. Error	t-Statistic	Prob.
C	0.544485	0.202907	2.683422	0.017

R-squared	0	Mean dependent var	0.544485
Adjusted R-squared	0	S.D. dependent var	1.239716
S.E. of regression	1.239716	Akaike info criterion	3.328104
Sum squared resid	23.05345	Schwarz criterion	3.376391
Log likelihood	-25.6248	Hannan-Quinn criter.	3.330576
Durbin-Watson stat	2.289506		

In this case, we see that the forecast from the model is statistically superior to the naïve forecast.