

Seminar 4: ARMA modelling of the forward premium: Box Jenkins methodology and forecasting.

The objectives of this seminar are to:

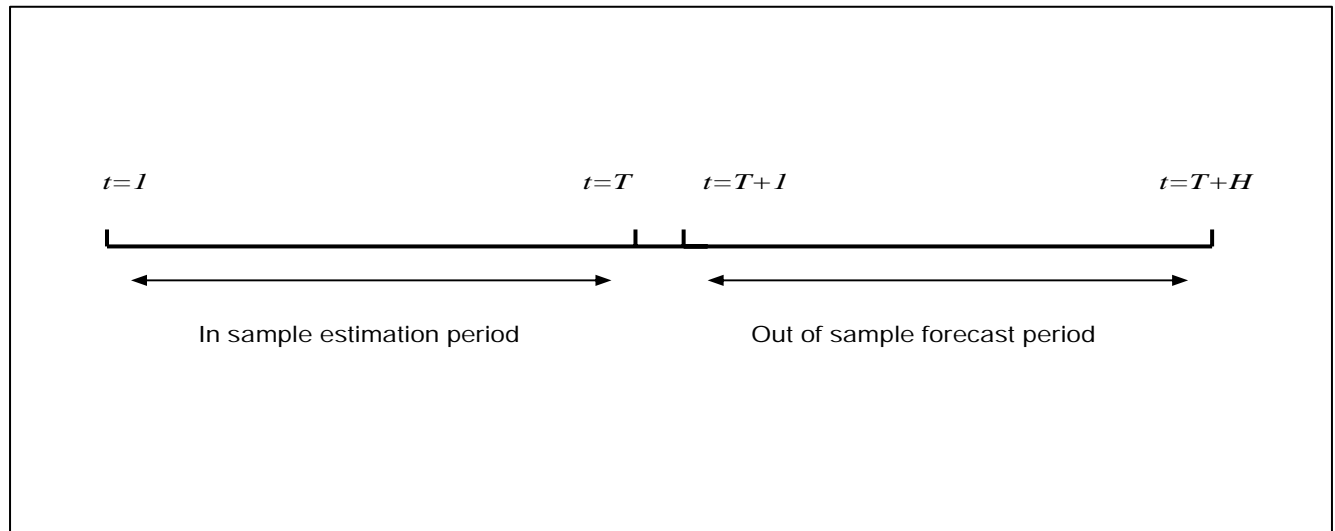
- Identify, estimate and test an ARMA model for the forward premium (Box-Jenkins methodology).
- Forecast the forward premium using this ARMA model and using the CIP (structural) model estimated in Seminar 3.
- Evaluate the forecasts from both models to see which provides better out-of-sample forecasts of the forward premium.

The learning outcomes from this seminar will be to develop your understanding of:

- ARMA modelling (identification, estimation and testing) and forecasting in *Eviews*.
- Forecast evaluation.
- The utility of ARMA models as forecasting tools.

The analysis is carried out in this handout for models of the forward premium (using the same workfile, **cip_sem3.wf1**, as used in Seminar 3). Try carrying out a similar analysis in your own time comparing the forecasts from a structural model of the 3 month holding period returns on sterling (UIP) with those from an ARMA model of this variable.

We are going to estimate the models for the period 5/09/2001 – 9/30/2004 ($t = 1, \dots, T$). Then we are going to forecast the forward premium out of sample for the period 10/1/2004 – 9/30/2005 ($t = T + 1, \dots, T + H$).




1. Forecasting with the CIP model

Firstly we need to re-estimate the GMM CIP model (see seminar 3) for the in sample estimation period 5/09/2001 – 9/30/2004.

Open the equation **cip_gmm** (which you saved in the last seminar) and click **Estimate** on the equation toolbar. Change the **Estimation settings/sample** to: 5/09/2001 – 9/30/2004. Then click OK.

Now forecast the forward premium for the period 10/1/2004 – 9/30/2005.

Click **Forecast** on the equation toolbar. Change the **Forecast sample** to: 10/1/2004 – 9/30/2005. Click OK:



Forecast: FP_3MF
Actual: FP_3M
Forecast sample: 10/01/2004 9/30/2005
Included observations: 261

Root Mean Squared Error	0.004045
Mean Absolute Error	0.003422
Mean Abs. Percent Error	26.21858
Theil Inequality Coefficient	0.138642
Bias Proportion	0.560762
Variance Proportion	0.049786
Covariance Proportion	0.389452

Freeze this forecast view for comparison later with the ARMA model forecasts: Click **freeze** on the equation toolbar. Then name the resulting graph: **cip_fcst**

These forecasts are static forecasts. These forecasts are made using the:

- Actual values of the explanatory variables (in this case, the interest differential) in the forecast period.
- Estimated coefficients from the in-sample estimation period ($\hat{\beta}$):

$$\hat{y}_{T+h} = x'_{T+h} \hat{\beta}, \quad h = 1, \dots, H$$

The forecast error is given by:

$$\begin{aligned}\hat{e}_{T+h}^S &\equiv y_{T+h} - \hat{y}_{T+h} \\ &= x'_{T+h}\beta + \varepsilon_{T+h} - x'_{T+h}\hat{\beta} \\ &= \varepsilon_{T+h} + x'_{T+h}(\beta - \hat{\beta})\end{aligned}$$

The forecast error depends on both residual uncertainty (the first term) and parameter uncertainty (the second term).

The forecast error variance is therefore:

$$\begin{aligned}\text{var}(\hat{e}_{T+h}^S) &= \text{var}(\varepsilon_{T+h}) + x'_{T+h} \text{var}(\hat{\beta})x_{T+h} \\ &= \sigma^2(1 + x'_{T+h}[X'X]^{-1}x_{T+h})\end{aligned}$$

From the above, a 95% forecast interval is given by:

$$\hat{y}_{T+h} \pm 1.96\sqrt{\text{var}(\hat{e}_{T+h}^S)}$$

The point forecasts (\hat{y}_{T+h}) and the 95% forecast interval are plotted in the *Eviews* output.

Forecast evaluation (see Brooks 5.12.8 and *Eviews* help index under 'Forecast, interval')

What is the total amount of error associated with the forecast? A forecast with a smaller total error is statistically more accurate. Clearly, a simple sum of the forecast errors wouldn't work: positive and negative errors would cancel out. Alternatives include summing the squared forecast errors or the absolute forecast errors:

Root - Mean - Squared - Error (RMSE)

$$RMSE = \sqrt{\frac{1}{H} \sum_{t=T+1}^{T+H} (\hat{y}_t - y_t)^2}$$

Compare these statistics over different models for y . The model with the smallest values for RMSE/MAE provides the most accurate forecasts.

Mean - Absolute - Error (MAE)

$$MAE = \frac{1}{H} \sum_{t=T+1}^{T+H} |\hat{y}_t - y_t|$$

RMSE is based on a quadratic loss function (i.e., it squares the forecast errors). It therefore penalizes large forecast errors (outliers) more heavily than MAE.

Mean - Absolute - Percentage - Error (MAPE)

$$MAPE = \frac{100}{H} \sum_{t=T+1}^{T+H} \left| \frac{\hat{y}_t - y_t}{y_t} \right|$$

The MAE is often reported in % (scale-free) terms.

Theil – Inequality – Coefficient (TIC)

$$TIC = \frac{RMSE}{\sqrt{\frac{1}{H} \sum_{t=T+1}^{T+H} \hat{y}_t^2 + \frac{1}{H} \sum_{t=T+1}^{T+H} y_t^2}}$$

TIC lies between zero and one.
A perfect forecast would imply
TIC=0 (RMSE=0).

It's also useful to see how well the forecast predicts:

- i) The mean of the actual series
- ii) The variance of the actual series

For example, the forecast may predict one of these well but not the other. In this context, *Eviews* reports a 'bias proportion' and 'variance proportion' which is informative about the accuracy of the forecasts of the mean and variance respectively of the series in the forecast period. Small bias and variance proportions indicate better forecasting of the mean and variance respectively. The covariance proportion captures the remaining unsystematic forecast error. This proportion will be large relative to the bias and variance proportions if the forecasts of the mean and variance are accurate (the bias, variance and covariance proportions sum to unity).

For the GMM CIP model:

- 56% of the total forecast error is due to inaccuracy in forecasting the mean of the forward premium (bias proportion).
- About 5% of the total forecast error is due to inaccuracy in forecasting the variance of the forward premium (variance proportion).
- The remaining 39% is due to unsystematic forecast error (covariance proportion).

This suggests that quite a large percentage of the total forecast error is due to inaccuracy in forecasting the mean of the forward premium. Can we build a model with more accurate forecasts of the forward premium? We will now go on to estimate and forecast an ARMA model for the forward premium.

2. Building an ARMA model for the forward premium

The three stages in the Box-Jenkins approach to ARMA modelling are:

1. Identification
2. Estimation
3. Testing

(see lecture 5 and Brooks 5.7)

Identification and estimation

Firstly we look at the ACF and PACF of the forward premium to help select an appropriate ARMA model.

Click on the *Sample* button (workfile toolbar) and re-set the sample to the in-sample estimation period: 5/09/2001 – 9/30/2004

Now open fp_3m and click *View/Correlogram/Level*:

Autocorrelation	Partial Correlation	AC	PAC	Q-Stat	Prob	
. *****	. *****	1	0.783	0.783	545.22	0.000
. *****	. ***	2	0.787	0.449	1096.6	0.000
. *****	. ***	3	0.806	0.372	1675.6	0.000
. *****	. *	4	0.784	0.195	2223.2	0.000
. *****	. **	5	0.800	0.223	2794.1	0.000
. *****	. *	6	0.794	0.159	3358.1	0.000
. *****	. *	7	0.789	0.124	3915.3	0.000
. *****	. *	8	0.787	0.091	4470.7	0.000
. *****	. *	9	0.793	0.112	5034.3	0.000
. *****	. .	10	0.772	0.013	5570.0	0.000

The first thing to note is the slow decay in the ACF. This suggests the 3-month forward premium may be non-stationary! So far, in accordance with finance theory, we've been assuming the forward premium is stationary (implying the forward and spot rates move together in the long-run).

For the moment, suppose we try and interpret the above ACF and PACF. There are 9 spikes in the PACF (and a slow decay in the ACF) so we might try fitting an AR(9) model. (This is subjective; other interpretations are possible – try fitting ARMA models for comparison).

On the main toolbar click **Quick/Estimate Equation** (Method=Least Squares) and enter:

fp_3m c ar(1) ar(2) ar(3) ar(4) ar(5) ar(6) ar(7) ar(8) ar(9)

Dependent Variable: FP_3M
 Method: Least Squares
 Date: 02/11/07 Time: 12:14
 Sample (adjusted): 5/23/2001 9/30/2004
 Included observations: 877 after adjustments
 Convergence achieved after 4 iterations

Variable	Coefficient	Std. Error	t-Statistic	Prob.
C	0.020211	0.004986	4.053455	0.0001
AR(1)	0.023083	0.03363	0.686383	0.4927
AR(2)	0.054673	0.033428	1.635527	0.1023
AR(3)	0.15936	0.033157	4.80619	0
AR(4)	0.071129	0.033303	2.135805	0.033
AR(5)	0.164254	0.03292	4.989539	0
AR(6)	0.130829	0.033304	3.928378	0.0001
AR(7)	0.124539	0.033154	3.75642	0.0002
AR(8)	0.108111	0.033383	3.238513	0.0012
AR(9)	0.138852	0.033595	4.133057	0
R-squared	0.782298	Mean dependent va	0.016244	
Adjusted R-squared	0.780038	S.D. dependent var	0.006525	
S.E. of regression	0.00306	Akaike info criterion	-8.729292	
Sum squared resid	0.00812	Schwarz criterion	-8.674828	
Log likelihood	3837.795	F-statistic	346.1678	
Durbin-Watson stat	2.011531	Prob(F-statistic)	0	
Inverted AR Roots				
	.09-.77i	.56-.58i	.56+.58i	.09+.77i
	-.71-.28i	-.43-.66i	-.43+.66i	-.71+.28i

Our suspicions are confirmed – there is one unit root in the AR polynomial indicating the forward premium is non-stationary.

The non-stationarity of the forward premium suggests the forward and spot rates are unrelated (they drift apart) in the long-run! This finding is counter-intuitive from a finance perspective: if the spot and forward rates are driven by the same fundamentals then we would expect them to be related in the long-run (implying the forward premium is stationary). However, it is not uncommon in empirical studies to find that the forward premium is non-stationary. This paradox may be due to ‘near unit root’ or long memory behaviour in the forward premium. This is behaviour which is ‘close’ to being non-stationary so that standard tests/models incorrectly indicate that the series is non-stationary (see Cuthbertson and Nitzsche, 2004, 25.3: Time-Series Properties for a discussion of the time-series properties of the forward premium). We will return to the issue of modelling long-memory in financial data in a later class.

For now we will proceed on the basis that the forward premium has one unit root. On this basis we need to first difference the forward premium to make it stationary ('remove' the unit root). Then we can identify, estimate and test an ARMA model for the differenced (stationary) series.

Firstly, inspect the correlogram of the first difference of the forward premium:

Open fp_3m and click *View/Correlogram/1st difference*:

Autocorrelation	Partial Correlation	AC	PAC	Q-Stat	Prob	
**** .	**** .	1	-0.514	-0.514	234.18	0.000
. .	*** .	2	-0.035	-0.406	235.28	0.000
. *	** .	3	0.097	-0.222	243.74	0.000
* .	** .	4	-0.090	-0.248	250.99	0.000
. .	* .	5	0.051	-0.183	253.30	0.000
. .	* .	6	0.002	-0.145	253.30	0.000
. .	* .	7	-0.014	-0.117	253.47	0.000
. .	* .	8	-0.013	-0.140	253.63	0.000
. .	. .	9	0.056	-0.046	256.46	0.000
. .	. .	10	-0.050	-0.056	258.74	0.000

Firstly note that differencing the series has indeed resulted in a stationary series (explain why). There are 3-4 spikes in the ACF and the PACF exhibits a decaying pattern. This suggests that an MA(4) model may be appropriate for the first difference of the forward premium (see lecture 5).

On the main toolbar click *Quick/Estimate Equation* (Method=Least Squares) and enter:

d(fp_3m) c ma(1) ma(2) ma(3) ma(4)

Dependent Variable: D(FP_3M)
 Method: Least Squares
 Date: 02/11/07 Time: 13:21
 Sample (adjusted): 5/11/2001 9/30/2004
 Included observations: 885 after adjustments
 Convergence achieved after 14 iterations
 Backcast: 5/07/2001 5/10/2001

Variable	Coefficient	Std. Error	t-Statistic	Prob.
C	1.93E-05	9.20E-06	2.100871	0.0359
MA(1)	-1.006855	0.0337	-29.87718	0
MA(2)	0.018719	0.047703	0.392414	0.6948
MA(3)	0.105833	0.047703	2.218604	0.0268
MA(4)	-0.027969	0.033755	-0.828612	0.4075
R-squared	0.50812	Mean dependent va	2.09E-05	
Adjusted R-squared	0.505885	S.D. dependent var	0.004297	
S.E. of regression	0.00302	Akaike info criterion	-8.761287	
Sum squared resid	0.008028	Schwarz criterion	-8.73425	
Log likelihood	3881.87	F-statistic	227.2639	
Durbin-Watson stat	2.003221	Prob(F-statistic)	0	
Inverted MA Roots	0.89	.24-.18i	.24+.18i	-0.36

The coefficients on the second and fourth MA terms are statistically insignificant. We will therefore drop these terms to increase the efficiency (reduce the variance) of our estimates (including irrelevant variables reduces efficiency). Click *Estimate* on the equation toolbar and enter:

d(fp_3m) c ma(1) ma(3)

Dependent Variable: D(FP_3M)
 Method: Least Squares
 Date: 02/11/07 Time: 13:33
 Sample (adjusted): 5/11/2001 9/30/2004
 Included observations: 885 after adjustments
 Convergence achieved after 7 iterations
 Backcast: 5/08/2001 5/10/2001

Variable	Coefficient	Std. Error	t-Statistic	Prob.
C	1.93E-05	9.46E-06	2.034317	0.0422
MA(1)	-0.996568	0.023738	-41.98282	0
MA(3)	0.088954	0.023803	3.737046	0.0002
R-squared	0.507603	Mean dependent va	2.09E-05	
Adjusted R-squared	0.506486	S.D. dependent var	0.004297	
S.E. of regression	0.003018	Akaike info criterion	-8.764755	
Sum squared resid	0.008036	Schwarz criterion	-8.748533	
Log likelihood	3881.404	F-statistic	454.6183	
Durbin-Watson stat	2.026844	Prob(F-statistic)	0	
Inverted MA Roots	0.88	0.38	-0.27	

Try estimating alternative AR/MA/ARMA models for the differenced forward premium. Compare the SC to see which is the 'best' model (choose the model with the smallest SC). Note that AIC tends to over-fit the number of terms (it chooses models with too many AR/MA terms)

Testing

Now we need to test whether this MA model has ‘mopped up’ all the dynamics in the data. To do this look at the correlogram for the model residuals:

Click *View/Residual Tests/Correlogram Q – statistics*

Autocorrelation	Partial Correlation	AC	PAC	Q-Stat	Prob	
. .	. .	1	-0.014	-0.014	0.1639	
. .	. .	2	0.014	0.014	0.3337	
. .	. .	3	0.031	0.031	1.1707	0.279
* .	* .	4	-0.062	-0.062	4.6284	0.099
. .	. .	5	0.025	0.023	5.1880	0.159
. .	. .	6	0.008	0.009	5.2412	0.263
. .	. .	7	-0.003	0.001	5.2473	0.386
. .	. .	8	-0.000	-0.006	5.2473	0.513
. .	. .	9	0.034	0.037	6.3011	0.505
. .	. .	10	-0.042	-0.041	7.8602	0.447

The ACF and PACF are compatible with the residuals following a white noise process (be sure you understand why). This indicates the MA model is an adequate representation of the first difference of the forward premium.

3. Forecasting the forward premium with the ARMA model

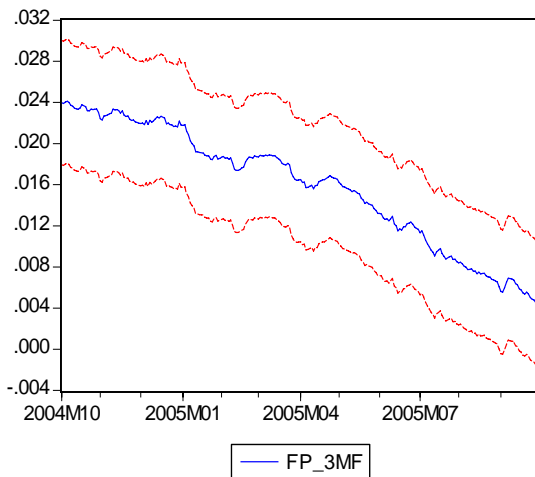
The final piece of analysis is to forecast the forward premium using the MA model for the differenced forward premium and to compare these forecasts with those from the CIP model.

Click **Forecast** on the equation toolbar.
 Select fp_3m as the **Series to forecast**.
 Change the **Forecast sample** to: 10/1/2004 – 9/30/2005.
 Change the **Method** to Static forecast.

Click OK:

The 'static' option will calculate one-step ahead forecasts for the ARMA model. There is an alternative option for this model to make dynamic forecasts. Dynamic forecasts use only information available up to time T (the end of the estimation period) to make the forecast. This requires the values of the explanatory variables themselves to be forecasted in the out-of-sample period. An ARMA model generates forecasts of its own explanatory variables since these are simply lagged values of the dependent variable (and error terms). Accordingly the model can be forecasted recursively (e.g., the forecast for y_{T+1} can be used subsequently to forecast y_{T+2} and so on up to y_{T+H}). This is sometimes called multi-step or h-step-ahead forecasting.

In contrast the static forecasts reported for the CIP model used the actual values of the explanatory variable (interest differential) to make forecasts. There was no option to make dynamic forecast in that case because the CIP model does not generate forecasts of the interest differential.



Forecast:	FP_3MF
Actual:	FP_3M
Forecast sample:	10/01/2004 9/30/2005
Included observations:	261
Root Mean Squared Error	0.002733
Mean Absolute Error	0.002096
Mean Abs. Percent Error	26.03801
Theil Inequality Coefficient	0.082788
Bias Proportion	0.137034
Variance Proportion	0.051547
Covariance Proportion	0.811418

Q: Compare the forecasts of the forward premium from the MA model with those from the CIP (structural) model. Which model provides the more accurate forecasts?

On all criteria the time series model out-performs the CIP model in the accuracy of its forecasts: the RMSE, MAE, MAPE and TIC are all smaller for the ARMA model than the CIP model. Notice also that the bias proportion, at around 14%, is much smaller for the ARMA model than the CIP model (which was around 56%). This suggests that the ARMA model forecasts the mean of the forward premium in the out-of-sample period much more accurately than the CIP model.

Conclusions

ARMA models are a-theoretical, statistical models and are of little use for policy formation. However they are very useful for forecasting financial time series and often do this better than structural/theory based models. A Box-Jenkins analysis led to the estimation of an MA model for the first difference of the forward premium. The forecasts of the forward premium from this MA model provided more accurate forecasts than a CIP (structural) model for the forward premium.

The ARMA modelling highlighted a paradox in that the forward premium is apparently non-stationary (it has a unit root). This finding is in conflict with the theoretical model (which predicts the forward premium is stationary). This finding is not a problem for the statistical (ARMA) model since first-differencing removes the unit root; a stationary ARMA model can simply be built for the first difference of the forward premium. However, whilst this solution is statistically valid, it is not a wholly satisfactory outcome as it leaves unresolved the theoretical paradox of a non-stationary forward premium. Perhaps this failing is due to inadequacies in our statistical (ARMA) model to describe a wider range of 'near' unit root or long memory behaviour in the forward premium. We will explore this issue in a future class.

References

- Cuthbertson and Nitzsche (2004). *Quantitative Financial Economics*. Second Edition. Wiley: Chichester.
- Brooks (2002). *Introductory Econometrics for Finance*. CUP: Cambridge