



BANK OF ENGLAND

CENTRE FOR CENTRAL BANKING STUDIES

ECONOMIC MODELLING AND FORECASTING

A short-term inflation forecast for the UAE: Part I

by

Andy Blake and Ole Rummel
Centre for Central Banking Studies
Bank of England

6 February 2015

1 Short-term inflation forecasting

This exercise provides an example of a short-term inflation forecast (STIF) in EViews. The aim of the exercise is to construct a STIF using disaggregated consumer price index (CPI) data for the United Arab Emirates (UAE). In particular, it generates a STIF for the composite CPI (and hence the monthly and/or annual headline inflation rate) up to six months ahead, which can then be compared to the actual inflation numbers.¹ The methodology of the STIF is extremely straightforward: we construct a time-series model for each CPI component at a chosen level of disaggregation. Each of these models is then used to produce a short-term forecast for the respective CPI component. The individual forecasts are then (re-)aggregated into the CPI index, using the (published) weights of each component in the composite CPI index. The forecast for the CPI index in levels can then be used to calculate a (short-term) inflation forecast.

Although the final STIF for the UAE works reasonably well, this exercise is purely **illustrative** and no claim is made that this is the only – or, indeed, the best – way of generating a short-term inflation forecast.

As we will see, generating a STIF is both a science and an art, and relies to a great extent on the knowledge of local price pressures and dynamics, the (short-run) behaviour of economic agents in the economy, the institutional set-up regarding national statistics offices and, hence, an understanding of the construction of the CPI index. We will return to these points as we go along.

2 Data and EViews programme

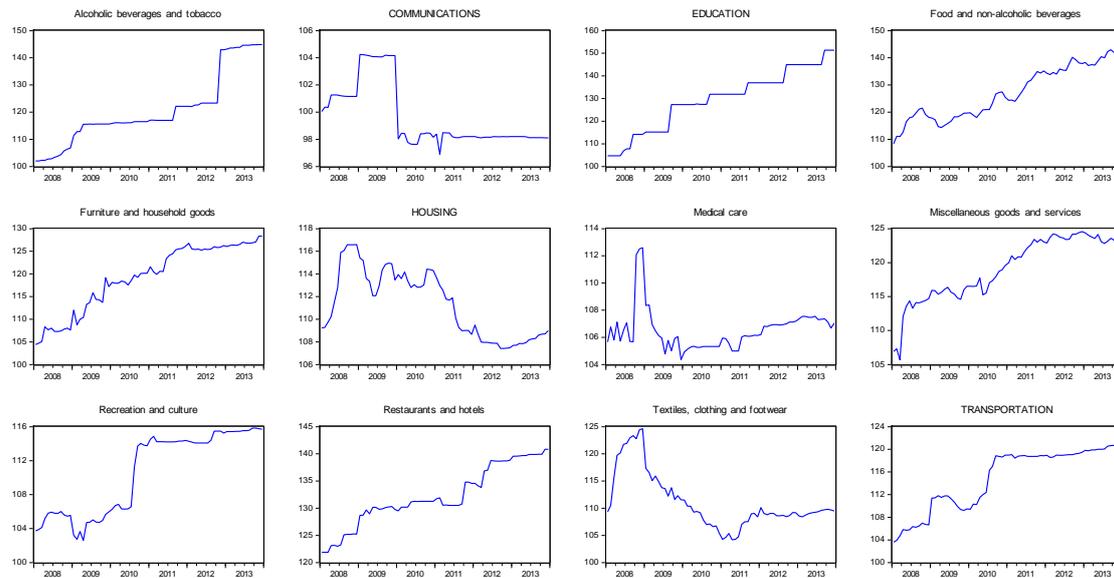
All the underlying files can be found in a folder called **STIF – UAE** in the current course **Data** folder.

The STIF used in this exercise is based on a monthly dataset of twelve disaggregated UAE CPI components, spanning the period from January 2008 (2008M01 in EViews notation) to June 2014 (2014M06) for a total of 78 monthly observations. For the purposes of the exercise, we will assume that we have data up to December 2013, i.e., 2013M12. The forecasts will then be constructed for the six-month period from January 2014 to June 2014 and compared to the actual inflation data over the same period.

Open the workfile **stif_uae.wf1** and look at the group **group_all**. These are the twelve main components of the UAE's composite CPI index. As is quite customary in CPI indices in many countries, the three components with the largest weights are housing (39.33 per cent); food and non-alcoholic beverages (13.94 per cent) and transportation (9.94 per cent). Altogether, these three components account for slightly more than 63 per cent of the UAE composite CPI index. The three components with the lowest weights are recreation and culture (3.07 per cent), medical care (1.12 per cent) and alcoholic beverages and tobacco (0.22 per cent). Time series of the price levels for the twelve components are given in Figure 1.

¹ In principle, a STIF can be performed for any forecast horizon, although a **short-term** inflation forecast based on simple times-series models that are meant to capture the significant persistence in inflation data is traditionally not really performed beyond six months into the future.

**Figure 1: Disaggregated UAE composite CPI components
(January 2008 – December 2013)**



Many of the series show some form of (upward) trend over the time period (or at least large parts of it), with some displaying additional idiosyncratic features (such as `medical_care`, `housing` and `communications`). Amongst other things, this raises the question of what time period to use for the estimation of the underlying models. For example, would you want to use the entire sample period for `transportation`? In a similar vein, `education` looks quite deterministic, with regular increases in August, and we will return to this point below. `Food_soft_drinks`, on the other hand, is clearly upward sloping, but there are substantial short-term movements around the trend. Without further analysis, we cannot say whether these movements are seasonal or not. Based on visual inspection alone, there are – apart from February 2011 for `communications` and the period of 2008 Q4 for `medical_care` – no other obvious outliers (aberrant observations) that need to be accounted for. Even so, we may want to think about the step-ups (level interventions) in `beverages_tobacco` (September 2011 and November 2012) and `recreation_culture` (September 2009 as well as September 2012) a bit more.

Open the programme `uae_component.prg`, which contains prototype models for each of the twelve components. These are **prototype** models in that they are initial attempts (arrived at by – more or less – trial and error) to model the twelve underlying series. We make no claim that these models will be good models for either data description or forecasting. Indeed, we may want to try and modify these models as we go along. The programme looks as follows:²

```
smp1 2008m05 2013m12
equation eqn_food.ls food_soft_drinks c @trend
food_soft_drinks(-1 to -4) d2010m12
```

² The different EViews programmes that we will encounter in this exercise are discussed in more detail in the Appendix.

```

smpl 2009m01 2013m12
equation eqn_educ.ls d(education) c dm09

smpl 2009m05 2013m12
equation eqn_bevs.ls beverages_tobacco c @trend
beverages_tobacco(-1 to -3) d2011m09 d2012m11

smpl 2010m09 2013m12
equation eqn_trans.ls transportation c @trend
transportation(-1 to -3)

smpl 2011m03 2013m12
equation eqn_rec.ls recreation_culture c
recreation_culture(-1) d2012m09

smpl 2011m07 2013m12
equation eqn_comms.ls communications c @trend
communications(-1 to -2)

smpl 2011m09 2013m12
equation eqn_misc.ls miscellaneous_goods c
miscellaneous_goods(-1) d2013m06

smpl 2011m10 2013m12
equation eqn_text.ls textiles_clothing c @trend
textiles_clothing(-1)

smpl 2012m02 2013m12
equation eqn_furn.ls d(furniture_household) c @trend
d(furniture_household(-1)) d(furniture_household(-2))

smpl 2012m02 2013m12
equation eqn_med.ls medical_care c @trend
medical_care(-1) d2013m11

smpl 2012m06 2013m12
equation eqn_rest.ls restaurants_hotels c @trend
restaurants_hotels(-1)

smpl 2012m11 2013m12
equation eqn_hous.ls housing c @trend housing(-1)

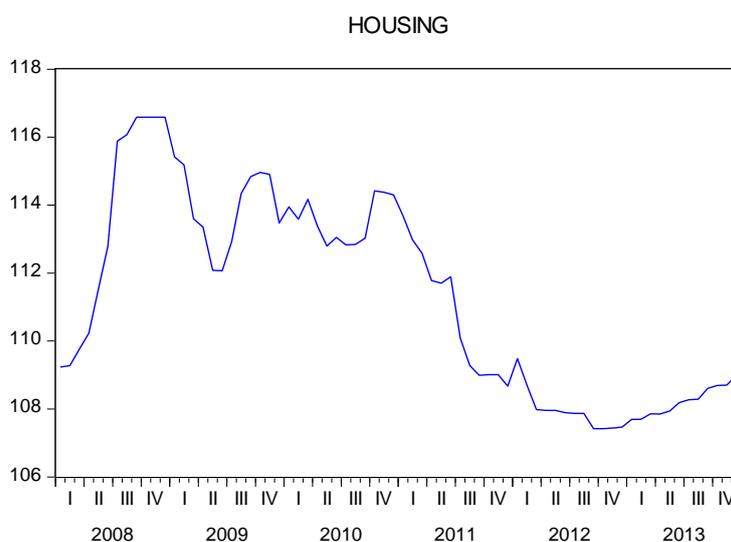
```

For this exercise, we have chosen to model each individual component index series as a simple autoregressive function of a constant, one or more lagged values of the underlying component index, a linear time trend and – where appropriate – seasonal dummies as well as impact dummy variables for specific months during the sample period. You will also note that all but three series appear in levels, while education, furniture_household and medical_care enter in first differences, i.e., changes.

The average R^2 of the twelve individual models is equal to 0.768, with the R^2 's of the individual models ranging from 0.226 for `textiles_clothing` to an extremely high 0.999 for `beverages_tobacco`.

Where do the prototype models come from? The example of `housing` may serve to illustrate the steps that we followed in terms of building these models. Figure 2 replicates the time series of the underlying CPI component over the period from January 2008 to December 2013. We note the sharp rise over 2008, followed by a levelling off, volatility to the end of 2010 and a marked fall over the course of 2011 and the first half of 2012. An upward trend in the series only starts in earnest towards the end of 2012, which is why we decided to shorten the sample period for estimation significantly from November 2012 to December 2013.

Figure 2: Time series of housing (January 2008 – December 2013)



Having decided to use an autoregressive model with p lags ($AR(p)$), where p will obviously vary from component to component index, the first question concerns the number of autoregressive lags to include in the equation to be estimated. We follow a **general-to-specific** approach, estimating an AR model with a ‘large’ number of lags and then testing the statistical significance of each individual lag, eliminating insignificant lags sequentially in an attempt to reduce the lag order of the AR model. We start off with an initial number of lags, p_{max} . A formula for p_{max} that has been shown to perform well in applied work is $p_{max} = \text{int}(T^{1/3})$, where T is the number of observations and $\text{int}(\cdot)$ refers to the integer part of $T^{1/3}$. In the case of `housing`, we have 14 observations over the shortened sample period, such that $p_{max} = \text{int}(14^{1/3}) = \text{int}(2.41) = 2$. We therefore start off by estimating an $AR(2)$ model on the level of `housing`. As this is a **general-to-specific** modelling approach, we start off with the most general model possible, so we also include a time trend (`@trend`) in the estimation. Table 1 contains the estimation results of our initial model.

Table 1: AR(2) model for housing (November 2012 – December 2013)

Dependent Variable: HOUSING
 Method: Least Squares
 Sample: 2012M11 2013M12
 Included observations: 14

Variable	Coefficient	Std. Error	t-Statistic	Prob.
C	86.06218	36.14619	2.380948	0.0385
@TREND	0.101095	0.037959	2.663242	0.0238
HOUSING(-1)	-0.018442	0.324698	-0.056797	0.9558
HOUSING(-2)	0.162431	0.270094	0.601387	0.5610
R-squared	0.975782	Mean dependent var		108.1200
Adjusted R-squared	0.968517	S.D. dependent var		0.489615
S.E. of regression	0.086875	Akaike info criterion		-1.813735
Sum squared resid	0.075473	Schwarz criterion		-1.631148
Log likelihood	16.69615	Hannan-Quinn criter.		-1.830637
F-statistic	134.3058	Durbin-Watson stat		2.253135
Prob(F-statistic)	0.000000			

We find that the second lag of housing (housing(-2)) is not significant, so we eliminate it. Results for this new model are given in Table 2.

Table 2: AR(1) model for housing (November 2012 – December 2013)

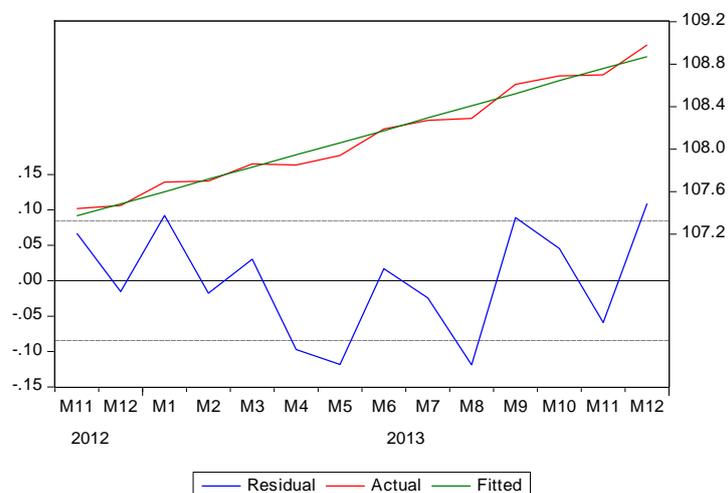
Dependent Variable: HOUSING
 Method: Least Squares
 Sample: 2012M11 2013M12
 Included observations: 14

Variable	Coefficient	Std. Error	t-Statistic	Prob.
C	96.81028	30.49341	3.174794	0.0088
@TREND	0.111471	0.032815	3.396921	0.0060
HOUSING(-1)	0.038144	0.301614	0.126465	0.9016
R-squared	0.974906	Mean dependent var		108.1200
Adjusted R-squared	0.970344	S.D. dependent var		0.489615
S.E. of regression	0.084317	Akaike info criterion		-1.921065
Sum squared resid	0.078202	Schwarz criterion		-1.784124
Log likelihood	16.44745	Hannan-Quinn criter.		-1.933741
F-statistic	213.6776	Durbin-Watson stat		2.078477
Prob(F-statistic)	0.000000			

The first lag is also statistically insignificant, so we could reduce the lag order, $p = 1$, further, resulting in a model with a constant, a time trend and no lags. As a consequence of eliminating the first lag of housing from the equation, though, the model would become purely deterministic, consisting only of a constant and a time trend. In order to allow for some dynamics, we may therefore want to retain housing(-1) in the equation, even though it is insignificant. In practical terms, this would permit deviations from the deterministic time trend. At this stage, we can take a look at the estimated residuals to get an impression of their dynamics. Clicking on **View/Actual,Fitted,Residual/Actual,Fitted,Residual Graph** generates a graph of

the actual versus the fitted values as well as the residuals, which are obviously the difference between the two. This graph is reproduced in Figure 3.

Figure 3: Actual, fitted and residual values for the AR(1) model for housing (November 2012 – December 2013)



As can be seen from Figure 3, there are some outliers of either positive or negative size that we can identify: May 2013, August 2013 and December 2013.³ In many instances in applied work, outliers come in pairs, i.e., a very large positive (negative) value in one period is followed by a negative (positive) value of equal size one (or two) period later. Such a clustering of outliers in adjacent time periods would be indicative of an impact dummy. In the case of Figure 3, the one period where we **may** want to consider including a dummy is the last one, in the sense that to forecast ‘well’, the starting point for the forecast should be relatively close to the actual values. But the augmented AR(1) model for housing with the addition of an impact dummy for December 2013 (d2013m12) did not have a material impact on the forecast process so, to preserve degrees of freedom, we decided to leave it out of the model. The final prototype forecasting model for housing, which has an R^2 of 0.975, is therefore the one shown in Table 2 above, consisting of a constant, a linear time trend and the first lag of housing.

In addition, it appears that many seasonally observed economic time series display seasonality in the sense that the observations in certain seasons have properties that differ from those data points in other seasons. If the time period under observation were long enough, evidence of residual spikes that recur in the same period every year would be indicative of a **seasonal** dummy. There is no such evidence for any of the twelve models that we consider in this exercise.

Each of the twelve models can obviously be improved upon from an econometric standpoint. To begin with, there is no reason why we should use simple autoregressive models. A large number of STIFs in other central banks, including the Bank of England, are formulated more generally as either structural time-series

³ At the same time, we should keep in mind that when working with 95 per cent confidence bands, as displayed in Figure 3, we would still expect to see residual values falling outside the upper and lower bands in 5 per cent of cases. In essence, we are looking for very large and marked outliers in a case of ‘you know them when you see them’.

models (STSMs) or unobserved component (UC) models and estimated using the state-space representation and the Kalman filter.⁴

Now open and run the programme **uae_component.prg** by clicking on **Run** on the programme window. After completion, you will find that EViews has created – and saved – twelve equation objects in the workfile, one for each CPI component.

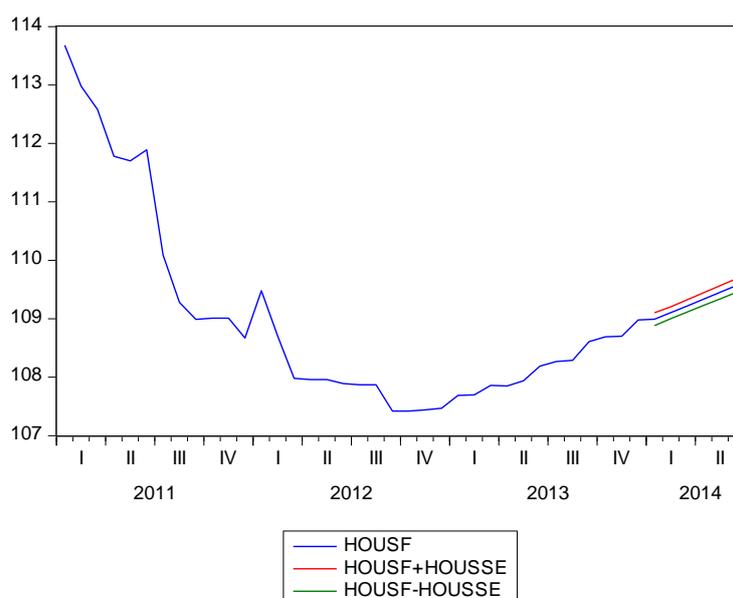
3 Forecasting the individual components

In the next step, we will use each of these twelve individual equation objects to produce a forecast over the period from 2014M01 to 2014M06. Open the programme **uae_component_forecast.prg**. When run (again, click on **Run** in the programme window), this automatically forecasts all the component models over the period from 2014M01 to 2014M06 and graphs the output using twelve versions of the following sample code for the `housing` series (with the individual programmes amended as appropriate):

```
eqn_hous.forecast housf housse
group ho housf (housf+housse) (housf-housse)
ho.line
```

An example of one of these graphs (for `housing`), showing the mean forecast (denoted `housf`) as well as the plus and minus one standard error forecast interval (given by `housf+housse` and `housf-housse` respectively), is shown in Figure 4 below. We can see that the level of the housing CPI is forecast to increase steadily over the next six months. At first glance, there is little to suggest that this forecast is a bad one.

Figure 4: Forecast of the housing equation with one standard-error bounds (January 2014 – June 2014)



⁴ In fact, Part II of this exercise illustrates an approach using structural time-series models estimated using the state-space representation and the Kalman filter.

We should note in passing that the underlying forecasting model for education (as well as `furniture_household` and `medical_care`) is estimated on the first difference of the respective variable, e.g., `d(education)`, yet the unreported graph for education shows the forecast for the **level** – rather than the first difference – of education. The reason for this is that we have re-written the dependent variable in the equation we have estimated as `d(education)` (you can check this for yourself in the EView’s equation object `eqn_edu` in the workfile).⁵ EViews always provides you with the option of forecasting an expression involving the dependent variable. If this expression can be normalised (i.e., solved for the (first) series in the expression), EViews also provides you with the option to forecast the normalised series. In short, this means that in the case at hand, EViews will give us the option of either forecasting the series in first differences, i.e., `d(education)`, or in levels, i.e., `education`.⁶

Do you want to revise any of the models in light of their forecasts? Do we even have to try and improve the individual models? An argument can be made that, *a priori*, we should probably spend more time and effort in improving the models that have greater weight in the all-items CPI index. On the other hand, the large body of evidence on forecast combination and/or aggregation agrees on the benefits of combining forecasts from different models (see Timmermann (2006) or Section 8.5 in Elliott and Timmermann (2008)). The reason for this is that particular biases or shortcomings in one model are offset by biases going the other way in another model.

At the same time, some of the individual forecasts throw up some interesting questions. To begin with, the mark of a good forecasting model is that the forecasted values converge towards their long-run (unconditional) mean in the data, which is called the **attractor**. This will be the case for dynamic models, especially those containing autoregressive terms. In practical terms, you may find with some data-sets that the forecast path goes the other way from the historical trend before the forecast period, i.e., the forecast changes direction: the forecast of a hitherto upward trending time-series may suddenly fall and the forecast of a declining time-series may suddenly increase. Both of these facts are due to the forecast mechanically taking the series back to the attractor, i.e., the unconditional mean of the series in the data. But there is little evidence of this happening within the particular dataset we are using.

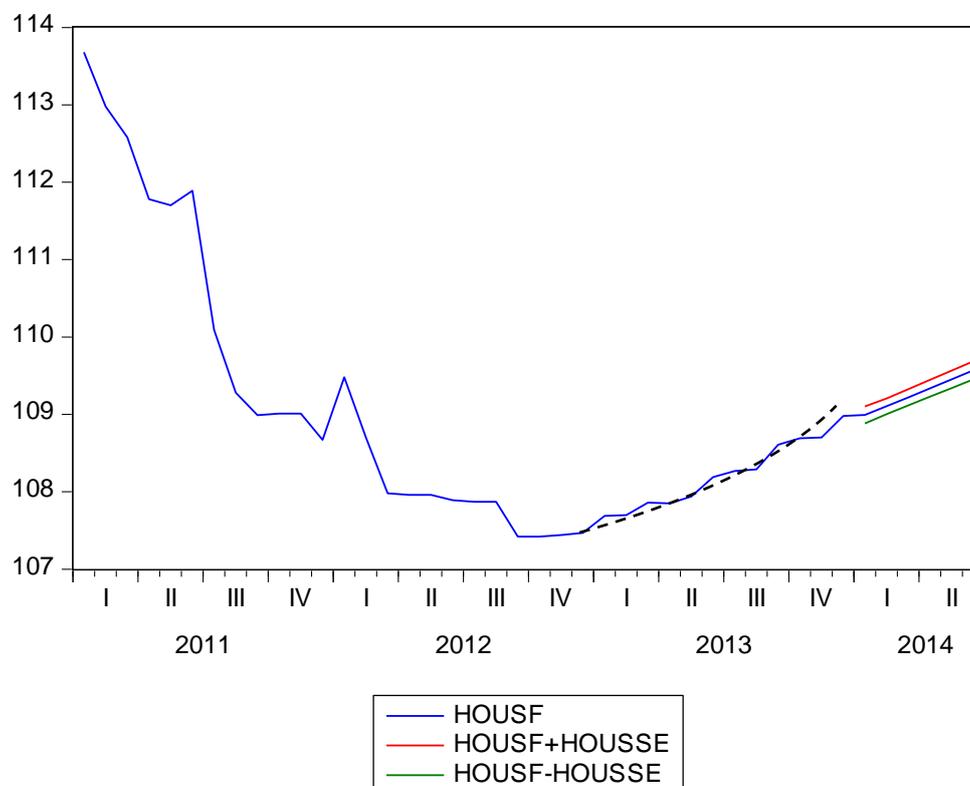
But another issue closer to hand is that some forecasts do not quite ‘smell’ right, in the sense that the forecast path could still be inconsistent with their previous trend. Without wanting to belabour the point, this underscores the important role of intuition and judgement in the construction of a STIF, where subjectivity is as – if not more – important than objectivity. An example of this may be provided by the forecast for `housing`, which is shown in Figure 5 below. Imposing a very rough (mental) curved trendline since the end of 2012 over the blue line shows that the forecasted path of `housing`, which is essentially a continuation of the estimated linear trendline, may be too low compared to the previous period. This particular path of the forecast

⁵ The alternative would be to generate the first difference of `education` manually as a new variable, using `deducation = d(education)`, and then estimating the regression equation with `deducation`. The drawback of going about estimation in this way is that the forecast will be for the growth rate of the underlying education price index, and **not** for the level of the underlying index.

⁶ EViews’ default option is to do the latter, i.e., forecast the underlying series (`education`) rather than the particular expression involving the dependent variable (`d(education)`).

may therefore be at odds with the recent past, so we may want to think about using a different model. In fact, Part II of this exercise will do exactly that.

Figure 5: Forecast of the housing equation with one standard-error bounds (January 2014 – June 2014)



4 Forecasting the headline CPI index

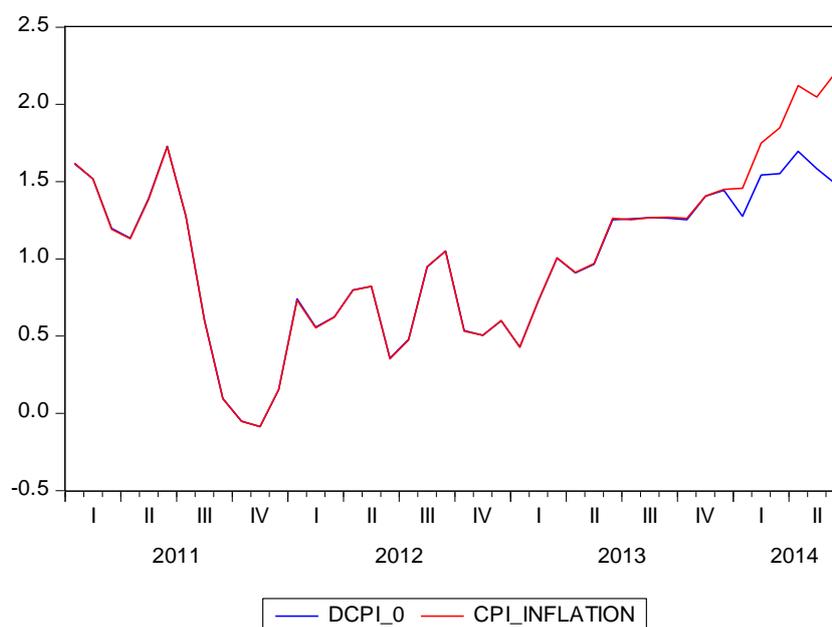
The headline CPI forecast is now straightforward: we simply weight together the twelve components using the published CPI weights as a model of the composite CPI inflation rate. The programme `uae_cpi_forecast.prg` defines the weights, specifies the model and solves for all the forecasts and then calculates and plots the inflation rate. The underlying weights are defined in a vector called `weights`:⁷

```
vector(12) weights
weights.fill 13.936, 0.218, 7.580, 39.334, 4.208, 1.124,
9.941, 6.932, 3.067, 4.004, 4.348, 5.308
```

Have a look at the programme and the way it creates the model and then run it by clicking on **Run** in the programme window. The output in Figure 6 consists of two series: `cpi_inflation`, which is the rate of inflation based on the actual underlying data, and `dcpi_0`, which is the disaggregated inflation forecast. Is the forecast any good?

⁷ The weights are taken from Section 4.4 in the UAE's National Bureau of Statistics' *Consumer Price Indices Annual Bulletin 2013*.

**Figure 6: Forecast of the composite CPI inflation rate
(January 2014 – May 2014)**

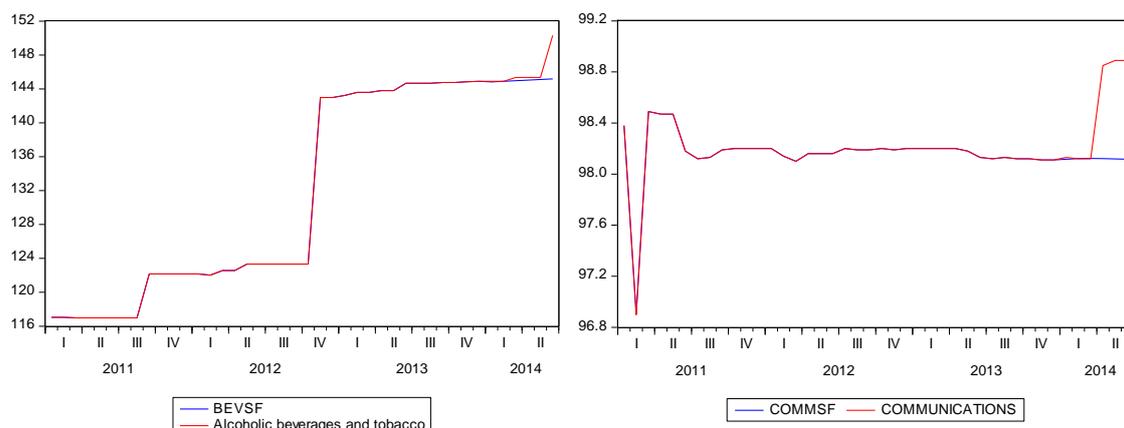


The actual composite CPI annual inflation rates for January through June 2014 are 1.46 per cent, 1.75 per cent, 1.85 per cent, 2.12 per cent, 2.05 per cent and 2.20 per cent respectively. The forecasted inflation rates over the same period (January through June 2014) are 1.27 per cent, 1.54 per cent, 1.55 per cent, 1.70 per cent, 1.58 per cent and 1.49 per cent. This is clearly **not** a good (short-term) forecast.

Some of the reasons for the poor forecast performance have to do with ‘known unknowns’, such as changes to administered prices, which can be taken into account in the forecasting process, while others have to do with ‘unknown unknowns’, such as shocks hitting particular CPI components over the forecast period, which are impossible to address. Finally, it could simply be that our forecast model is wrong, in the sense that it is an inappropriate representation of the true data-generating process. There seems to be evidence for all three of these happening in the dataset.

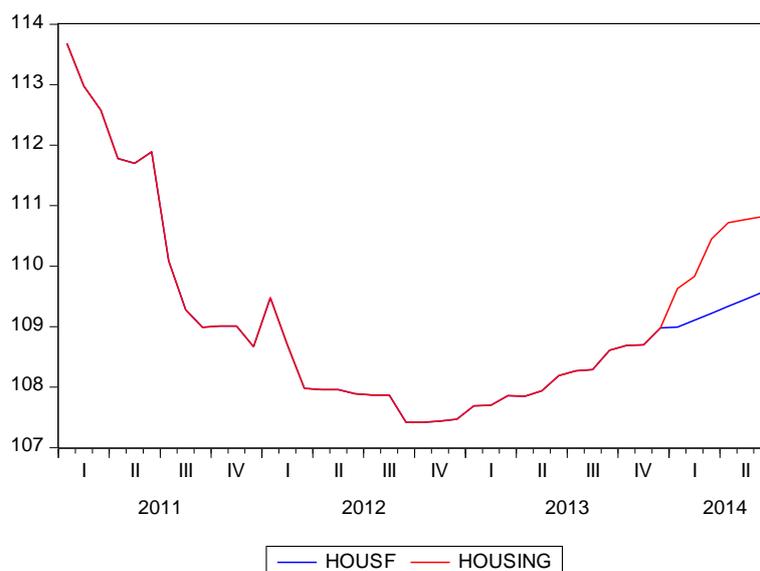
To begin with, the series for `beverages_tobacco` as well as `communications` show step-ups in May and March 2014 respectively (Figure 7), which resemble either a rise in taxes or an increase in administered prices. Both of these will be known to the authorities in advance and can be communicated to the central banks so that they can be incorporated into the STIF.

Figure 7: Forecast of beverages_tobacco and communication with actual outturns (January 2014 – June 2014)



The forecast for housing seems to be a case of where the model may be inappropriate. Figure 8 shows the forecasted values as well as the actual outturns over the forecast period.

Figure 8: Forecast of housing with actual outturns (January 2014 – June 2014)

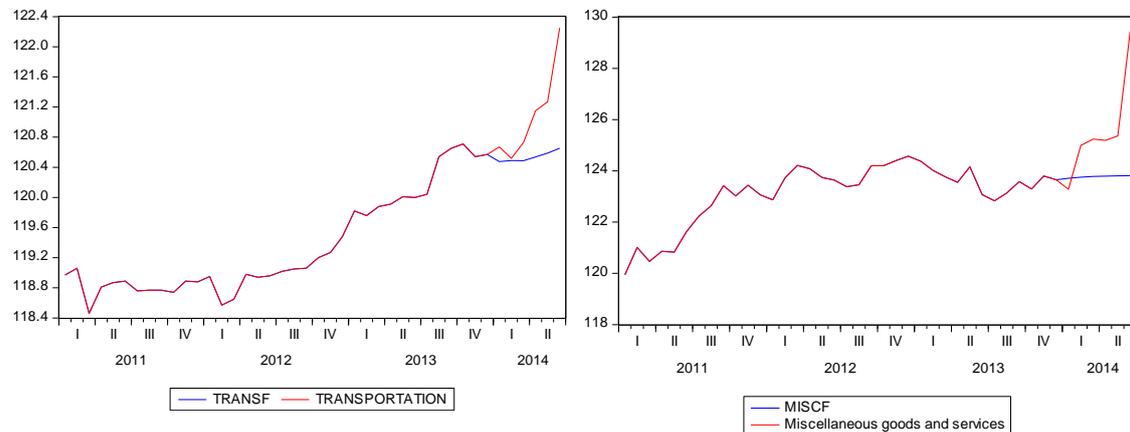


As discussed in Section 3 above, the estimated model, which consists of a constant, a linear time trend and an insignificant first lag of housing, is incapable of capturing the slight acceleration in the non-linear trend over the period since the end of 2012. Remember that housing enters the CPI index with a weight of 39 per cent, meaning that a bad forecast for housing will have a major impact on the overall forecasting performance in the STIF methodology.⁸

⁸ The second largest weight in the index is *food_soft_drinks*, the forecast for which is – *ex post* and frankly speaking – also a bit underwhelming.

Finally, the actual data outturns for transportation and miscellaneous_goods are rather at odds with the forecasts coming from their respective models. Figure 9 shows the graphs of the actual data versus the forecasts for these three series.

Figure 9: Forecast of transportation and miscellaneous_goods with actual outturns (January 2014 – June 2014)



Here, the actual data outturns over the first half of 2014 are very much at odds with historical price developments. These episodes are very difficult to capture with any model. On the other hand, the underlying forecasting models can be re-estimated quite easily, meaning that as soon as the new data comes in, the latest price developments can be included quite quickly.

References and further reading

- Elliott, G and Timmermann, A (2008)**, ‘Economic forecasting’, *Journal of Economic Literature*, Vol. 46, No. 1, pages 3-56.
- Timmermann, A (2006)**, ‘Forecast combinations’, Chapter 4 in Elliott, G, Granger, C W J and Timmermann, A (eds), *Handbook of economic forecasting*, Vol. 1, Amsterdam, Elsevier, pages 135-196.

Appendix

The individual components are estimated using the EViews programme called **uae_component.prg**. The programme is divided into several steps, which will be described below.

The first part estimates an autoregressive model by OLS for the food and non-alcoholic beverages CPI component series (`food_soft_drinks`) over the sample period from May 2008 to December 2013. The estimated equation is given a corresponding name for use later on, equal to `eqn_food`. The equations for the remaining components will be denoted by `eqn_X`, where X represents an appropriate label for each of the other components.

```
smpl 2008m05 2013m12
equation eqn_food.ls food_soft_drinks c @trend
food_soft_drinks(-1 to -4) d2010m12
```

The next part estimates an (autoregressive) model by OLS on the first difference of the education CPI component series (`d(education)`) over the period from January 2009 to December 2013. Again, the estimated equation is given a corresponding name, in this case `eqn_edu`.

```
smpl 2009m01 2013m12
equation eqn_edu.ls d(education) c dm09
```

The next part estimates autoregressive models by OLS for another ten CPI component series (alcoholic beverages and tobacco (`beverages_tobacco`); transportation (`transportation`); recreation and culture (`recreation_culture`); communications (`communications`); miscellaneous goods and services (`miscellaneous_goods`); textiles, clothing and footwear (`textiles_clothing`); furniture and household goods (`furniture_households`); medical care (`medical_care`) and restaurants and hotels (`restaurants_hotels`)) over their relevant sample periods, with the starting date ranging from May 2009 to June 2012. Again, each estimated equation is given a corresponding name.

```
smpl 2009m05 2013m12
equation eqn_bevs.ls beverages_tobacco c @trend
beverages_tobacco(-1 to -3) d2011m09 d2012m11
```

```
smpl 2010m09 2013m12
equation eqn_trans.ls transportation c @trend
transportation(-1 to -3)
```

```
smpl 2011m03 2013m12
equation eqn_rec.ls recreation_culture c
recreation_culture(-1) d2012m09
```

```
smpl 2011m07 2013m12
```

```

equation eqn_comms.ls communications c @trend
communications(-1 to -2)

smpl 2011m09 2013m12
equation eqn_misc.ls miscellaneous_goods c
miscellaneous_goods(-1) d2013m06

smpl 2011m10 2013m12
equation eqn_text.ls textiles_clothing c @trend
textiles_clothing(-1)

smpl 2012m02 2013m12
equation eqn_furn.ls d(furniture_household) c @trend
d(furniture_household(-1)) d(furniture_household(-2))

smpl 2012m02 2013m12
equation eqn_med.ls medical_care c @trend medical_care(-
1) d2013m11

smpl 2012m06 2013m12
equation eqn_rest.ls restaurants_hotels c @trend
restaurants_hotels(-1)

```

The twelfth and final OLS estimation is for the housing CPI component series (housing). This autoregressive model uses the shortest estimation period, which runs from November 2012 to December 2013. Again, the estimated equation is given a corresponding name, in this case `eqn_hous`.

```

smpl 2012m11 2013m12
equation eqn_hous.ls housing c @trend housing(-1)

```

The individual component equations that we have just estimated are then forecast over the period from January 2014 to June 2014 using the EViews programme called **uae_component_forecast.prg**. The programme is again divided into several steps, which will be described below. Note that the forecasts could be more efficiently produced in a loop. Each of the twelve forecasts generates a point mean forecast, which will be labelled by the suffix `f`, as well as the forecast standard error, which is labelled by the suffix `se`. Using the EViews command `group`, we then create a group consisting of the point mean forecast as well as the forecast confidence interval, defined as the point mean forecast plus and minus one standard error. Each group is given a short name, such as `be` for `beverages_tobacco`, and the final line associated with each CPI component forecast plots the contents of each group, i.e, the point mean forecast and the plus and minus one-standard-error bands, as a line graph.

```

smpl 2014m01 2014m06

eqn_bevs.forecast bevsf bevsse
group be bevsf (bevsf+bevsse) (bevsf-bevsse)

```

be.line

eqn_edu.forecast eduf eduse
group ed eduf (eduf+eduse) (eduf-eduse)
ed.line

eqn_comms.forecast commsf commsse
group co commsf (commsf+commsse) (commsf-commsse)
co.line

eqn_food.forecast foodf foodse
group fo foodf (foodf+foodse) (foodf-foodse)
fo.line

eqn_hous.forecast housf housse
group ho housf (housf+housse) (housf-housse)
ho.line

eqn_furn.forecast furnf furnse
group fu furnf (furnf+furnse) (furnf-furnse)
fu.line

eqn_trans.forecast transf transse
group tr transf (transf+transse) (transf-transse)
tr.line

eqn_med.forecast medf medse
group me medf (medf+medse) (medf-medse)
me.line

eqn_misc.forecast miscf miscse
group mi miscf (miscf+miscse) (miscf-miscse)
mi.line

eqn_rec.forecast recf recse
group rel recf (recf+recse) (recf-recse)
rel.line

eqn_rest.forecast restf restse
group re2 restf (restf+restse) (restf-restse)
re2.line

eqn_text.forecast textf textse
group te textf (textf+textse) (textf-textse)
te.line

smpl 2011m01 2014m06

The individual forecasts for the twelve components that we have just produced are then combined using the EViews programme called **uae_cpi_forecast.prg**. The programme is again divided into several steps, which will be described below. Note that the individual forecasts are combined using the original CPI weights. The programme defines the weights, specifies the model and solves for the forecast and then calculates and plots the inflation rate we just generated (`dcpi_0`) together with the actual composite CPI inflation rate (`cpi_inflation`).

```
vector(12) weights
weights.fill 13.936, 0.218, 7.580, 39.334, 4.208, 1.124,
9.941, 6.932, 3.067, 4.004, 4.348, 5.308
```

We define an empty vector consisting of twelve elements using the EViews command `vector(12)` and call it `weights`. We then collect the twelve weights of the individual components and fill the vector `weights` with them. The use of the vector `weights` will become more obvious below.

```
model stif
```

Forecasting is most easily done in an EViews object called a **model**. The above line creates a new EViews model object called `stif`.

```
stif.append composite = (weights(1)*foodf +
weights(2)*bevsv + weights(3)*textf +
weights(4)*housf + weights(5)*furnf +
weights(6)*medf + weights(7)*transf +
weights(8)*commsf + weights(9)*recf +
weights(10)*eduf + weights(11)*restf +
weights(12)*miscf)/100
```

The section above generates the model called `stif` as consisting of the individual forecasts weighted by the respective entries of the `weights` vector. The advantage of using `weights` is that we do not have to hard code the weights in the programme. For example, the model could have alternatively been defined by the following command:

```
stif.append composite = (13.936*foodf + 0.218*bevsv +
7.580*textf + 39.334*housf + 4.208*furnf +
1.124*medf + 9.941*transf + 6.932*commsf +
3.067*recf + 4.004*eduf + 4.348*restf +
5.308*miscf)/100
```

But CPI weights are regularly updated. Using the `weights` vector, we only need to update the CPI weights once in the `weights` vector at the beginning of the programme rather than throughout.

```
smpl 2010m01 2014m06
stif.solve
```

The `stif` model is then ‘solved’ over the period from January 2010 to June 2014. This will generate a new series called `composite_0`, which is the ‘forecast’ of the `composite` series defined above over the period from January 2010 to June 2014. As such, we would expect this series to equal the actual inflation rate over that part of the sample for which actual data is available.

```
smp1 2011m01 2014m06
genr dcpi_0 = 100*(composite_0-composite_0(-
12))/composite_0(-12)
genr cpi_inflation = 100*(all_items_cpi - all_items_cpi(-
12))/all_items_cpi(-12)
```

This part of the programme calculates the annual inflation rates from January 2011 to June 2014 (`dcpi_0`) using the forecasted composite CPI series (`composite_0`) based on the twelve underlying CPI components. It also generates the actual (headline) inflation rate (`cpi_inflation`) from the all-items CPI index over the same period.

```
group inflation dcpi_0 cpi_inflation
inflation.line
```

The final part of the programme groups the series we have just created (`dcpi_0`) and the actual inflation rate (`cpi_inflation`) into a group called `inflation` and plots them. The end result will be Figure 6 above.