Lecture Note of Bus 41202, Spring 2008: Univariate Volatility Models. Mr. Ruey Tsay

Out-of-Sample Forecasts

Model selection depends not only on the *goodness-of-fit* of a model to the data, but also on the *objective* of the analysis. For forecasting, a model that is best in the in-sample fitting does not necessarily provide more accurate forecasts. Thus, many people use the performance of *out-of-sample* forecasts to aid the selection of a statistical model. By *out-of-sample*, I meant that the data used in model fitting are different from those used in forecasting evaluation. Typically, one divides the data into two subperiods. Thr first period is for model fitting and is called *estimation subsample* whereas the second period consists of data used to evaluate the forecasting performance and is called *forecasting subsample*.

Given T data points, say $r_1, \ldots, r_T$, we divide the data as $\{r_1, \ldots, r_n\}$ and $\{r_{n+1}, \ldots, r_T\}$, where n is the initial forecast origin. A reasonable choice is $n = T/2$ for large T or $n = 2T/3$, otherwise. Suppose there are two competing models, say M_1 and M_2. Let h be the forecast horizon, i.e. we are interested in 1-step to h-step ahead forecasts. The out-of-sample forecasting evalaution works as follows:

1. Let $m = n$ be the initial forecast origin. Fit model M_1 and M_2 using the data $\{r_1, \ldots, r_m\}$. For each fitted model, compute 1-step to h-step ahead forecasts at the forecast origin m.

   (a) For model M_1, denote the forecasts as $r_{1,m}(1), \cdots, r_{1,m}(h)$. Compute the forecast errors as

   $$e_{1,m}(1) = r_{m+1} - r_{1,m}(1), e_{1,m}(2) = r_{m+2} - r_{1,m}(2), \cdots,$$

$$e_{1,m}(h) = r_{m+h} - r_{1,m}(h).$$

(b) For model M_2, denote the forecasts as $r_{2,m}(1), \cdots, r_{2,m}(h)$. Compute the forecast errors as

$$e_{2,m}(1) = r_{m+1} - r_{2,m}(1), e_{2,m}(2) = r_{m+2} - r_{2,m}(2), \cdots,$$

$$e_{2,m}(h) = r_{m+h} - r_{2,m}(h).$$

2. Advance the forecast origin by 1, i.e. $m = m+1$, and go to Step 1.

3. The iteration stops when the forecast orign is $m = T$.

In this way, we will have $(T-n-1)$ 1-step ahead forecast errors for each model, and $(T-n-2)$ 2-step ahead forecast errors for each model, and so on. We then compute the root-mean squared forecast errors for ℓ-step ahead forecasts of the model M_j as

$$\text{RSME}_j(\ell) = \sqrt{\sum_{i=m}^{T-\ell} e_{j,i}^2(\ell)}$$

where $\ell = 1, 2, \ldots, h$ and $j = 1, 2$.
For 1-step ahead forecasts, we select model M_1 over model M_2 if $\text{RMSE}_1(1) < \text{RMSE}_2(1)$; otherwise, model M_2 is preferred. For 2-step ahead forecasts, we select model M_1 over model M_2 if $\text{RMSE}_1(2) < \text{RMSE}_2(2)$; otherwise, M_2 is preferred. The same rule applies to other forecast horizons.

I wrote two **R** scripts to perform out-of-sample forecasts. The first script is called **r-fore.txt**, which perform forecasts of a fitted ARIMA model at a specified forecast origin. The second script is called **r-backtest.txt**, which compute the RMSE of out-of-sample forecasts for a fitted ARIMA model with a given initial forecast origin n.

These two scripts are available on the course web page. You may download and store in your working directory (just as you stored the data sets). You can then use the command

```
source(‘‘r-fore.txt’’)
source(‘‘r-backtest.txt’’)
```

to compile the scripts in R. The commands are called **forecast** and **backtest**, respectively. For demonstration, consider the weekly mortgage rate series of Q1 in HW#3.

```
> x=read.table("w-mortg.txt")
> mort=x[,4]

> source("r-fore.txt")
> length(mort)
[1] 222

> m1=arima(mort,order=c(1,1,0))
> predict(m1,4)
$pred
Time Series:
Start = 223
End = 226
Frequency = 1
[1] 5.886363 5.887713 5.887999 5.888060

$se
Time Series:
Start = 223
End = 226
```

```
Frequency = 1
[1] 0.07958213 0.12505259 0.16014524 0.18922831

> forecast(m1,mort,200,4)
Time Series:
Start = 201
End = 204
Frequency = 1
[1] 6.240670 6.235333 6.233859 6.233452
Time Series:
Start = 201
End = 204
Frequency = 1
[1] 0.06734154 0.10917899 0.14217644 0.16961281
>
> source("r-backtest.txt")
> backtest(m1,mort,200,4)
[1] "RMSE of out-of-sample forecasts"
[1] 0.1518742 0.2279446 0.2950994 0.3531029

> m2=arima(mort,order=c(0,1,1))
> backtest(m2,mort,200,4)
[1] "RMSE of out-of-sample forecasts"
[1] 0.1510807 0.2254471 0.2902807 0.3430948
```