

Forecasting: principles and practice

Exercises: Set 10

20 November 2013

Before doing any exercises in R, load the **fpp** package using `library(fpp)`.

1. Choose one of the following seasonal time series: `condmilk`, `hsales`, `uselec`
 - (a) Do the data need transforming? If so, find a suitable transformation.
 - (b) Are the data stationary? If not, find an appropriate differencing which yields stationary data.
 - (c) Identify a couple of ARIMA models that might be useful in describing the time series. Which of your models is the best according to their AIC values?
 - (d) Estimate the parameters of your best model and do diagnostic testing on the residuals. Do the residuals resemble white noise? If not, try to find another ARIMA model which fits better.
 - (e) Forecast the next 24 months of data using your preferred model.
 - (f) Compare the forecasts obtained using `ets()`.

2. For the same time series you used in Q1, try using a non-seasonal model applied to the seasonally adjusted data obtained from STL. The `stlf()` function will make the calculations easy (with `method="arima"`). Compare the forecasts with those obtained in Q1. Which do you think is the best approach?

3. For the time series you selected from the retail data set in Exercise 7.3, develop an appropriate seasonal ARIMA model, and compare the forecasts with those you obtained earlier.

Obtain up-to-date data from January 2008 onwards from the ABS website (www.abs.gov.au) (Cat. 8501.0, Table 11), and compare your forecasts with the actual numbers. How good were the forecasts from the two models?