Forecasting: principles and practice

Exercises: Set 4 30 October 2013

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

1. For each of the following time series: `bricksq`, `ibmclose`, `hsales`.
   (a) Use the preferred forecasting method you identified for that time series in the Exercise Set 2 and apply it to the full data set.
   (b) Compute the residuals and plot the ACF. Do the residuals appear to be white noise?
   (c) Do a Ljung-Box test on the residuals. What do the results mean?
2. The data below represent the monthly sales (in thousands) of product A for a plastics manufacturer for years 1 through 5 (data set `plastics`).

	1	2	3	4	5
Jan	742	741	896	951	1030
Feb	697	700	793	861	1032
Mar	776	774	885	938	1126
Apr	898	932	1055	1109	1285
May	1030	1099	1204	1274	1468
Jun	1107	1223	1326	1422	1637
Jul	1165	1290	1303	1486	1611
Aug	1216	1349	1436	1555	1608
Sep	1208	1341	1473	1604	1528
Oct	1131	1296	1453	1600	1420
Nov	971	1066	1170	1403	1119
Dec	783	901	1023	1209	1013

   (a) Plot the time series of sales of product A. Can you identify seasonal fluctuations and/or a trend?
   (b) Use an STL decomposition to calculate the trend-cycle and seasonal indices. (Experiment with having fixed or changing seasonality.)
   (c) Do the results support the graphical interpretation from part (a)?
   (d) Compute and plot the seasonally adjusted data.
   (e) Change one observation to be an outlier (e.g., add 500 to one observation), and recompute the seasonally adjusted data. What is the effect of the outlier?
   (f) Does the outlier have an effect if you use `robust=TRUE` with the `stl` function?
   (g) Does it make any difference if the outlier is near the end rather than in the middle of the time series?
   (h) Use a random walk to produce forecasts of the seasonally adjusted data.
   (i) Reseasonalize the results to give forecasts on the original scale.

   [Hint: you can use the `stlf` function with `method="naive"`.]