

In HW 4, you were asked to identify a small set of candidate  $ARIMA(p, d, q)$  models (perhaps some of you even made a “final choice”) for each of the following data sets:

- **ibm**: daily closing IBM stock prices (dates not given)
- **internet**: number of users logged on to an Internet server each minute.
- **robot**: final horizontal position of an industrial robot put through a series of planned exercises.

Remembering your candidate models for each data set, fit and diagnose your model selections. That is, use the methods from Chapter 7 to fit your chosen models (for uniformity, you could just use maximum likelihood for each model fit). Then, diagnose your fitted model(s) by doing a thorough analysis of the residuals and implementing the overfitting technique (Chapter 8).

For each data set and for each model you entertained in HW 4, what do you think now? Would you like to suggest another model for further investigation? Or, are you satisfied with your HW 4 model selections?

- If your original model choices in HW 4 are “reasonable,” convince me that they are.
- If your original model choices are deemed “not reasonable,” use the information from your diagnoses to specify another model. Then, evaluate the merit of this new model using the methods from Chapter 7 and Chapter 8.

Your goal is to come up with **one final model** for each data set—the “best” one. Convince me that your final model does a good job at explaining the variability in the data, but also adhere to the Principle of Parsimony. There are no “right” answers here, but there are certainly bad answers (stay away from these).

**Important:** Remember your final model for each data set.