Exercise: exponential smoothing I

Earlham College
BUS 323 - Fall 2025 - Labadie

In this activity, you will practice employing exponential smoothing methods using datasets in the `fpp3` package, and evaluating forecast accuracy based on forecast errors. Please turn in your code with answers in comments, or a document with your answers and plots (Word doc or a PDF generated from Markdown, for example).

1. We will use the `WWWusage` dataset, which contains number of users connected to a website server per minute over 100 minutes.

   (a) Use the following code to make the data workable: `www_usage <- as_tsibble(WWWusage)`

   (b) Plot the historical data.

   (c) Use `stretch_tsibble()` to create cross-validation training sets with initial size of 10.

   (d) Estimate a simple exponential smoothing model, a trended exponential smoothing model, and a damped trended exponential smoothing model.

   (e) Evaluate the accuracy of each model based on 1-step forecasts using cross-validation. Which model performs best?

   (f) Produce a 10-step forecast using the model of your choice. Produce a plot including your forecast and prediction interval.