Exercise: simple forecasting methods

In this exercise, you'll use some of the simple forecasting methods discussed in class 9/25 (mean forecasting, naïve forecasting, seasonal naïve forecasting, and drift forecasting) with datasets included in the `fpp3` package.

1. Use the `aus_production` dataset to produce forecasts for Australian beer production.
    a. First, filter the data down to reduce the scope of the training data to 1992 Q1 to 2006 Q4 (hint: use `filter_index()`). Feel free to also select the variable `Beer` to further reduce the scope of the training dataset (hint: use `select()`).
    b. Fit a mean model; a naïve model; and a seasonal naïve model to the training data (hint: use `model(MEAN())`, `model(NAIVE())`, and `model(SNAIVE())`).
    c. Generate forecasts for 14 future quarters (hint: use `forecast()`).
    d. Plot all three forecasts with the time plot for the historical data. Add in the actual observed values for 2007 Q1 and beyond. Which forecast fits the observed data best?
2. Use the `gafa_stock` dataset to produce forecasts for Google's daily closing stock price.
    a. First, filter the data down to reduce the scope of the training data to 2015 and only Google stocks (hint: use `filter()`).
    b. Fit a mean model; a naïve model; and a drift model to the training data (hint: use `model(MEAN())`, `model(NAIVE())`, and `model(RW())`).
    c. Generate forecasts for one month (so just for January 2016) (hint: use `forecast()`).
    d. Plot all three forecasts with the time plot for the historical data. Add in the actual observed values for January 2016. Which forecast fits the observed data best?