Time Series II:
Seasonally Adjusted Models and Smoothing Techniques

Alan B. Gelder
06E:071, The University of Iowa[1]

1 A Review of Forecasting

In the previous lecture, we introduced AR models as a tool for improving forecasts when we have autocorrelation. There, we focused on an example that had trend and seasonal patterns. However, we can use AR models to help improve forecasts *whenever* there is autocorrelation, even when the underlying patterns are different.

1.1 Blackberry Sales

[This example is adapted from Dr. Whitten's notes.]

For this example, we will use weekly sales data on the number of Blackberry Storm phones that were sold during the first 16 weeks after its release in 2008.[2]

- Make a time series plot of the data. Describe any patterns that you find.

Define the mean square deviation (MSD) as follows:

$$MSD = \frac{1}{n}\sum_{i=1}^{n}(y - \hat{y})^2$$

In words, MSD is equal to the average squared difference between the actual and the predicted values of y. The MSD is a measurement that we can use for assessing how closely a regression model matches our actual data; a lower MSD indicates a closer fit.

- Compare the MSD for linear, quadratic, and exponential time series models of the Blackberry Storm data. Which model is the best?

[1]The notes for this class follow the general format of Blake Whitten's lecture notes for the same course.

[2]The data is from Research in Motion, Ltd.

- Stat > Time Series > Trend Analysis > Variable: Blackberry Sales > Model Type: (Linear, Quadratic, Exponential Growth) > OK

- Check and store the residuals for your chosen model. Is there autocorrelation? How can you tell?
  - (Return to Trend Analysis) > Graphs > Residual Plots: Residuals verses order > OK > Storage > Residuals > OK

- Make a partial autocorrelation plot for the residuals.
  - Time Series > Partial Autocorrelation > Series: (Residuals) > OK

- Find the appropriate AR model. Which AR models should you check? (Use $\alpha = 0.05$)
  - Time Series > ARIMA > Series: (Residuals) > *Unselect* include constant term > Autoregressive: (Enter number of lags) > OK

- Make forecasts of the residuals for the next 5 weeks.
  - (Return to ARIMA) > Forecasts > Lead: 5 > OK > OK

- Make forecasts of $\hat{y}$ for the next 5 weeks.
  - (Return to Trend Analysis) > Generate Forecasts > Number of Forecasts: 5 > OK > OK

- Forecast the number of Blackberry Storm sales in week 19.

- Give numerical evidence that the forecast for week 19 is less precise than the forecast in week 17.

- Forecast sales for week 19 with 95% certainty. Interpret your answer.

2 Seasonally Adjusted Models

Additive method for seasonality: There are two main ways to account for seasonality in time series models. One of these two ways is to include a set of binary seasonal variables in the model, as we saw in the previous lecture. This method can be referred to as the addition method because the seasonal indicator variables enter the model additively. For quarterly data, our model is

$$\underbrace{y}_{\text{Actual}} = \underbrace{\beta_0 + \beta_1 x}_{\text{Trend}} + \underbrace{\gamma_1 Q1 + \gamma_2 Q2 + \gamma_3 Q3}_{\text{Seasonality}} + \varepsilon$$

The addition method adjusts the intercept of the model up and down with the binary seasonal variables.

Multiplicative method for seasonality: Instead of accounting for seasonality with a set of binary variables, we can begin with the trend only model and then multiply the trend by a seasonal factor:

$$\underbrace{y}_{\text{Actual}} = \underbrace{(\beta_0 + \beta_1 x)}_{\text{Trend}} \times \underbrace{(\text{Seasonal Factor})}_{\text{Seasonality}}$$

We can estimate this model as follows:

$$y = \hat{y} \times (\text{Seasonal Factor})$$

Here, $\hat{y}$ is the predicted value from the trend only model. A seasonal factor greater than 1 will increase the predicted value (e.g. in peak seasons), and a seasonal factor less than 1 will decrease the predicted value (e.g. in off-seasons). A seasonal factor of 1 can be thought of as the average seasonal effect. (Note that this coincides with the trend only model since it predicts the *average* value of y conditional on time.)

Rearranging terms identifies the seasonal factor as the ratio of the actual values to the predicted values:

$$\frac{y}{\hat{y}} = \text{Seasonal Factor}$$

This ratio will be the basis for calculating seasonal factors.

Deseasonalized values and forecasts: Whenever we have data with seasonal patterns, we want our forecasts to account for seasonality. Both the additive method and the multiplicative method allow us to *include* seasonality in our model. **(How?)**

One big advantage that the multiplicative method has over the addition method is that seasonality can also be *taken out* of the data, leaving us with seasonally adjusted or deseasonalized values. This can be done by dividing the actual values (y) by the seasonal factor.

Seasonally adjusted values enables us to see the underlying pattern of the data. Since seasonal factors in peak seasons are greater than 1, seasonally adjusted values are smaller than the actual values. Likewise, since seasonal factors in off-seasons are less than 1, seasonally adjusted values are greater than the actual values. Thus, seasonally adjusted values put the data from peak seasons on the same level as data from off-seasons, which allows us to make level comparisons across different seasons.

In summary, we want to *remove* seasonality (by dividing) in order to compare data from different seasons, and we want to *include* seasonality (by multiplying) in making forecasts.

$$\begin{aligned} \text{Seasonally Adjusted Values:} &\quad y/\,(\text{Seasonal Factor}) \\ \text{Forecasts:} &\quad \hat{y} \times (\text{Seasonal Factor}) \end{aligned}$$

Analogy with inflation: Since \$1 in 1950 is not the same as \$ 1 today, we would need to adjust for inflation if we wanted to do an accurate comparison of wealth over those different points in time. Without adjusting for inflation, a family that was well to do in 1950 may appear to be near the poverty line today. Only after we adjust for inflation are we able to compare family income in 1950 with family income today.

Adjusting for inflation enables us to equate \$1 in 1950 with \$1 today. Similarly, we can adjust for seasonality in order to compare \$1 in February retail sales with \$1 in December retail sales. The multiplicative method of accounting for seasonality enables us to equate the units that we are measuring across seasons.

2.1 Steps for Deseasonalized Data

1. Estimate the trend only model and store the fitted values.
    - Stat > Regression > Regression > Response: (Time Series Variable) > Predictors: (Time) > Storage > Fits > OK > OK
2. Divide the actual values by the fitted values (i.e. $y/\hat{y}$).

3. Average the $y/\hat{y}$ ratios for each seasonal period. This will give you the seasonal factors.
4. Divide the actual values (y) by the seasonal factors. This will give you the deseasonalized values.

2.2 Oranges

The oranges dataset contains the monthly price of oranges from January 1991 through December 2000. The prices have been adjusted for inflation so that 100 represents the average price of oranges between 1982 and 1984 (see p. 19 in the text). The price of oranges fluctuates throughout the year as they go in and out of season.

- Plot the orange prices. At what time of year are oranges the cheapest? When are they the most expensive?

The estimated trend only model for the price of oranges is

$$\hat{y} = 182.3758 + 0.8547 \times Time$$

From this estimated model, we can obtain the fitted values ($\hat{y}$) for each month. The next step is to divide the actual price of oranges by the predicted price of oranges for each month (i.e. $y/\hat{y}$). A portion of the data is given below:

Time	Month	Year	Price of Oranges	$\hat{y}$	$y/\hat{y}$
1	Jan	1991	205.7	183.2305	1.1226295
2	Feb	1991	224	184.0852	1.2168276
3	Mar	1991	235.4	184.9399	1.2728456
4	Apr	1991	245.5	185.7947	1.3213512
5	May	1991	244.4	186.6494	1.3094071
6	Jun	1991	270.8	187.5041	1.4442352
⋮	⋮	⋮	⋮	⋮	⋮
119	Nov	2000	227	284.0864	0.7990526
120	Dec	2000	214.7	284.9411	0.7534889

Now we need to average the $y/\hat{y}$ ratios for each month. Starting in January, we have the following:

Month	Year	$y/\hat{y}$
Jan	1991	1.1226295
Jan	1992	0.9695739
Jan	1993	0.7592877
Jan	1994	0.7588781
Jan	1995	0.8195966
Jan	1996	0.8182909
Jan	1997	0.7656176
Jan	1998	0.7893305
Jan	1999	1.0486922
Jan	2000	0.8728337
Average:		0.87247307

The seasonal factor for January is 0.8725. We can interpret this number as follows: the price of oranges in January is typically 87.25% of the average monthly price of oranges.

The seasonal factors for each month are given below:

Month	Seasonal Factor
January	0.8725
February	0.8832
March	0.8810
April	0.9143
May	0.9797
June	1.0471
July	1.0939
August	1.1812
September	1.2488
October	1.1486
November	0.9326
December	0.8295

Note that the average of the seasonal factors is 1.0010. Since 1 represents the average seasonal effect, the seasonal factors should always have an average that is roughly equal to 1.

The final step for obtaining seasonally adjusted orange prices is to divide the actual orange prices by the seasonal factor. This is shown in the following table:

Time	Month	Year	Price of Oranges	Seasonal Factor	Deseasonalized Price
1	Jan	1991	205.7	0.8725	235.76
2	Feb	1991	224	0.8832	253.62
3	Mar	1991	235.4	0.881	267.20
4	Apr	1991	245.5	0.9143	268.51
5	May	1991	244.4	0.9797	249.46
6	Jun	1991	270.8	1.0471	258.62
⋮	⋮	⋮	⋮	⋮	⋮
119	Nov	2000	227	0.9326	243.41
120	Dec	2000	214.7	0.8295	258.83

The seasonally adjusted data is plotted below:

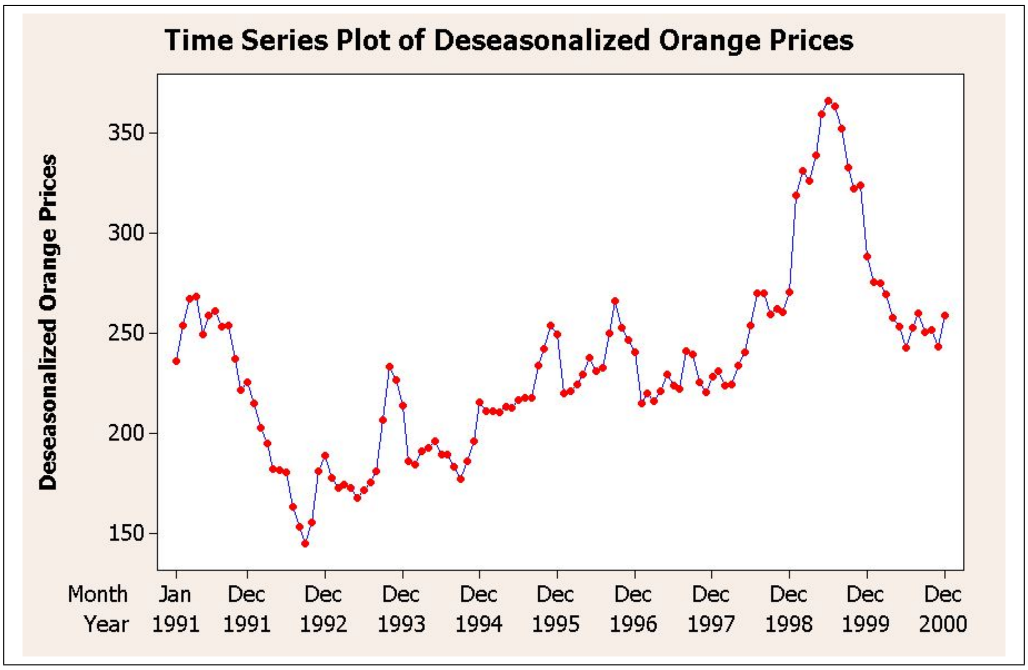

The seasonally adjusted data reveals that orange prices were exceptionally low during 1992 and quite high during 1999. These two price variations are dramatic enough that they can be seen in the original data (although not as clearly as in the deseasonalized data). One thing that the seasonally adjusted data captures that is not clear in the original data is that the price spike at the end of 1993 was unseasonably large. This price spike looks like most of the others price spikes in the original data. However, the seasonally adjusted data *also* shows a spike, which means that the underlying price of oranges actually went up.

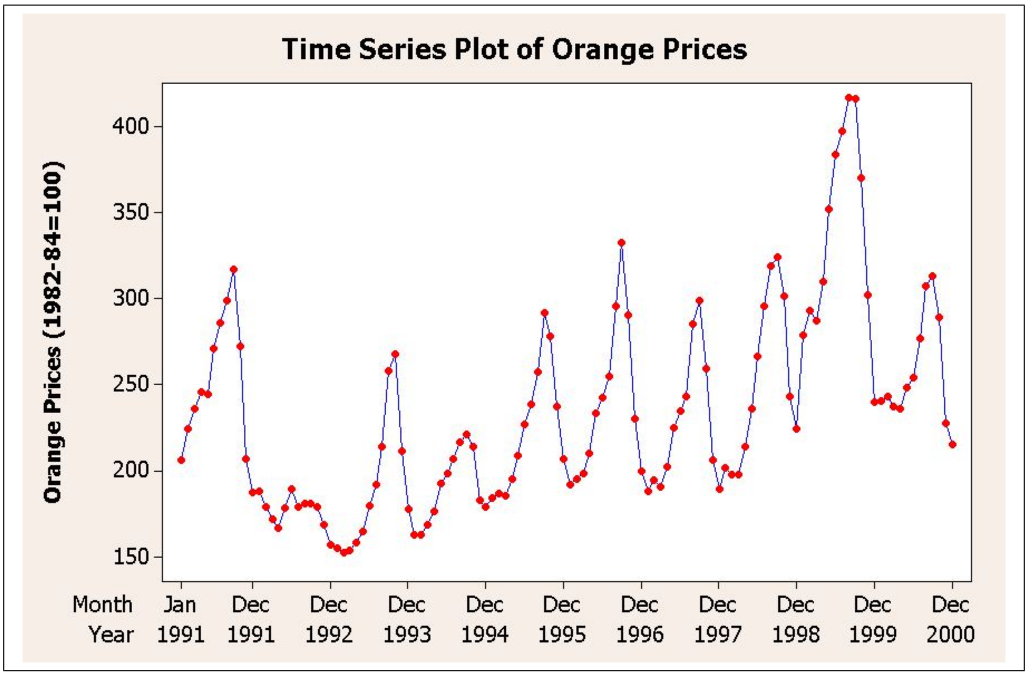

2.3 Time Series Decomposition

Minitab has a procedure for plotting seasonally adjusted values. Unfortunately, **Minitab has a different method for computing seasonal factors**. Although the method is different, Minitab can quickly produce similar results. The commands are as follows:

- Stat > Time Series > Decomposition > Variable: (Enter time series variable) > Seasonal length: (12 for monthly data, 4 for quarterly data, etc.) > OK

The output for the Time Series Decomposition contains the original data, the deseasonalized data, the detrended data, and the data that is left over after both the trend and the seasonality have been removed (i.e. the residuals of the trend and seasonal model). The next figure contains the Time Series Decomposition for the orange price data:

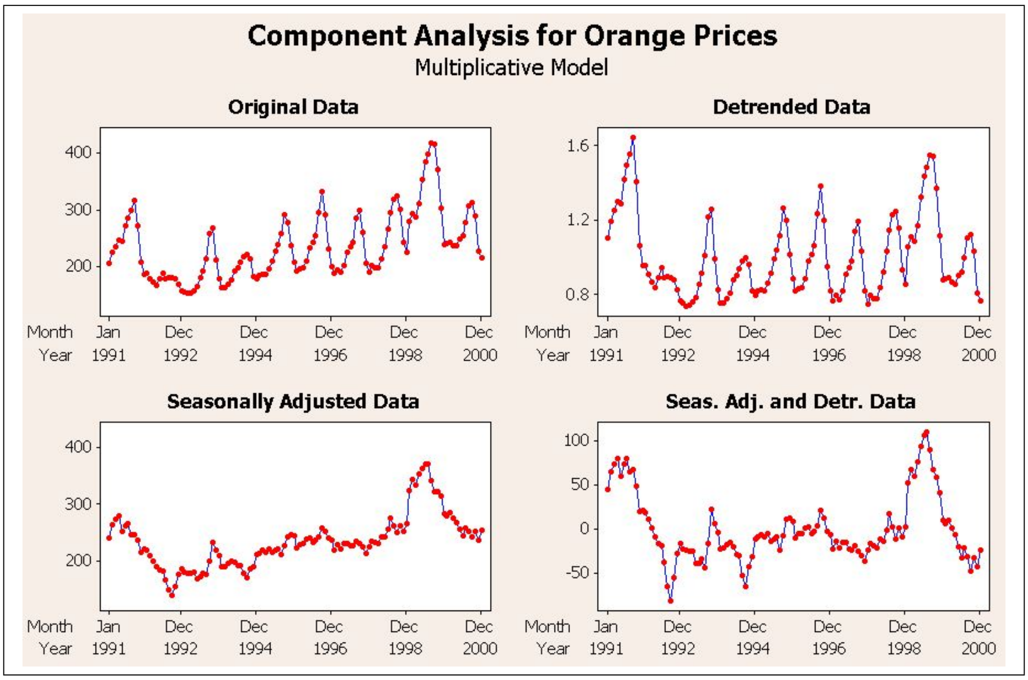

- Compare the seasonal factors from Minitab for the orange price data with the seasonal factors in these notes.

3 Smoothing Techniques

In addition to forecasting, another common objective of time series is to identify long run patterns. Time series data is often "noisy" in the sense that there are lots of fluctuations over time that can mask long run patterns. We will discuss two smoothing techniques that can be used to uncover long run patterns: moving averages and exponential smoothing.

3.1 Moving Averages

Moving averages is probably one of the easiest statistical techniques that we cover in this course. We first pick our smoothing length. For monthly data, a smoothing length of 12 indicates that we our smoothing over the course of one year. Next we make moving average forecasts for each month (beginning at the 13th month of our data) according to the following formula:

$$\hat{y}_t = \frac{y_{t-1} + y_{t-2} + y_{t-3} + \ldots + y_{t-11} + y_{t-12}}{12}$$

This is referred to as a 12 month moving average. The forecast for month 13, $\hat{y}_{13}$, is equal to the average value of y for the first 12 months. The forecast for month 14, $\hat{y}_{14}$, is equal to the average value of y from months 2 through 13. Similarly, $\hat{y}_{15}$ equals the average of y from month 3 through month 14. For any month, the forecast $\hat{y}$ is equal to the average of the previous 12 months.

In general, the formula for moving averages is

$$\hat{y}_t = \frac{y_{t-1} + y_{t-2} + y_{t-3} + \ldots + y_{t-k}}{k}$$

where k is the smoothing length.

Forecasting with moving averages: Moving averages can be used for forecasting. However, it often does a poor job forecasting since the values for y from periods in the distant past are weighted the same as periods in the recent past (i.e. $1/k$ of $\hat{y}_t$ comes from y_{t-k}, while another $1/k$ of $\hat{y}_t$ comes from y_{t-1}). This prevents moving average forecasts from closely following any sudden changes (such as price spikes) in the data. However, as a smoothing technique, sudden changes in the data are precisely the thing that moving averages glosses over. If a change is persistent over a long period of time, then moving averages will account for it.

3.2 Exponential Smoothing

Exponential smoothing is another technique that can be used to identify long run patterns in time series data. Unlike moving averages, exponential smoothing forecasts can place more weight on the recent past and less weight on the distant past. Another feature of exponential smoothing is that forecasts take into account *all* previous periods (not just the k previous periods).

The formula for moving averages is

$$\hat{y}_t = \omega y_{t-1} + (1 - \omega)\hat{y}_{t-1}$$

ω is called the *smoothing constant* and is a number between 0 and 1. The smoothing constant determines how much weight should be placed on the previous period (a higher ω indicates more weight on y_{t-1}). The remainder of the weight is given to the forecast for the previous period, $\hat{y}_{t-1}$, which is in turn made up of the actual value for $t-2$ and the forecast for $t-2$. (Then, of course, the forecast for $t-2$ is made up of the actual value for $t-3$ and the forecast for $t-3$, etc.)

We can see how exponential smoothing forecasts incorporate all of the previous observations as follows:

$$
\begin{aligned}
\hat{y}_t &= \omega y_{t-1} + (1-\omega)\hat{y}_{t-1} \\
&= \omega y_{t-1} + (1-\omega)(\omega y_{t-2} + (1-\omega)\hat{y}_{t-2}) \\
&= \omega y_{t-1} + (1-\omega)(\omega y_{t-2} + (1-\omega)(\omega y_{t-3} + (1-\omega)\hat{y}_{t-3})) \\
&= \omega y_{t-1} + (1-\omega)(\omega y_{t-2} + (1-\omega)(\omega y_{t-3} + (1-\omega)(\omega y_{t-4} + (1-\omega)\hat{y}_{t-4})))
\end{aligned}
$$

$$
\Rightarrow \hat{y}_t = \omega y_{t-1} + \omega(1-\omega)y_{t-2} + \omega(1-\omega)^2 y_{t-3} + \omega(1-\omega)^3 y_{t-4} + (1-\omega)^4\hat{y}_{t-4}
$$

A pattern is clearly emerging here. However, we need to make an assumption about what the exponential smoothing forecast is for the very first time period in our data. We will assume that

$$\hat{y}_1 = y_1$$

In other words, our forecast for the first period is equal to the actual value of the first period.

Now we can write a general formula for the exponential smoothing forecast at time t that incorporates all of the previous periods. If there are n periods in our data before period t (i.e. t is period $n+1$), then the exponential smoothing forecast at period t is given as follows:

$$\hat{y}_t = \left(\sum_{i=1}^{n} \omega(1-\omega)^{i-1} y_{t-i} \right) + (1-\omega)^n y_{t-n}$$

3.3 Winter Wheat

[This example is adapted from problems 13.28 and 13.29]

The winter wheat data contains the monthly price of a bushel of Montana winter wheat from July 1929 through October 2002.

- Plot the winter wheat data. Is there a trend? Is there seasonality? Are other patterns present in the data?

- Find and plot a 12 month moving average and a 120 month moving average for the winter wheat data.

  – Stat > Time Series > Moving Average > Variable: Dollars-per-bushel > MA length: (12 or 120) > Storage > Fits (one period ahead forecasts) > OK > OK

- What elements of the original data does each moving average capture? What features does each moving average smooth?

- Find and plot the exponential smoothing models with smoothing constants 0.2, 0.5, and 0.8. What do you notice about the different smoothing constants?

  – Stat > Time Series > Single Exp Smoothing > Variable: Dollars-per-bushel > Weight to Use in Smoothing > Use: (Enter smoothing constant) > Storage > Fits (one period ahead forecasts) > OK > OK

- Make forecasts for November 2002 using the 12 month moving average, the 120 month moving average, and the 0.5 exponential smoothing model. Looking at the data, which forecast seems to be the most reasonable?

  – Moving Average OR Single Exp Smoothing > Generate Forecasts > Number of forecasts: 1 > Storage > Forecasts > OK > OK