Chapter **Three**

Extrapolation 1. Moving Averages and Exponential Smoothing

©VLADGRIN/Getty Images

Consider the situation facing a manager who must periodically forecast the inventories for hundreds or thousands of products. Each day or week or month, updated forecasts for the many inventories are required within a short time period. While it might well be possible to develop sophisticated forecasting models for each of the items, in many cases some very simple short-term forecasting tools are adequate for the job.

A manager facing such a task is likely to use some form of time-series *smoothing*. All the time-series smoothing methods use a form of weighted average of past observations to smooth up-and-down movements, that is, some statistical method of suppressing short-term fluctuations. The assumption underlying these methods is that the fluctuations in past values represent random departures from some smooth curve that, once identified, can plausibly be extrapolated into the future to produce a forecast or series of forecasts.

We will examine five basic smoothing techniques in this chapter. All five of these have the common characteristic that only a past history of the time series to be forecast is necessary to produce the forecast. Further, all are based on the concept that there is some underlying pattern to the data; that is, all time-series data to be forecast are assumed to have some cycles or fluctuations that tend to recur. The five methods, to be examined in turn, are:

1. Moving averages
2. Simple exponential smoothing
3. Holt's exponential smoothing
4. Winters' exponential smoothing
5. Adaptive–response-rate single exponential smoothing

LEARNING OBJECTIVES

After studying this chapter, you should be able to:

1. Explain what a moving average is and how to use one to make a forecast.
2. Distinguish between simple moving averages and exponential smoothing models.
3. Explain why the term *exponential* is used in naming exponential smoothing models.
4. Explain when simple exponential smoothing would be an appropriate forecast method.
5. Explain when Holt's exponential smoothing would be an appropriate forecast method.
6. Explain when Winter's exponential smoothing would be an appropriate forecast method.
7. Explain some methods that can be used to forecast new products.
8. Explain what an event model is and when such a model might be used to make a forecast.

MOVING AVERAGES

The simple statistical method of moving averages may mimic some data better than a complicated mathematical function.

The simple statistical method of moving averages may mimic some data better than a complicated mathematical function. Figure 3.1 shows the exchange rate between the Japanese yen and the U.S. dollar from 1980Q1 through 2016Q4. Figure 3.1 does not exhibit a simple linear, exponential, or quadratic trend similar to those we will examine in Chapters 4 and 5. Instead, the series appears to show

FIGURE 3.1 **Exchange Rate of the Japanese Yen Against the U.S. Dollar** (c3t1)

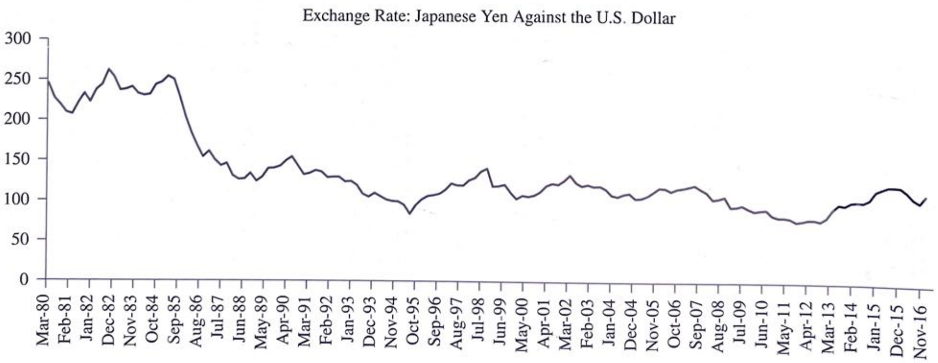

substantial randomness, which we may be able to eliminate with a technique that averages the most recent values.

To illustrate how a moving average is used, consider Table 3.1, which contains data for the exchange rate between the Japanese yen and one U.S. dollar, as shown in Figure 3.1. To calculate the three-quarter moving average first requires

TABLE 3.1 **Exchange Rate: Japanese Yen Against the US Dollar** (c3t1)

Date	Exchange Rate: Japanese Yen Against the US Dollar	Three Quarter Moving Average (MA3)	Forecast Based on MA3	Five Quarter Moving Average (MA5)	Forecast Based on MA5
Mar-80	245.52	NA	NA	NA	NA
Jun-80	227.27	NA	NA	NA	NA
Sep-80	219.17	230.65	NA	NA	NA
Dec-80	209.79	218.74	230.65	NA	NA
Mar-81	207.73	212.23	218.74	221.89	NA
Jun-81	221.45	212.99	212.23	217.08	221.89
Sep-81	233.40	220.86	212.99	218.31	217.08
Dec-81	223.13	225.99	220.86	219.10	218.31
Mar-82	238.09	231.54	225.99	224.76	219.10
Jun-82	244.35	235.19	231.54	232.08	224.76
Sep-82	262.49	248.31	235.19	240.29	232.08
Dec-82	253.79	253.54	248.31	244.37	240.29
Mar-83	237.39	251.22	253.54	247.22	244.37
Jun-83	238.59	243.26	251.22	247.32	247.22
Sep-83	241.59	239.19	243.26	246.77	247.32
Dec-83	233.26	237.81	239.19	240.92	246.77
Mar-84	231.00	235.28	237.81	236.37	240.92
Jun-84	232.01	232.09	235.28	235.29	236.37
Sep-84	244.23	235.75	232.09	236.42	235.29
Dec-84	247.43	241.22	235.75	237.59	236.42
Mar-85	255.07	248.91	241.22	241.95	237.59
Jun-85	250.67	251.06	248.91	245.88	241.95
Sep-85	230.02	245.25	251.06	245.48	245.88
Dec-85	204.66	228.45	245.25	237.57	245.48
Mar-86	184.33	206.34	228.45	224.95	237.57
Jun-86	167.99	185.66	206.34	207.53	224.95
Sep-86	154.63	168.99	185.66	188.33	207.53
Dec-86	161.31	161.31	168.99	174.59	188.33
Mar-87	150.38	155.44	161.31	163.73	174.59
Jun-87	143.64	151.78	155.44	155.59	163.73
Sep-87	146.09	146.70	151.78	151.21	155.59
Dec-87	131.18	140.30	146.70	146.52	151.21
Mar-88	126.62	134.63	140.30	139.58	146.52

TABLE 3.1 (continued)

Date	Exchange Rate: Japanese Yen Against the US Dollar	Three Quarter Moving Average (MA3)	Forecast Based on MA3	Five Quarter Moving Average (MA5)	Forecast Based on MA5
Jun-88	127.29	128.37	134.63	134.97	139.58
Sep-88	134.04	129.32	128.37	133.05	134.97
Dec-88	124.30	128.54	129.32	128.69	133.05
Mar-89	129.61	129.32	128.54	128.37	128.69
Jun-89	139.76	131.22	129.32	131.00	128.37
Sep-89	140.68	136.68	131.22	133.68	131.00
Dec-89	142.94	141.13	136.68	135.46	133.68
Mar-90	150.03	144.55	141.13	140.60	135.46
Jun-90	154.50	149.16	144.55	145.58	140.60
Sep-90	143.30	149.28	149.16	146.29	145.58
Dec-90	132.33	143.38	149.28	144.62	146.29
Mar-91	134.33	136.65	143.38	142.90	144.62
Jun-91	137.82	134.83	136.65	140.46	142.90
Sep-91	135.86	136.00	134.83	136.73	140.46
Dec-91	128.76	134.14	136.00	133.82	136.73
Mar-92	129.34	131.32	134.14	133.22	133.82
Jun-92	129.14	129.08	131.32	132.18	133.22
Sep-92	123.35	127.27	129.08	129.29	132.18
Dec-92	124.19	125.56	127.27	126.95	129.29
Mar-93	119.21	122.25	125.56	125.04	126.95
Jun-93	108.26	117.22	122.25	120.83	125.04
Sep-93	105.11	110.86	117.22	116.02	120.83
Dec-93	109.76	107.71	110.86	113.31	116.02
Mar-94	105.55	106.81	107.71	109.58	113.31
Jun-94	101.59	105.64	106.81	106.06	109.58
Sep-94	99.30	102.15	105.64	104.26	106.06
Dec-94	98.72	99.87	102.15	102.98	104.26
Mar-95	94.93	97.65	99.87	100.02	102.98
Jun-95	83.90	92.51	97.65	95.69	100.02
Sep-95	94.81	91.21	92.51	94.33	95.69
Dec-95	102.06	93.59	91.21	94.88	94.33
Mar-96	106.11	100.99	93.59	96.36	94.88
Jun-96	107.35	105.17	100.99	98.85	96.36
Sep-96	108.99	107.48	105.17	103.86	98.85
Dec-96	114.48	110.27	107.48	107.80	103.86
Mar-97	122.05	115.17	110.27	111.79	107.80
Jun-97	119.19	118.57	115.17	114.41	111.79
Sep-97	119.39	120.21	118.57	116.82	114.41
Dec-97	125.89	121.49	120.21	120.20	116.82

(continued on next page)

TABLE 3.1 (continued)

Date	Exchange Rate: Japanese Yen Against the US Dollar	Three Quarter Moving Average (MA3)	Forecast Based on MA3	Five Quarter Moving Average (MA5)	Forecast Based on MA5
Mar-98	129.06	124.78	121.49	123.12	120.20
Jun-98	137.03	130.66	124.78	126.11	123.12
Sep-98	140.36	135.49	130.66	130.35	126.11
Dec-98	118.01	131.80	135.49	130.07	130.35
Mar-99	119.03	125.80	131.80	128.70	130.07
Jun-99	120.47	119.17	125.80	126.98	128.70
Sep-99	110.75	116.75	119.17	121.73	126.98
Dec-99	103.07	111.43	116.75	114.27	121.73
Mar-00	107.47	107.10	111.43	112.16	114.27
Jun-00	106.23	105.59	107.10	109.60	112.16
Sep-00	107.88	107.19	105.59	107.08	109.60
Dec-00	111.65	108.58	107.19	107.26	107.08
Mar-01	119.19	112.91	108.58	110.48	107.26
Jun-01	122.34	117.73	112.91	113.46	110.48
Sep-01	120.95	120.83	117.73	116.40	113.46
Dec-01	125.74	123.01	120.83	119.97	116.40
Mar-02	133.10	126.60	123.01	124.26	119.97
Jun-02	123.78	127.54	126.60	125.18	124.26
Sep-02	119.65	125.51	127.54	124.64	125.18
Dec-02	121.16	121.53	125.51	124.68	124.64
Mar-03	118.65	119.82	121.53	123.27	124.68
Jun-03	119.19	119.67	119.82	120.49	123.27
Sep-03	115.99	117.94	119.67	118.93	120.49
Dec-03	108.49	114.55	117.94	116.70	118.93
Mar-04	106.56	110.35	114.55	113.78	116.70
Jun-04	109.41	108.15	110.35	111.93	113.78
Sep-04	110.77	108.91	108.15	110.24	111.93
Dec-04	103.98	108.05	108.91	107.84	110.24
Mar-05	105.12	106.62	108.05	107.17	107.84
Jun-05	107.95	105.68	106.62	107.45	107.17
Sep-05	112.16	108.41	105.68	108.00	107.45
Dec-05	117.48	112.53	108.41	109.34	108.00
Mar-06	117.13	115.59	112.53	111.97	109.34
Jun-06	113.58	116.06	115.59	113.66	111.97
Sep-06	116.51	115.74	116.06	115.37	113.66
Dec-06	117.55	115.88	115.74	116.45	115.37
Mar-07	119.22	117.76	115.88	116.80	116.45
Jun-07	121.49	119.42	117.76	117.67	116.80
Sep-07	116.72	119.14	119.42	118.30	117.67
Dec-07	112.37	116.86	119.14	117.47	118.30

TABLE 3.1 (continued)

Date	Exchange Rate: Japanese Yen Against the US Dollar	Three Quarter Moving Average (MA3)	Forecast Based on MA3	Five Quarter Moving Average (MA5)	Forecast Based on MA5
Mar-08	103.71	110.93	116.86	114.70	117.47
Jun-08	105.04	107.04	110.93	111.87	114.70
Sep-08	107.10	105.28	107.04	108.99	111.87
Dec-08	94.41	102.18	105.28	104.53	108.99
Mar-09	95.05	98.85	102.18	101.06	104.53
Jun-09	96.57	95.34	98.85	99.63	101.06
Sep-09	92.56	94.73	95.34	97.14	99.63
Dec-09	89.94	93.03	94.73	93.71	97.14
Mar-10	90.86	91.12	93.03	93.00	93.71
Jun-10	91.38	90.73	91.12	92.26	93.00
Sep-10	84.66	88.97	90.73	89.88	92.26
Dec-10	82.06	86.03	88.97	87.78	89.88
Mar-11	82.18	82.97	86.03	86.23	87.78
Jun-11	81.11	81.79	82.97	84.28	86.23
Sep-11	76.88	80.06	81.79	81.38	84.28
Dec-11	78.20	78.73	80.06	80.09	81.38
Mar-12	79.53	78.20	78.73	79.58	80.09
Jun-12	79.43	79.05	78.20	79.03	79.58
Sep-12	78.05	79.00	79.05	78.42	79.03
Dec-12	82.74	80.08	79.00	79.59	78.42
Mar-13	92.38	84.39	80.08	82.43	79.59
Jun-13	99.16	91.43	84.39	86.35	82.43
Sep-13	97.99	96.51	91.43	90.07	86.35
Dec-13	101.90	99.68	96.51	94.83	90.07
Mar-14	102.29	100.73	99.68	98.74	94.83
Jun-14	101.76	101.98	100.73	100.62	98.74
Sep-14	105.32	103.12	101.98	101.85	100.62
Dec-14	116.00	107.69	103.12	105.45	101.85
Mar-15	119.10	113.47	107.69	108.89	105.45
Jun-15	121.59	118.90	113.47	112.75	108.89
Sep-15	121.69	120.80	118.90	116.74	112.75
Dec-15	121.48	121.59	120.80	119.97	116.74
Mar-16	115.12	119.43	121.59	119.80	119.97
Jun-16	107.08	114.56	119.43	117.39	119.80
Sep-16	102.60	108.27	114.56	113.60	117.39
Dec-16	111.41	107.03	108.27	111.54	113.60
Mar-17	**112.73**		107.03		111.54
			MAPE = 5.79%		**MAPE = 7.57%**

that we sum the first three observations (245.52, 227.27, and 219.17). This three-quarter total is then divided by 3 to obtain 230.65, which is the first number in the "Three-Quarter Moving Average (MA3)" column. This "smoothed" number, 230.65, becomes the forecast for 1980Q4, which is designated as Dec-80. Often when using quarterly data, we select one month in the quarter to represent the quarter. In this example, we are using the end month of each quarter to represent the quarter. Thus, March (Mar) represents the first quarter, June (Jun) represents the second quarter, September (Sep) represents the third quarter, and December (Dec) represents the fourth quarter.

The final value in the "Three-Quarter Moving Average MA3" column (107.03) is the forecast for 2017Q1, designated as Mar-17. It was calculated by summing the final three values in the "Actual" column and then dividing by 3 (321.10/3 = 107.03).

The five-quarter moving averages displayed in the same table are calculated in like manner: the first moving average of 221.89 is calculated by summing the first five actual values and dividing by 5:

$$\frac{245.52 + 227.27 + 219.17 + 209.79 + 207.73}{5} = \frac{1{,}109.47}{5} = 221.89$$

Thus, 221.89 becomes the forecast for the next period (Jun-81).

The five entries from Dec-15 through Dec-16 in the "Actual" column are averaged to give the final five-quarter moving average:

$$\frac{121.48 + 115.12 + 107.08 + 102.60 + 111.41}{5} = \frac{557.70}{5} = 111.54$$

This final moving average serves as the forecast for Mar-17.

The choice of the interval for the moving average depends on the length of the underlying cycle or pattern in the original data.

Obviously, three- and five-quarter moving averages are not the only kinds of moving averages. We could calculate seven- or nine-quarter moving averages if we wished or eight- or ten-quarter averages and so on. The choice of the interval for the moving average depends on the length of the underlying cycle or pattern in the original data. If we believe the actual data to be exhibiting a cycle that recurs every four periods, we would choose a four-period moving average in order to best dampen the short-run fluctuation. The simplest naive model of Chapter 1 used each period's actual value as the forecast for the next period; you could correctly think of this model as a one-period moving average, that is, a special case of the model we are examining here.

In order to compute whether the three-quarter or five-quarter moving average is the better forecasting model, it is useful to compute the mean absolute percentage error (MAPE) as we calculated it in Chapter 1. Table 3.1 shows the MAPE for both forecasts at the bottom of the table. The MAPE of 5.79 percent for the three-quarter moving average is less than the MAPE of 7.57 percent calculated for the five-quarter case, and so we conclude that the better forecast in this particular case is generated by the three-quarter model.

FIGURE 3.2 **Three-Quarter Moving-Average Forecast of the U.S. Exchange Rate with Japan** (c3t1)

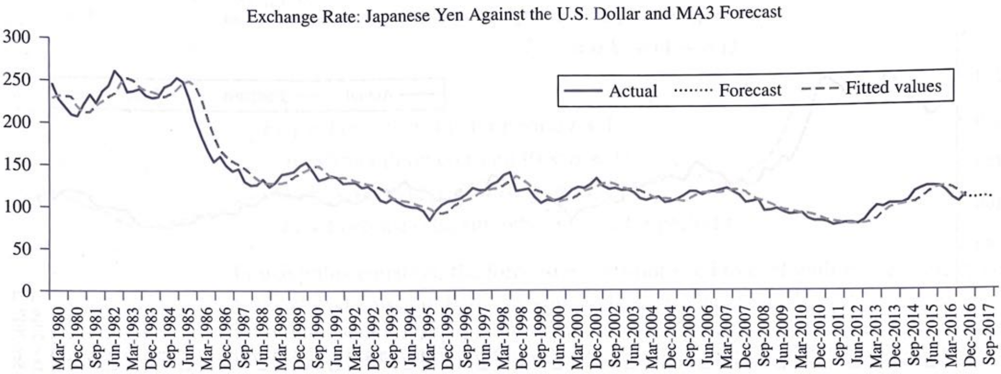

In preparing the forecasts for Mar-17, it was assumed that the actual value for that quarter was unknown. However, the actual value for that quarter is known in this situation and is shown in Table 3.1. Thus, we can see which of the two moving-average forecasts developed above was really the best for Mar-17. The forecast for the single quarter (Mar-17) shows that the five-quarter moving average was more accurate for this one quarter, even though for the entire time span the three-quarter moving average had the lower MAPE.

The three- and five-quarter moving averages are shown graphically in Figures 3.2 and 3.3, respectively. Notice in Figures 3.2 and 3.3 that the peaks and troughs of the actual series are different from those for either moving average. This failure of the moving averages to predict peaks and troughs is one of the shortcomings of moving-average models.

The moving-average forecasting method has fooled more than one forecaster by appearing to identify a cycle.

Moving averages are most likely to be successful in forecasting when the data are stationary. That is, when there is little, if any, positive or negative trend in the series.

One final and important observation: The moving-average forecasting method has fooled more than one forecaster by appearing to identify a cycle when, in fact, no cycle was present in the actual data. Such an occurrence can be understood if you think of an actual data series as being simply a series of random numbers. Since any moving average is serially correlated, because a number of contiguous periods have been averaged, *any* sequence of random numbers could appear to exhibit cyclical fluctuation.[1]

[1] This incorrect conclusion is sometimes called the *Slutsky-Yule* effect, named after Eugen Slutsky and G. Udny Yule, who first pointed out the possibility of making a mistake in this manner. See Eugen E. Slutsky, "The Summation of Random Causes as the Source of Cyclic Processes," *Econometrica* 5 (1937), pp. 105–46; and G. Udny Yule, "On a Method of Investigating Periodicities in Disturbed Series, with Special Reference to Wolfer's Sunspot Numbers," Royal Society of London, *Philosophical Transactions* (1927), pp. 267–98.

FIGURE 3.3 **Five-Quarter Moving-Average Forecast of the U.S. Exchange Rate with Japan** (c3t1)

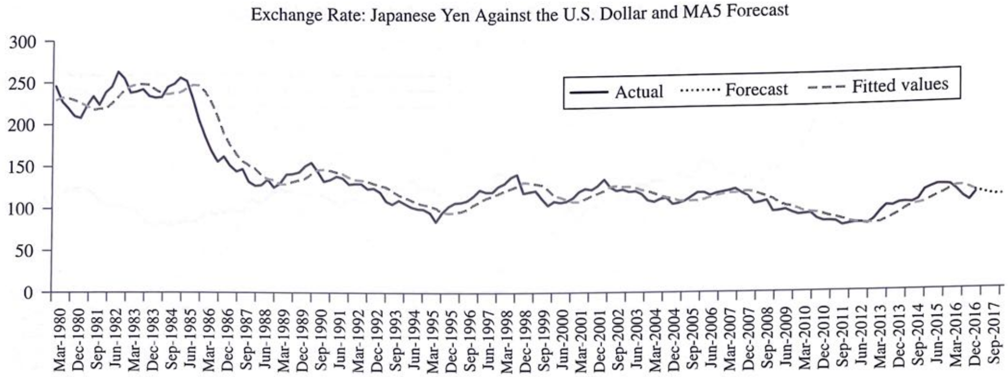

SIMPLE EXPONENTIAL SMOOTHING

With exponential smoothing, the forecast value at any time is a weighted average of all the available previous values.

Simple exponential smoothing, like moving averages, uses only past values of a time series to forecast future values of the same series and is properly employed when there is no trend or seasonality present in the data. With exponential smoothing, the forecast value at any time is a weighted average of all the available previous values; the weights decline exponentially as you go back in time.

Moving-average forecasting gives equal weights to the past values included in each average. However, exponential smoothing gives more weight to the recent observations and less to the older observations. The weights are made to decline exponentially with the age of the observation to conform to the argument that the most recent observations contain the most relevant information. As a result, more recent values are accorded proportionately more influence than older observations.

The number we choose for α is called the *level smoothing constant*.

Exponential smoothing proceeds by smoothing past values of the series but in a different manner than moving averages. The calculations for producing exponentially smoothed forecasts can be expressed as an equation. The weight of the most recent observation is assigned by multiplying the observed value by α, the next most recent observation by $(1 - \alpha)\alpha$, the next observation by $(1 - \alpha)^2\alpha$, and so on. The number we choose for α is called the *level smoothing constant*.[2]

[2] Our notation throughout the chapter for exponential smoothing follows approximately the notation found in Everette S. Gardner, "Exponential Smoothing: The State of the Art," *Journal of Forecasting* 4, no. 1 (1985), pp. 1–28. This article contains a very complete description of different forms of smoothing that are in common use and explains (with advanced mathematics) that there may be theoretical advantages for employing smoothing in situations where it can be shown that certain assumptions concerning the probability distribution of the series are met.

FIGURE 3.3 **Five-Quarter Moving-Average Forecast of the U.S. Exchange Rate with Japan** (c3t1)

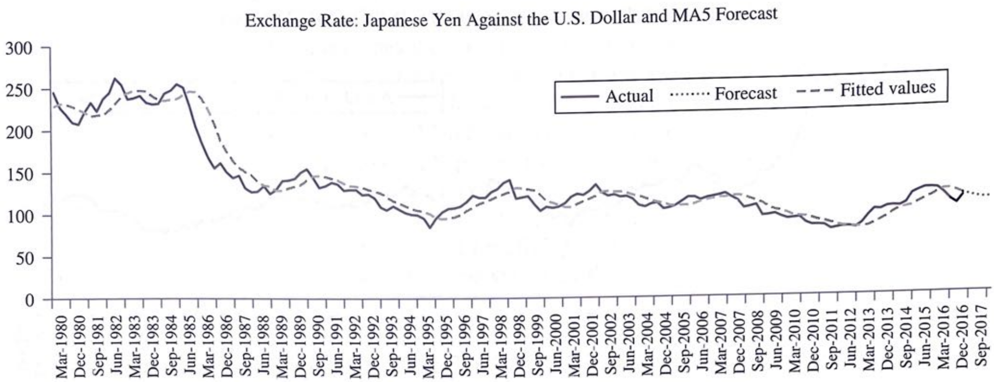

SIMPLE EXPONENTIAL SMOOTHING

With exponential smoothing, the forecast value at any time is a weighted average of all the available previous values.

Simple exponential smoothing, like moving averages, uses only past values of a time series to forecast future values of the same series and is properly employed when there is no trend or seasonality present in the data. With exponential smoothing, the forecast value at any time is a weighted average of all the available previous values; the weights decline exponentially as you go back in time.

Moving-average forecasting gives equal weights to the past values included in each average. However, exponential smoothing gives more weight to the recent observations and less to the older observations. The weights are made to decline exponentially with the age of the observation to conform to the argument that the most recent observations contain the most relevant information. As a result, more recent values are accorded proportionately more influence than older observations.

The number we choose for α is called the *level smoothing constant*.

Exponential smoothing proceeds by smoothing past values of the series but in a different manner than moving averages. The calculations for producing exponentially smoothed forecasts can be expressed as an equation. The weight of the most recent observation is assigned by multiplying the observed value by α, the next most recent observation by $(1 - \alpha)\alpha$, the next observation by $(1 - \alpha)^2\alpha$, and so on. The number we choose for α is called the *level smoothing constant*.[2]

[2] Our notation throughout the chapter for exponential smoothing follows approximately the notation found in Everette S. Gardner, "Exponential Smoothing: The State of the Art," *Journal of Forecasting* 4, no. 1 (1985), pp. 1–28. This article contains a very complete description of different forms of smoothing that are in common use and explains (with advanced mathematics) that there may be theoretical advantages for employing smoothing in situations where it can be shown that certain assumptions concerning the probability distribution of the series are met.

The simple exponential smoothing model can be written in the following manner:

$$F_{t+1} = \alpha X_t + (1 - \alpha) F_t \tag{3.1}$$

where:

F_{t+1} = Forecast value for period $t + 1$
α = Smoothing constant ($0 < \alpha < 1$)
X_t = Actual value now (in period t)
F_t = Forecast (i.e., smoothed) value for period t

In using this equation, the forecaster does not need to deal with every actual past value at every step; only the exponentially smoothed value for the last period and the actual value for this period are necessary. However, you will see that in practice all previous actual and forecast values play a role in the simple exponential smoothing forecast.

An alternative way of writing Equation 3.1 results from rearranging the terms as follows:

$$\begin{aligned} F_{t+1} &= \alpha X_t + (1 - \alpha) F_t \\ &= \alpha X_t + F_t - \alpha F_t \\ &= F_t + \alpha(X_t - F_t) \end{aligned} \tag{3.2}$$

From this form we can see that the exponential smoothing model "learns" from past errors. The forecast value at period $t + 1$ is increased if the actual value for period t is greater than it was forecast to be, and it is decreased if X_t is less than F_t.

We can show how all historical observations are included in a simple exponential smoothing forecast as follows:

$$F_{t+1} = \alpha X_t + (1 - \alpha) F_t \tag{3.3}$$

and $\quad F_t = \alpha X_{t-1} + (1 - \alpha) F_{t-1}$

therefore $\quad F_{t+1} = \alpha X_t + (1 - \alpha)\alpha X_{t-1} + (1 - \alpha)^2 F_{t-1}$

and $\quad F_{t-1} = \alpha X_{t-2} + (1 - \alpha) F_{t-2}$

thus, $\quad F_{t+1} = \alpha X_t + (1 - \alpha)\alpha X_{t-1} + (1 - \alpha)^2 \alpha X_{t-2} + (1 - \alpha)^3 F_{t-2}$

We could continue this expansion to include X terms as far back as we have data, but this is probably far enough to help you see how the weights for previous time periods become smaller and smaller at a rate that depends on the value of α, as will be shown in the following tables for two alternative values of α.

The value of the level smoothing constant α is constrained to be in the range of zero to one. If a value close to 1 is chosen, recent values of the time series are weighted heavily relative to those of the distant past when the smoothed values are calculated. Likewise, if the value of α is chosen close to 0, then the values of the time series in the distant past are given weights comparable to those given the recent values. The rate at which the weights decrease can be seen from their values for an α of 0.1:

Time	$\alpha = 0.1$ Calculation	Weight
t		0.1
$t-1$	0.9×0.1	0.090
$t-2$	$0.9 \times 0.9 \times 0.1$	0.081
$t-3$	$0.9 \times 0.9 \times 0.9 \times 0.1$	0.073
⋮		⋮
Total		1.000

Regardless of the smoothing constant chosen, the weights will eventually sum to 1. Whether the sum of the weights converges on 1 quickly or slowly depends on the smoothing constant chosen. If, for example, we choose a smoothing constant of 0.9, the sum of the weights will approach 1 much more rapidly than when the level smoothing constant is 0.1:

Time	$\alpha = 0.9$ Calculation	Weight
t		0.9
$t-1$	0.1×0.9	0.09
$t-2$	$0.1 \times 0.1 \times 0.9$	0.009
$t-3$	$0.1 \times 0.1 \times 0.1 \times 0.9$	0.0009
⋮		⋮
Total		1.000

In practice, relatively small values of alpha (α) generally work best when simple exponential smoothing is the most appropriate model.

As a guide in choosing α, select values close to 0 if the series has a great deal of random variation; select values close to 1 if you wish the forecast values to depend strongly on recent changes in the actual values. The mean absolute percentage error (MAPE) is often used as the criterion for assigning an appropriate smoothing constant; the smoothing constant giving the smallest MAPE would be selected as the model likely to produce the smallest error in generating additional forecasts. In practice, relatively small values of alpha (α) generally work best when simple exponential smoothing is the most appropriate model.

The following example will demonstrate the technique. Suppose we wish to forecast the University of Michigan Index of Consumer Sentiment for January 2017 based on monthly data from January 2016 through December 2016. These values are shown in the "X_t" column of Tables 3.2A and 3.2B for January 2016 through December 2016. Since no previous forecast is available for the first period (January 2016), we have to select a value. This value is called a seed value because it allows the process to begin, just as a plant seed allows the plant to begin growth. In the A part of Table 3.2, we use a seed value equal to that month's actual index value of 92. That is, we are assuming that the forecast for January 2016 was perfect. This process of choosing an initial value for the smoothed (forecast) series is called *initializing* the model, or *warming up* the model.[3] All the other values in the "Forecast" column were calculated by using Equation 3.1

TABLE 3.2A Simple Exponential Smoothing Forecast of the University of Michigan Index of Consumer Sentiment (c3t2A&B)

Date	UMICS = X_t	Forecast = $F_t = \alpha X_{t-1} + (1-\alpha)F_{t-1}$	Error ($X - F$)	Absolute Error	Absolute % Error (Absolute Error/X)*100
Jan-2016	92	92.00	0.00	0.00	0.00
Feb-2016	91.7	92.00	−0.30	0.30	0.33
Mar-2016	91	91.88	−0.88	0.88	0.97
Apr-2016	89	91.53	−2.53	2.53	2.84
May-2016	94.7	90.52	4.18	4.18	4.42
Jun-2016	93.5	92.19	1.31	1.31	1.40
Jul-2016	90	92.71	−2.71	2.71	3.02
Aug-2016	89.8	91.63	−1.83	1.83	2.04
Sep-2016	91.2	90.90	0.30	0.30	0.33
Oct-2016	87.2	91.02	−3.82	3.82	4.38
Nov-2016	93.8	89.49	4.31	4.31	4.59
Dec-2016	98.2	91.21	6.99	6.99	7.11
Jan-2017	NA	94.01			
				MAPE =	2.86

Table 3.2A. Simple Exponential Smoothing Forecast of the University of Michigan Index of Consumer Sentiment and calculation of the MAPE (c3t2). A seed value of 92 has been used and an alpha of 0.4 has been arbitrarily selected.

(continued on next page)

[3] The choice of a starting value in exponential smoothing models has been a matter of some discussion, with little empirical evidence favoring any particular approach. R. G. Brown first suggested using the mean of the data for the starting value, and this suggestion has been quite popular in actual practice. A linear regression (like that described in Chapter 4) is sometimes used when selecting starting values for seasonal factors, and time-series decomposition (as discussed in Chapter 6) has also been used. If the data include a trend, backcasting is sometimes used to select a starting value; but if the trend is erratic, this sometimes leads to negative starting values, which make little sense. A discussion of the various alternatives (including using the first value in the series or using the mean of the series, which are both popular in practice) appears in the Gardner article (footnote 2).

TABLE 3.2B

Date	UMICS = X_t	Forecast = F_t = $\alpha X_{t-1} + (1 - \alpha)F_{t-1}$	Error $(X - F)$	Absolute Error	Absolute % Error (Absolute Error/X)*100
Jan-2016	92	0.00	92.00	92.00	100.00
Feb-2016	91.7	36.80	54.90	54.90	59.87
Mar-2016	91	58.76	32.24	32.24	35.43
Apr-2016	89	71.66	17.34	17.34	19.49
May-2016	94.7	78.59	16.11	16.11	17.01
Jun-2016	93.5	85.04	8.46	8.46	9.05
Jul-2016	90	88.42	1.58	1.58	1.75
Aug-2016	89.8	89.05	0.75	0.75	0.83
Sep-2016	91.2	89.35	1.85	1.85	2.03
Oct-2016	87.2	90.09	−2.89	2.89	3.32
Nov-2016	93.8	88.93	4.87	4.87	5.19
Dec-2016	98.2	90.88	7.32	7.32	7.45
Jan-2017	NA	93.81			
				MAPE =	14.67

Table 3.2B. Here a seed value of 0 has been used and alpha has been kept at 0.4. Note how little the forecast for January 2017 has changed.

with a level smoothing constant (α) of 0.4. The actual and forecast values are shown in Figure 3.4.

In part B of Table 3.2, we have changed the seed value to zero. This is certainly a dramatic difference from the perfect forecast of 92 assumed in part A of that table. Notice that, even with a fairly short data series of 12 observations, the

FIGURE 3.4
A Simple Exponential Smoothing Forecast of the University of Michigan Index of Consumer Sentiment
(c3t2A&B)

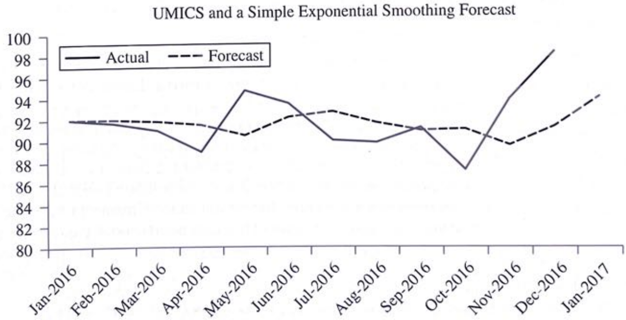

affect on the January 2017 forecast is quite small. The difference is only 0.20. The MAPE for the model in Table 3.2B is 14.67 percent compared to a MAPE of 2.86 percent for the model in Table 3.2A. You can see that this is strongly influenced by the very low early errors when the first forecast was assumed to be perfect compared to the large early errors when the first forecast was assumed to be zero.

Let us illustrate the calculation of the forecast value for February 2016 by using Equation 3.1 as follows:

$$F_{t+1} = \alpha X_t + (1 - \alpha) F_t$$
$$F_2 = \alpha X_1 + (1 - \alpha) F_1$$
$$F_2 = 0.4(92) + (1 - 0.4)(92) = 92$$

This smoothed value of 92 is the forecast for February 2016 ($t = 2$). Once actual data for February become available, the model is used to forecast March, and so on.

The error for February 2016 forecast (rounded) is calculated as:

$$e_2 = X_2 - F_2 = 91.7 - 92 = -0.30$$

The predominant reason for using simple smoothing is that it requires a limited quantity of data and it is simpler than most other forecasting methods. Its limitations, however, are that its forecasts lag behind the actual data, and it has no ability to adjust for any trend or seasonality in the data. Like moving averages, simple exponential smoothing is best used when the data have little or no positive or negative trend. That is, simple exponential smoothing is best applied for data that are stationary. In the following section, you will see that a modification to the simple exponential smoothing model allows for the use of data that does have a significant trend.

HOLT'S EXPONENTIAL SMOOTHING

Holt's two-parameter exponential smoothing method (called "Double Exponential Smoothing Holt" in ForecastX™) is an extension of simple exponential smoothing; it adds a growth factor (or trend factor) to the smoothing equation as a way of adjusting for the trend.

Two further extensions of the smoothing model can be used in order to bring the forecast values closer to the values observed if the data series exhibits a trend and/or seasonality (the first extension is discussed in this section, and the second in the following section). In real-world situations, one or both of these techniques are often used because real-world data often are not stationary and often include a seasonal pattern.

The first extension is to adjust the smoothing model for a trend in the data; with a trend in the data, the simple smoothing model will have large errors that often move from positive to negative or vice versa. When a trend exists, the forecast may be improved by adjusting for the trend by using a form of smoothing named after its originator, C. C. Holt. Holt's two-parameter exponential smoothing method is an extension of simple exponential smoothing; it adds a growth factor (or trend factor) to the smoothing equation as a way of adjusting for the trend. Three equations and two smoothing constants are used in the model.

$$F_{t+1} = \alpha X_t + (1 - \alpha)(F_t + T_t) \tag{3.4}$$

$$T_{t+1} = \gamma(F_{t+1} - F_t) + (1 - \gamma)T_t \tag{3.5}$$

$$H_{t+m} = F_{t+1} + mT_{t+1} \tag{3.6}$$

where:

F_{t+1} = Smoothed value for period $t + 1$

α = Smoothing constant for the level ($0 < \alpha < 1$)

X_t = Actual value now (in period t)

F_t = Forecast (i.e., smoothed) value for time period t

T_{t+1} = Trend estimate

γ = Smoothing constant for the trend estimate ($0 < \gamma < 1$)

m = Number of periods ahead to be forecast

H_{t+m} = Holt's forecast value for period $t + m$

Equation 3.4 adjusts F_{t+1} for the growth of the previous period, T_t, by adding T_t to the smoothed value of the previous period, F_t. The trend estimate is calculated in Equation 3.5, where the difference of the last two smoothed values is calculated. Because these two values have already been smoothed, the difference between them is assumed to be an estimate of trend in the data. The second smoothing constant, γ in Equation 3.5, is arrived at by using the same principle employed in simple exponential smoothing. The most recent trend ($F_{t+1} - F_t$), is weighted by γ, and the last previous smoothed trend, T_t, is weighted by ($1 - \gamma$). The sum of the weighted values is the new smoothed trend value T_{t+1}.

Equation 3.6 is used to forecast m periods into the future by adding the product of the trend component, T_{t+1}, times the number of periods to forecast, m, to the current value of the smoothed data F_{t+1}.

Holt's method accounts for any linear trend in the data.[4] As an example, consider personal consumption expenditures (PCE) in the United States. A graph of

[4] All trends, of course, do not have to be linear, and there are smoothing models that can account for nonlinear trends. In this chapter, we are examining only a subset of the number of possible smoothing models. For a listing of smoothing models, see Carl C. Pegels, "Exponential Forecasting: Some New Variations," *Management Science* 15, no. 12 (1969), pp. 311–15, or the Gardner article (1985). Both of these articles cover many smoothing models, including some that are very rarely used in actual practice.

FIGURE 3.5 **Personal Consumption Expeditures in the United States. Quarterly Data in Millions of Constant 2009 Dollars** (c3t3)

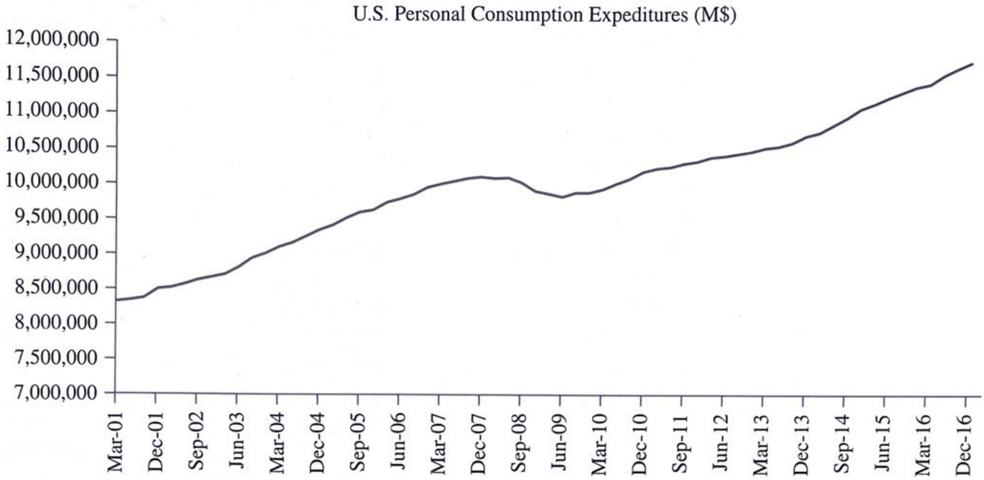

PCE (in millions of constant dollars) is shown in Figure 3.5. In this graph, it is clear that over time PCE have been rising. That is, the data have a clear positive trend (they are not stationary). The downturn in the U.S. economy is evident in 2008, when PCE dropped for a short period. Note that in Figure 3.5, the vertical axis has been adjusted to start at 7,000,000 million rather than at zero. This has been done to better illustrate the drop during 2008.

Table 3.3 illustrates the application of Holt's model to U.S. personal consumption expenditures, using the ForecastX™ software. The two smoothing constants are $\alpha = 0.88$ and $\gamma = 0.68$. Two starting values are needed: one for the first smoothed value and another for the first trend value. The initial smoothed value is often a recent actual value available; the initial trend value is often 0.00 if no past data are available (see footnote 3). The following naming conventions for model constants are used by ForecastX™ for all smoothing models (simple, Holt's, and Winters'):

ForecastX™ Naming Conventions for Smoothing Constants

Alpha (α)	=	the level smoothing constant
Gamma (γ)	=	the trend smoothing constant
Beta (β)	=	the seasonal smoothing constant (Winters' only)

For the personal consumption data, Equations 3.4 through 3.6 can be used to calculate the Holt's forecast for June 2001. To do so, we will arbitrarily select the first actual value as our initial smoothed value ($F_1 = 8{,}319{,}403$) and 53,000 as our

Date	U.S. Personal Consumption Expenditures (M$)	Fitted and Forecast Values	Error (Actual - Predicted)
Mar-01	8,319,403	8,347,587	−28,184
Jun-01	8,340,761	8,339,189	1,572
Sep-01	8,371,247	8,357,981	13,266
Dec-01	8,499,132	8,395,073	104,059
Mar-02	8,524,579	8,574,887	−50,307
Jun-02	8,568,126	8,588,268	−20,142
Sep-02	8,628,043	8,616,129	11,915
Dec-02	8,674,353	8,679,435	−5,082
Mar-03	8,712,527	8,724,691	−12,165
Jun-03	8,809,507	8,756,378	53,129
Sep-03	8,939,387	8,877,635	61,751
Dec-03	9,008,814	9,043,673	−34,858
Mar-04	9,096,415	9,103,518	−7,103
Jun-04	9,155,468	9,183,567	−28,099
Sep-04	9,243,001	9,228,182	14,820
Dec-04	9,337,837	9,319,571	18,266
Mar-05	9,409,222	9,424,996	−15,774
Jun-05	9,511,451	9,490,903	20,548
Sep-05	9,585,233	9,601,215	−15,982
Dec-05	9,621,339	9,669,688	−48,349
Mar-06	9,729,225	9,680,507	48,718
Jun-06	9,781,025	9,806,265	−25,240
Sep-06	9,838,106	9,851,600	−13,493
Dec-06	9,938,409	9,899,171	39,237
Mar-07	9,990,656	10,016,871	−26,215
Jun-07	10,024,604	10,061,063	−36,459
Sep-07	10,069,158	10,074,270	−5,113
Dec-07	10,081,798	10,112,047	−30,249
Mar-08	10,060,966	10,109,444	−48,477
Jun-08	10,077,941	10,061,580	16,361
Sep-08	10,005,097	10,080,750	−75,654
Dec-08	9,884,724	9,973,223	−88,499
Mar-09	9,850,832	9,801,088	49,744
Jun-09	9,806,377	9,780,826	25,552
Sep-09	9,865,864	9,754,609	111,255
Dec-09	9,864,805	9,870,956	−6,152
Mar-10	9,917,689	9,880,052	37,637
Jun-10	9,998,389	9,950,425	47,964
Sep-10	10,063,083	10,058,782	4,301
Dec-10	10,166,127	10,131,219	34,909
Mar-11	10,217,123	10,251,666	−34,543
Jun-11	10,237,676	10,290,064	−52,387
Sep-11	10,282,234	10,281,185	1,049

TABLE 3.3 (continued)

Date	U.S. Personal Consumption Expenditures (M$)	Fitted and Forecast Values	Error (Actual - Predicted)
Dec-11	10,316,776	10,320,067	−3,291
Mar-12	10,379,022	10,353,141	25,881
Jun-12	10,396,630	10,427,523	−30,893
Sep-12	10,424,119	10,433,235	−9,117
Dec-12	10,453,205	10,452,665	540
Mar-13	10,502,300	10,480,936	21,364
Jun-13	10,523,928	10,540,435	−16,507
Sep-13	10,573,135	10,556,596	16,539
Dec-13	10,662,222	10,611,859	50,362
Mar-14	10,712,811	10,727,272	−14,461
Jun-14	10,813,346	10,776,807	36,539
Sep-14	10,912,871	10,893,319	19,553
Dec-14	11,036,376	11,006,620	29,755
Mar-15	11,102,370	11,146,833	−44,463
Jun-15	11,181,347	11,194,821	−13,474
Sep-15	11,255,893	11,262,028	−6,135
Dec-15	11,319,286	11,332,015	−12,729
Mar-16	11,365,214	11,388,523	−23,309
Jun-16	11,484,859	11,421,667	63,192
Sep-16	11,569,014	11,569,145	−130
Dec-16	11,655,004	11,660,696	−5,692
Mar-17		11,743,915	
Jun-17		11,832,155	
Sep-17		11,920,394	
Dec-17		12,008,634	
Mar-18		12,096,873	
Jun-18		12,185,113	
Sep-18		12,273,352	
Dec-18		12,361,592	
	Alpha = 0.88		
	Gamma = 0.68		
	MAPE = 0.30%		

initial trend (T_1 = 53,000). The initial trend estimate of 53,000 was determined by calculating the average quarter to quarter change for the entire series. The smoothed value for period 2 (June 2001) is calculated by:

$$\begin{aligned} F_{t+1} &= \alpha X_t + (1-\alpha)(F_t + T_1) \\ F_2 &= 0.88(8,319,403) + (1-0.88)(8,319,403 + 53,000) \\ &= 7,321,075 + 1,004,688 \\ &= 8,325,763 \end{aligned}$$

The trend estimate for period 2 is calculated as:

$$\begin{aligned} T_{t+1} &= \gamma(F_{t+1} - F_t) + (1-\gamma)T_t \\ T_2 &= 0.68(8{,}325{,}763 - 8{,}319{,}403) + (1-0.68)(53{,}000) \\ &= 0.68(6{,}360) + (0.32)(53{,}000) \\ &= 4{,}325 + 16{,}960 \\ &= 21{,}285 \end{aligned}$$

The forecast for period 2 is calculated as:

$$\begin{aligned} H_{t+m} &= F_{t+1} + mT_{t+1} \\ H_2 &= F_2 + 1\,T_2 \\ &= 8{,}325{,}763 + (1)(21{,}285) \\ &= 8{,}347{,}048 \end{aligned}$$

Our calculated forecast for June 2001 differs from what you see in Table 3.3. This is because our arbitrary selection of seed values differs from those selected by ForecastX™. Over the course of many quarters, the effect of differing seed values would diminish to almost nothing, and if we continued the hand calculations, our final forecasts would be virtually identical to those in Table 3.3. In Table c3t2A&B, the affect of changing the seed value for a simple exponential smoothing problem was demonstrated to have little affect on the final forecast. The same result is true for all exponential smoothing models. In the example for simple exponential smoothing, we had only 12 observations, while in the current situation, we have 64 observations. The more observations one has, the less the affect of the seed value on the final forecast.

Figure 3.6 shows a plot of the actual values, the fitted values, and the forecast values generated by this model. For 2001 through 2016 (the period for which data exist), the calculated values are called the "fitted" data. For 2017 through 2018 (the forecast period), the values are called "forecast" values.

FIGURE 3.6 **Personal Consumption Expenditures and Holt's Forecast** In this case, alpha (the level constant) was 0.88 and gamma (the trend constant) was 0.68. (c3t3)

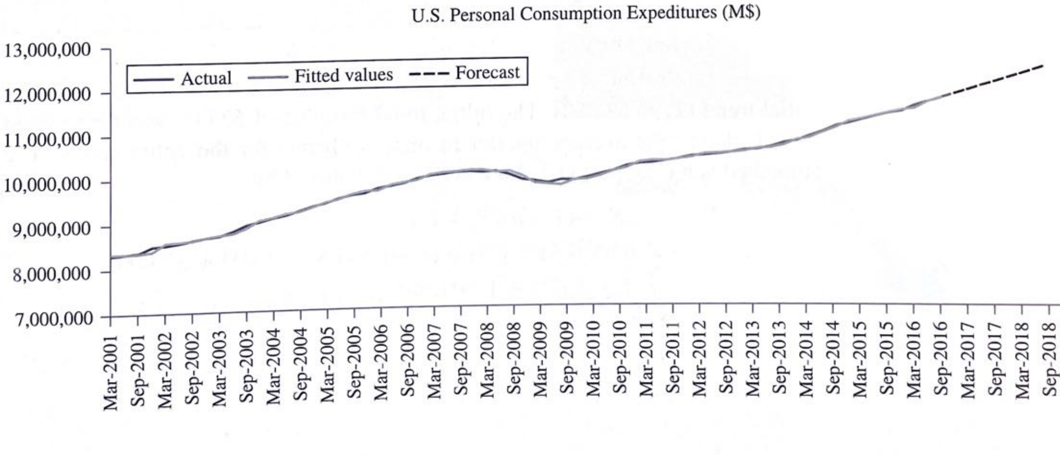

Some commercially available forecasting packages allow the forecaster to minimize the value of the MAPE (or some similar summary statistic) by automatically adjusting the smoothing constants (ForecastX™ automatically adjusts). This, of course, is preferable to making numerous adjustments by hand. In this example, the smoothing constants were determined by ForecastX™.

Holt's form of exponential smoothing is tedious to calculate by hand, or even using Excel. For Holt's exponential smoothing and other advanced smoothing methods, specialized forecasting software (such as ForecastX™) should be used. Holt's exponential smoothing is appropriate when the data show some linear trend but little or no seasonality. A descriptive name for Holt's smoothing might be *linear-trend smoothing.*

WINTERS' EXPONENTIAL SMOOTHING

Vinters' exponential moothing model is the econd extension of the asic smoothing model; t is used for data that xhibit both trend and easonality.

Winters' exponential smoothing model is the second extension of the basic smoothing model; it is used for data that exhibit both trend and seasonality. It is a three-parameter model that is an extension of Holt's model. Winters was a student of professor Holt and developed this modification as part of his graduate work. For this reason, the same model is referred to by two names: Winters' exponential smoothing and Holt-Winters' exponential smoothing.

In the Winters' model, an additional equation adjusts the model for the seasonal component. The four equations necessary for Winters' model are:

$$F_t = \alpha X_t / S_{t-p} + (1 - \alpha)(F_{t-1} + T_{t-1}) \quad \textbf{(3.7)}$$

$$S_t = \beta X_t / F_t + (1 - \beta) S_{t-p} \quad \textbf{(3.8)}$$

$$T_t = \gamma(F_t - F_{t-1}) + (1 - \gamma) T_{t-1} \quad \textbf{(3.9)}$$

$$W_{t+m} = (F_t + m T_t) S_{t+m-p} \quad \textbf{(3.10)}$$

where:

F_t = Smoothed value for period t

α = Smoothing constant for the level $(0 < \alpha < 1)$

X_t = Actual value now (in period t)

F_{t-1} = Average experience of series smoothed to period $t - 1$

T_{t+1} = Trend estimate

S_t = Seasonality estimate

β = Smoothing constant for seasonality estimate $(0 < \beta < 1)$

γ = Smoothing constant for trend estimate $(0 < \gamma < 1)$

m = Number of periods in the forecast lead period

p = Number of periods in the seasonal cycle

W_{t+m} = Winters' forecast for m periods into the future

Equation 3.7 updates the smoothed series for both trend and seasonality; note that the equation is only slightly different from Equation 3.4 in Holt's model. In Equation 3.7, X_t is divided by S_{t-p} to adjust for seasonality; this operation deseasonalizes the data or removes any seasonal effects left in the data. It is easy to see how this deseasonalizes the data if you consider what happens when S_{t-p} is greater than 1, as it would be when the value in period $t - p$ is greater than the average in its seasonality. Dividing X_t by S_{t-p} reduces the original value by a percentage equal to the percentage that the seasonality of the period was above the average. An opposite adjustment would take place if the period were below the average in terms of seasonality.

The seasonality estimate itself is smoothed in Equation 3.8, and the trend estimate is smoothed in Equation 3.9; each of these processes is exactly the same as in simple exponential smoothing. The final equation, 3.10, is used to compute the forecast for m periods into the future; the procedure is almost identical to that in Holt's model (Equation 3.6).

To illustrate Winters' exponential smoothing, we will use data for U.S. sales by new car dealers, in millions of dollars. The data are quarterly starting with the first quarter of 2001 and extending through the last quarter of 2016. The last month of each quarter is used to represent the quarter.

Sales by new car dealers are typically highest in the second and third quarters of the year. This would include the months of April, May, and June for the second quarter, then July, August, and September for the third quarter. As you can see by the dotted trend line in Figure 3.7., there has been an overall upward trend in the data since our 2001Q1 (Mar-01) starting point.

You have seen already how to apply the equations to do a few of the calculations for simple and Holt's exponential smoothing. We will not repeat that process for the Winters' model. Doing the calculations for a model like Winters' by hand

FIGURE 3.7 **U.S. Retail Sales: New Car Dealers: NAICS 44111: NSA: Millions of Dollars** (c3t4)

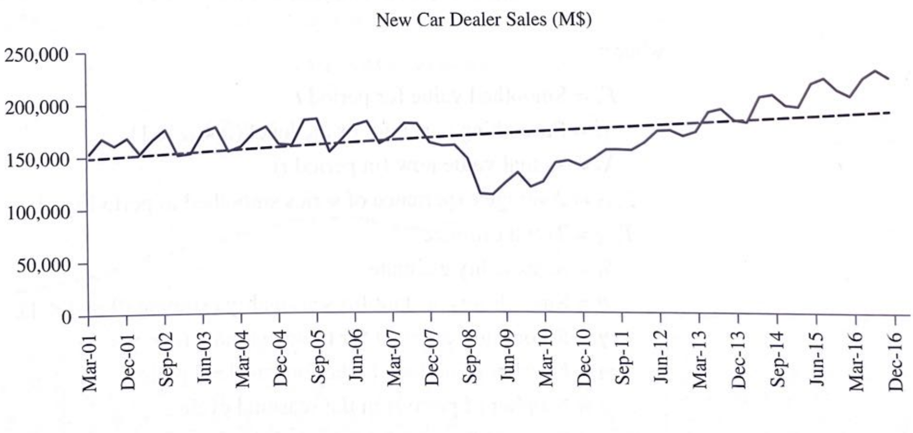

or in Excel is tedious. Professional forecasters use specialized forecasting software for such calculations.

Having ForecastX™ determine the parameters that would minimize the MAPE results in an alpha (level) constant of 0.86, a beta (seasonal) constant of 1.00, and a gamma (trend) constant of 0.06. The MAPE for the model is 3.65 percent.

As with simple and Holt's exponential smoothing, initial values must be selected to *initialize* or *warm up* the model. Over a long time period, such as in this example, the particular values selected have little effect on the forecast of sales by new car dealers for 2017 and 2018. These initial values are also determined within the software.

The results of the Winters' exponential smoothing forecast of sales by new car dealers are shown in Table 3.4 and in Figure 3.8. You can see, especially in the graph, that the model works quite well. The mean absolute percentage error (MAPE) of 3.65 percent suggests that one could have confidence in the Winters' forecast.

TABLE 3.4 **Winters' Three-Parameter Exponential Smoothing for New Car Dealer Sales** (c3t4)

Dates	Actual Sales (M$)	Fitted and Forecast Sales	Error	Dates	Actual Sales (M$)	Fitted and Forecast Sales	Error
Mar-2001	153,232	153,384	−152	Mar-2007	167,528	163,453	4,075
Jun-2001	167,504	169,134	−1,630	Jun-2007	180,006	182,930	−2,924
Sep-2001	160,865	169,026	−8,161	Sep-2007	179,216	182,976	−3,760
Dec-2001	167,812	146,934	20,878	Dec-2007	160,957	160,405	552
Mar-2002	152,641	164,119	−11,478	Mar-2008	158,360	164,023	−5,663
Jun-2002	166,036	170,007	−3,971	Jun-2008	158,742	172,743	−14,001
Sep-2002	175,863	166,830	9,033	Sep-2008	145,973	161,599	−15,626
Dec-2002	151,219	163,066	−11,847	Dec-2008	112,431	130,926	−18,495
Mar-2003	152,843	147,028	5,815	Mar-2009	111,787	113,834	−2,047
Jun-2003	173,360	168,732	4,628	Jun-2009	122,914	117,947	4,967
Sep-2003	178,086	175,191	2,895	Sep-2009	132,665	120,640	12,025
Dec-2003	154,916	163,076	−8,160	Dec-2009	119,052	114,509	4,543
Mar-2004	159,086	152,759	6,327	Mar-2010	124,075	120,746	3,329
Jun-2004	172,174	175,615	−3,441	Jun-2010	142,226	132,870	9,356
Sep-2004	178,077	174,724	3,353	Sep-2010	144,107	142,139	1,968
Dec-2004	161,216	161,390	−174	Dec-2010	138,852	125,933	12,919
Mar-2005	160,070	160,194	−124	Mar-2011	145,940	141,270	4,670
Jun-2005	183,594	176,212	7,382	Jun-2011	154,269	158,973	−4,704
Sep-2005	184,878	186,328	−1,450	Sep-2011	154,364	156,131	−1,767
Dec-2005	153,405	167,976	−14,571	Dec-2011	153,993	137,555	16,438
Mar-2006	164,864	153,908	10,956	Mar-2012	160,738	155,929	4,809
Jun-2006	178,363	180,943	−2,580	Jun-2012	172,010	174,567	−2,557
Sep-2006	181,615	180,771	844	Sep-2012	172,592	175,235	−2,643
Dec-2006	160,780	162,530	−1,750	Dec-2012	167,210	157,336	9,874

(continued on next page)

each Winters' model. For the new car dealer sales model, the seasonal indices were calculated as:

Seasonal Indices	Value
Index 1	0.95
Index 2	1.04
Index 3	1.05
Index 4	0.96

Since our data set began with the first quarter of the year 2001 (i.e., January, February, and March), index 1 above refers to this first quarter of the year as well. The remaining three quarters also match the calendar quarters. These indices may be easily interpreted as percentages. Index 1 (0.95) is interpreted as indicating that quarter one sales are usually about 5 percent below an "average quarter." On the other hand, the second quarter index of 1.04 indicates that typically second-quarter sales are 4 percent higher than an "average quarter." An average quarter is the result of adding all four quarters together and dividing by four.

With this interpretation in mind, it becomes easy to see that new car dealer sales are relatively high in quarters two and three. Some products and services will exhibit very strong seasonality, while others may be affected only to a minor degree. When working with business and economic data, it is usually a good assumption to expect the data to be seasonal. Computing a Winters' model for the data will help the researcher determine the magnitude of the seasonality and identify precisely when above-average and below-average occurrences take place.

ADAPTIVE–RESPONSE-RATE SINGLE EXPONENTIAL SMOOTHING

An interesting variant on simple smoothing called *adaptive–response-rate single exponential smoothing (ADRES)* has an important advantage over normal smoothing models because of the manner in which the smoothing constant is chosen. In ADRES smoothing, there is no requirement to actually choose an α value! This is an attractive feature if what you need is a very low-cost method of forecasting requiring no sophisticated knowledge of the technique. Real-world situations requiring the frequent forecasting of many items (perhaps thousands) would be ideal candidates for ADRES smoothing forecasts.

Adaptive-response smoothing does not use one single α value like the simple exponential smoothing model does.

Adaptive-response smoothing does not use one single α value like the simple exponential smoothing model does. The word *adaptive* in its name gives a clue to how the model works. The α value in the ADRES model is not just a single number, but rather *adapts* to the data. When there is a change in the basic pattern of the data, the α value adapts.

For instance, suppose that some data to be forecast fluctuate around a mean value of m. The best estimate of the next observation of the data might then be that mean value (m). But suppose further that after some time an outside force changes the mean value of m and the new value is now m'. The data then fluctuate around the new mean value of m'. If we had a way of adapting to the new mean of m', we could then use that adapted estimate as the forecast for future values of the data. In fact, we would like to be able to adapt each time the mean value of the data changed; sometimes we would adapt very often, if the mean changed frequently, and at other times we would adapt very rarely, if the data changed only infrequently.

Because of the simplicity of the ADRES smoothing model and its ability to adapt to changing circumstances, it is quite often used in actual practice. Keep in mind, however, that it is a variant of the simple smoothing model and so assumes that the data to be forecast have little trend or seasonality (or that the trend or seasonality in the data has been removed).

The ADRES model looks very much like the simple smoothing model presented earlier:

$$F_{t+1} = \alpha_t X_t + (1 - \alpha_t) F_t \quad \text{(ADRES equation)} \tag{3.11}$$

where:

$$\alpha_t = \left| \frac{S_t}{A_t} \right| \tag{3.12}$$

$$S_t = \beta e_t + (1 - \beta) S_{t-1} \quad \text{(Smoothed error)} \tag{3.13}$$

$$A_t = \beta |e_t| + (1 - \beta) A_{t-1} \quad \text{(Absolute smoothed error)} \tag{3.14}$$

$$e_t = X_t - F_t \quad \text{(Error)} \tag{3.15}$$

Note carefully the subscripts on the α term! There may now be a different α value for each period.

The ADRES equation is the same as the one for simple exponential smoothing with the exception of the manner in which the α value is chosen. In the simple exponential smoothing model, we chose the α value by selecting the value that minimized the mean absolute percentage error associated with the model. But in simple smoothing, we were allowed to choose only a single value for α. In the ADRES smoothing model, we may allow the α value to adapt as the data change.

The smoothing value (α) is now given as the absolute value of the smoothed error divided by the absolute smoothed error. The smoothed error is itself a smoothed value, with a smoothing factor of β. The absolute smoothed error is also a smoothed value, again using the smoothing constant β. In most cases, β is assigned a value of either 0.1 or 0.2. Thus, the first term of both the smoothed error and absolute smoothed error equations has a lighter weight than the second term.

To explain ADRES smoothing, consider Table 3.5, which lists 12 values of an observed data series. We would like to model the series using an adaptive–response-rate smoothing model. Note that the first six values of the series average about 100; the last six values in the series average about 125. This is a situation similar to that described in the preceding paragraphs and one conducive to the use of this technique. An adaptive–response-rate model should do quite well in modeling these data.

TABLE 3.5 Adaptive-Response Example (c3t5)

Period	Observed	Forecast	Error	Smoothed Error	Absolute Smoothed Error	α
1	100					
2	96	100.000	−4.00	−0.800	0.800	1.000
3	107	96.000	11.00	1.560	2.840	0.549
4	98	102.042	−4.04	0.440	3.080	0.143
5	103	101.464	1.53	0.659	2.771	0.238
6	99	101.830	−2.83	−0.039	2.783	0.014
7	126	101.790	24.21	4.811	7.068	0.681
8	128	118.267	9.73	5.795	7.601	0.762
9	122	125.687	−3.69	3.899	6.818	0.572
10	130	123.579	6.42	4.403	6.739	0.653
11	125	127.774	−2.77	2.968	5.946	0.499
12	124	126.390	−2.39	1.896	5.235	0.362

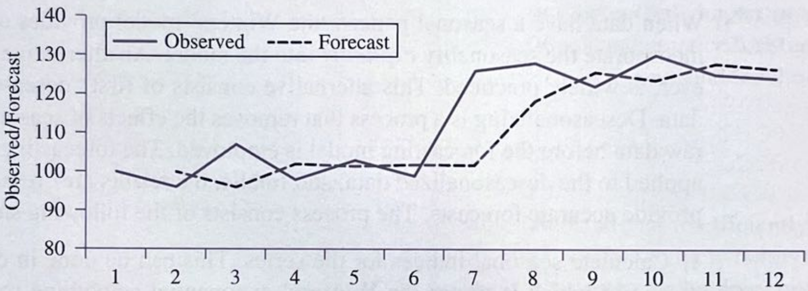

For period 5, the computations are as follows (with some rounding difference in the third decimal place):

$$\begin{aligned} F_5 &= \alpha_4 X_4 + (1-\alpha_4)F_4 \\ &= (0.143)(98) + (1-0.143)(102.042) \\ &= 14.014 + 87.450 \\ &= 101.464 \end{aligned}$$

Once the observed value of 103 becomes available for period 5, it is possible to make the following computations (assuming $\beta = .2$):

$$\begin{aligned} e_5 &= 103 - 101.464 = 1.536 \\ S_5 &= (0.2)(1.536) + (1-0.2)(0.440) = 0.659 \\ A_5 &= (0.2)(|1.536|) + (1-0.2)(3.080) = 2.771 \end{aligned}$$

and finally

$$\alpha_5 = \left|\frac{0.659}{2.771}\right| = 0.238$$

The process continues iteratively for all the remaining values in the example. In ForecastX™, you will get somewhat different results due to its use of a somewhat different algorithm.

Perhaps the most important consideration in adaptive–response-rate single exponential smoothing is the selection of the appropriate β factor. The β factor is usually set near 0.1 or 0.2 because these values reduce the effects of previous errors (i.e., they allow adaptation) but the values are small enough that the adaptation takes place gradually.

The ADRES model has no explicit way to handle seasonality. There are ways of using the ADRES model, however, with seasonal data. In fact, simple smoothing, Holt's smoothing, and the ADRES smoothing model may all be used with seasonal data. An example follows in the next section.

USING SINGLE, HOLT'S, OR ADRES SMOOTHING TO FORECAST A SEASONAL DATA SERIES

When data have a seasonal pattern, the Winters' model provides an easy way to incorporate the seasonality *explicitly* into the model. An alternative method, however, is widely practiced. This alternative consists of first "deseasonalizing" the data. Deseasonalizing is a process that removes the effects of seasonality from the raw data before the forecasting model is employed. The forecasting model is then applied to the deseasonalized data, and finally, the results are "reseasonalized" to provide accurate forecasts. The process consists of the following steps:

1. Calculate seasonal indices for the series. This can be done in different ways, one of which is to use the **Winters'** exponential smoothing routine in ForecastX™. As you have seen, part of the output from Winters' (Holt Winters' in ForecastX™) includes seasonal indices.
2. **Deseasonalize** the original data by **dividing** each value by its corresponding seasonal index.
3. Apply a forecasting method (such as simple, Holt's, or adaptive-response exponential smoothing) to the deseasonalized series to produce an intermediate forecast of the deseasonalized data.
4. Reseasonalize the series by **multiplying** each deseasonalized forecast by its corresponding seasonal index.

Many forecasters have found this method more accurate than using Winters' smoothing to incorporate seasonality.

Many forecasters have found this method more accurate than using Winters' smoothing to incorporate seasonality. This method is more flexible than the Winters' method alone because it allows for the use of simple smoothing in situations without any trend whatsoever while allowing Holt's smoothing to be used if a trend is present. (Recall that Winters' model assumes that a trend is present.) Further, the ADRES model could be used in situations where some adaptation of the α factor is desirable.

To illustrate this approach to forecasting a seasonal series, let us use data on new houses sold in the United States between January 2011 and December 2016. Table 3.6 shows a complete Excel sheet to illustrate the process. Note

Date	Total New Houses Sold in the United States: Thousands: NSA (NHS)	Seasonally Adjusted or Deseasonalized NHS (DSNHS	Holt's Forecast of DSNHS	Reseasonalized Forecast for NHS	SI
Jan-11	21	24	24	21.10	0.88
Feb-11	22	23	24	23.28	0.96
Mar-11	28	25	24	25.99	1.10
Apr-11	30	26	25	28.49	1.15
May-11	28	25	26	29.10	1.13
Jun-11	28	25	25	28.35	1.12
Jul-11	27	26	25	26.58	1.05
Aug-11	25	25	26	25.33	0.98
Sep-11	24	26	26	24.24	0.94
Oct-11	25	26	26	24.88	0.96
Nov-11	23	27	26	22.67	0.87
Dec-11	24	28	27	23.02	0.87
Jan-12	23	26	27	24.08	0.88
Feb-12	30	31	27	25.98	0.96
Mar-12	34	31	29	32.36	1.10
Apr-12	34	30	30	34.98	1.15
May-12	35	31	30	34.21	1.13
Jun-12	34	30	31	34.43	1.12
Jul-12	33	32	31	32.27	1.05
Aug-12	31	32	31	30.86	0.98
Sep-12	30	32	32	29.81	0.94
Oct-12	29	30	32	30.86	0.96
Nov-12	28	32	31	27.13	0.87
Dec-12	28	32	32	27.80	0.87
Jan-13	32	36	32	28.53	0.88
Feb-13	36	37	35	33.60	0.96
Mar-13	41	37	36	40.13	1.10
Apr-13	43	37	37	42.71	1.15
May-13	40	35	38	42.56	1.13
Jun-13	43	39	37	40.93	1.12
Jul-13	33	32	38	39.67	1.05
Aug-13	31	32	35	34.08	0.98
Sep-13	31	33	33	31.23	0.94
Oct-13	36	37	33	32.06	0.96
Nov-13	32	37	36	31.03	0.87
Dec-13	31	36	37	31.77	0.87
Jan-14	33	38	36	32.05	0.88
Feb-14	35	36	37	35.97	0.96

Date	Total New Houses Sold in the United States: Thousands: NSA (NHS)	Seasonally Adjusted or Deseasonalized NHS (DSNHS	Holt's Forecast of DSNHS	Reseasonalized Forecast for NHS	SI
Mar-14	39	35	37	40.74	1.10
Apr-14	39	34	36	41.83	1.15
May-14	43	38	35	39.97	1.13
Jun-14	38	34	37	41.29	1.12
Jul-14	35	33	36	37.25	1.05
Aug-14	36	37	35	34.00	0.98
Sep-14	37	39	36	33.75	0.94
Oct-14	38	40	38	36.57	0.96
Nov-14	31	36	39	33.89	0.87
Dec-14	35	40	38	32.54	0.87
Jan-15	39	44	39	34.58	0.88
Feb-15	45	47	42	40.82	0.96
Mar-15	46	42	45	49.50	1.10
Apr-15	48	42	44	50.05	1.15
May-15	47	42	43	48.54	1.13
Jun-15	44	39	42	47.37	1.12
Jul-15	43	41	41	42.93	1.05
Aug-15	41	42	41	40.54	0.98
Sep-15	35	37	42	39.24	0.94
Oct-15	39	41	40	38.06	0.96
Nov-15	36	42	40	34.98	0.87
Dec-15	38	44	41	35.74	0.87
Jan-16	39	44	43	37.72	0.88
Feb-16	45	47	44	42.40	0.96
Mar-16	50	45	46	50.30	1.10
Apr-16	55	48	46	52.67	1.15
May-16	53	47	47	53.45	1.13
Jun-16	50	45	47	52.82	1.12
Jul-16	54	52	46	48.34	1.05
Aug-16	46	47	49	48.48	0.98
Sep-16	44	47	48	45.37	0.94
Oct-16	46	48	48	45.98	0.96
Nov-16	40	46	48	41.73	0.87
Dec-16	39	45	47	41.04	0.87
Jan-17			46.38	40.77	0.88
Feb-17			46.63	44.98	0.96
Mar-17			46.89	51.62	1.10
Apr-17			47.15	54.19	1.15
May-17			47.40	53.67	1.13
Jun-17			47.66	53.23	1.12

FIGURE 3.9
New House Sales and the Deseasonalized New House Sales

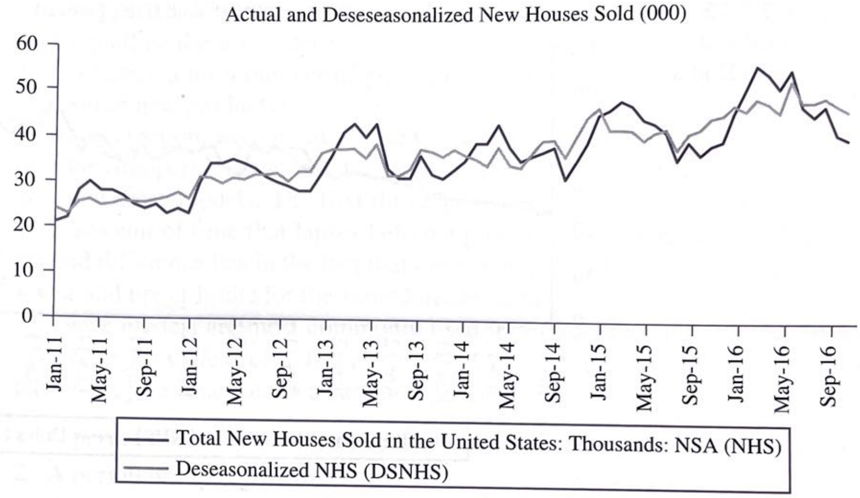

that in this example we are using monthly data. The first column contains the dates, as is always the case. The second column has the raw data for new houses sold (NHS). The third column has the data for NHS after the original data has been deseasonalized (DSNHS). To deseasonalize the NHS data, each value was divided by the seasonal index (SI) for the corresponding month. The seasonal indices are shown in the last column. Look at the seasonal indices to note that they repeat year after year; that is, all the January indices are the same. All the February indices are the same, and so on for all months. There are two good ways to get the seasonal indices. One is from the results of a Winters' exponential smoothing. Another way will be discussed in Chapter 6, which deals with time series decomposition. The raw NHS data and the deseasonalized values (DSNHS) are shown in Figure 3.9.

To forecast the deseasonalized data (DSNHS), we used a Holt's exponential smoothing model. The resulting forecast of DSNHS is shown in the fourth column of Table 3.6 and in Figure 3.10.

The seasonality was put back in the data by multiplying each value in the fourth column by the corresponding seasonal index from the last column. The results are in the fifth column and represent the forecast for new houses sold. The original values for new houses sold and the forecast are shown in the second and fifth of Table 3.6, as well as in Figure 3.11.

NEW-PRODUCT FORECASTING (GROWTH CURVE FITTING)

For new products, because they typically lack historical data, most forecasting techniques cannot produce satisfying results. For example, it is typically impossible for Holt's exponential smoothing to determine the trend since the data set

FIGURE 3.10
Deseasonalized NHS and a Holt's Forecast

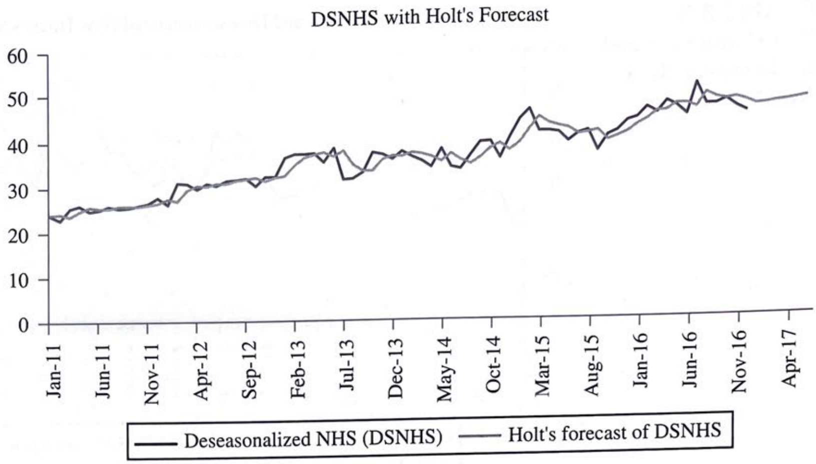

FIGURE 3.11
Actual New Houses Sold and the Forecast of NHS

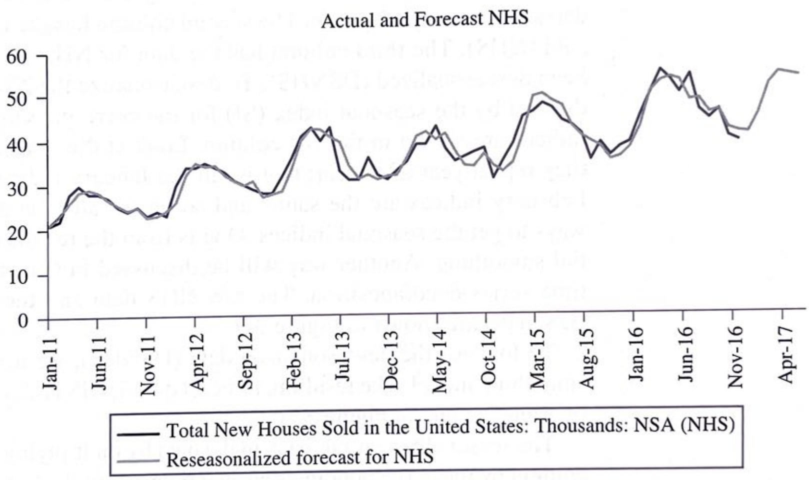

is too small. Alternatively, it may only predict a strong trend despite the fact that the new product has a growth limitation. To overcome this difficulty, forecasters use a number of models that generally fall in the category called *diffusion models* (probably because they described the manner in which technological innovations and new products "diffused" through an industry). These models are alternatively called S-curves, growth models, saturation models, or substitution curves. We have already seen one of these diffusion models in Chapter 1: the Bass model. An understanding of how to correctly use these models in the forecasting process can make them important tools for managerial decisions. These models as a group allow the forecaster to model the characteristic patterns that economists have identified for a number of processes (most importantly including the introduction of new products).

In this section, we present two new product models. The two diffusion models are the Gompertz curve and the logistic curve. There are two main differences between these models. The first difference is in the shapes of the product curve (i.e., amount of time that lapses before a product's growth curve stabilizes). The second difference lies in the fact that these new-product models may use different lower and upper limits for the same forecast data.

These models are most commonly used to forecast the sales of new products and technology life cycles. Just as new products have life cycles, technologies also have life cycles that follow a common pattern:

1. A period of slow growth just after introduction during an embryonic stage.
2. A period of rapid growth.
3. Slowing growth in a mature phase.
4. Decline.

The forecaster's task is to identify and estimate the parameters of such a pattern of growth using the same set of diagnostic statistics we have already learned to use with smoothing models in general.

Each new-product model has its own lower and upper limit. Expert opinion is needed to determine the correct upper and lower limits on the growth curves. In most instances, the lower limitation is 0 (e.g., sales cannot be below zero). Determining the upper limit is a more complicated undertaking. Regardless of the complication, diffusion models provide an important technique to use for forecasting when new products or technologies will replace existing products or technologies.

A significant benefit of using diffusion models in new-product forecasting is to identify and predict the timing of the four phases of the life cycle. In the late 1990s, the U.S. government decided to adopt a national standard for high-definition television (HDTV) and set a timetable for the changeover from analog to HDTV. The original plan called for broadcasters to begin broadcasting digital signals by 2002 and to turn off their analog transmitters altogether in 2006. This was a very ambitious plan and assumed that the adoption of HDTV by consumers would take place very quickly. Realizing that the elimination of analog transmissions would cause hardship if it occurred too early, another provision set forth by the Federal Communications Commission was that a market needed 85 percent penetration by HDTV before the analog signal could be eliminated.

Being able to forecast the growth and maturity of HDTV technology would allow broadcasters the opportunity to see if the 2006 "drop dead" date for analog television was reasonable. If 85 percent penetration was not reasonably achieved by this date, then broadcasters would be in the unenviable position of having to

FIGURE 3.12
Characteristic S-Curve

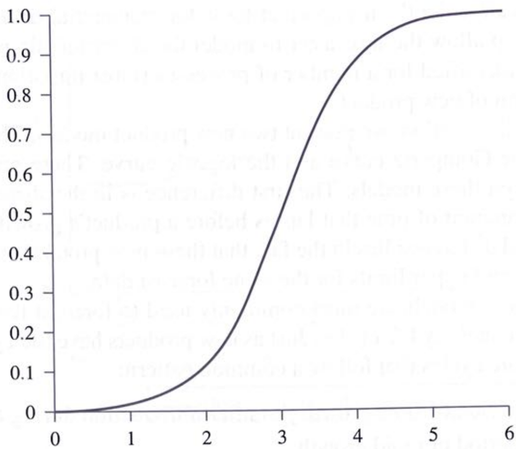

keep two transmitters functioning with, perhaps, two different and costly sets of programming. The costs in extra electricity, tower rental, and insurance would be substantial.

Analyses of many technology introductions (like HDTV) have shown that technology develops initially at a very slow growth rate. But it is also the case that these same technologies soon begin to grow in predictable patterns such as the S-curve shown in Figure 3.12.

The usual reason for the transition from very slow initial growth to rapid growth is often the result of solutions to technical difficulties and the market's acceptance of the new technology. But such growth curves also have their limits; the rapid growth cannot be sustained indefinitely. There are upper limits on the adoptions of new technology or the sales of new products. As the upper limit is reached, a maturity phase occurs in which growth slows and finally ceases. The economic law of diminishing marginal returns is usually at work in these processes.

Figure 3.13 contains information on the early years of HDTV shipments in the United States. During the entire time represented in the graph, there were very few broadcast stations operating in HDTV mode, but in each year, the number of HDTV broadcasters increased and the hours of HDTV programming available also increased. This entire four years of data represent the period of experimentation and slow growth characteristic of all new products and technologies.

When a new technology like HDTV enters the marketplace, we can expect an S-curve to accurately predict future shipments or sales. There is, however, more than one technique that could be used to model this S-curve.

Fortunately, ForecastX™ provides flexible settings for a multitude of situations. If you do not know what the upper limit should be, you can use ForecastX™ to

FIGURE 3.13
Shipments and Cumulative Shipments of HDTV Units (c3t7)

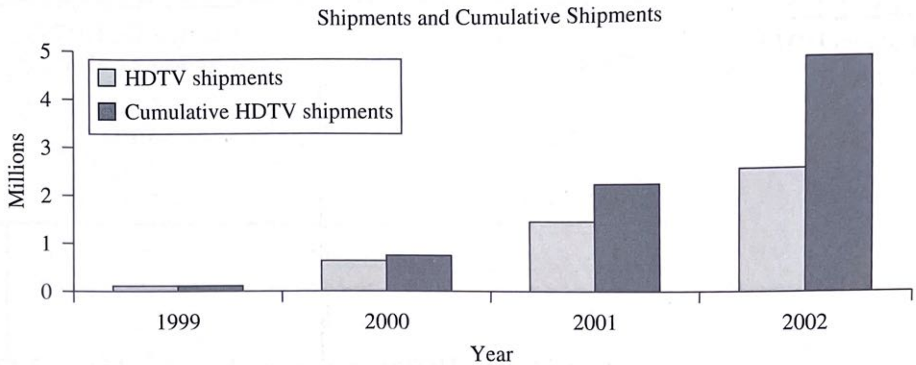

determine the best upper limit to fit your data. Of course, if you know the exact upper limit, ForecastX™ can use it to determine an optimal model.

The two most common forms of S-curves used in forecasting are the Gompertz curve and the logistics curve (also called the Pearl curve). A third useful model called the Bass model was discussed in Chapter 1; we will again cover that model and provide an example of its use.

Gompertz Curve

The Gompertz curve is named after its developer, Benjamin Gompertz, an English actuary. Gompertz applied calculus to actuarial questions and is most well known for *Gompertz's Law of Mortality*. Gompertz's law showed that the mortality rate increases in a geometric progression. Thus, when death rates are plotted on a logarithmic scale, a straight line known as the *Gompertz function* is obtained. The Gompertz curve is the most used actuarial function for investigating the process of aging. The slope of this function is known as the rate of actuarial aging, and differences in the longevity between species are the result in large part of differences in this rate of aging.

The Gompertz function is given as

$$Y_t = Le^{-ae^{-bt}}$$

where:

L = Upper limit of Y
e = Natural number 2.718282 . . .
a and b = coefficients describing the curve (estimated by ForecastX™)

The Gompertz curve will range in value from zero to L as t varies from $-\infty$ to ∞. The curve is widely used in the fields of biology and demography to model (i.e., forecast) the level of populations at a given point in time for plants and animals as well as many organisms. The Gompertz curve is an elegant way to summarize the growth of a population with just a few parameters.

TABLE 3.7
Data on HDTV Shipments (c3t7)

Date	Cumulative HDTV Shipments (millions)
12/31/1999	0.12
12/31/2000	0.77
12/31/2001	2.23
12/31/2002	4.76

Consider the HDTV shipments charted in Figure 3.13; the actual figures are given in Table 3.7.

The Gompertz curve estimated from this data (with the assumption that 248 million televisions is the upper limit) is shown in Figure 3.14. The assumption of 248 million television sets is used as the upper limit because in the year 2002 this was the total number of televisions of all types in use. This rather generous assumption reflects the opportunity for every existing television to be converted to an HDTV. An actual forecaster might choose a different upper limit if another rationale seemed more plausible.

Note the characteristic S-shape to the curve in Figure 3.14. In fitting the curve, we have used the first few data points and the known maximum value for television shipments to estimate a forecast of how HDTV shipments will progress through time. We have ample evidence that this Gompertz function will model the situation well; when color televisions were introduced in the 1960s, their adoption followed a very similar pattern. This form of curve fitting is often used, as it is here, to make forecasts far into the future. Unlike the moving-average and exponential smoothing models, growth curves are routinely used to make mid- to long-range forecasts.

In order to use ForecastX™ to make the estimate shown in Figure 3.14, the Method Selection dialog box would be filled out as shown in Figure 3.15.

FIGURE 3.14
An Estimate of HDTV Shipments Using Four Years of Data (c3t7)

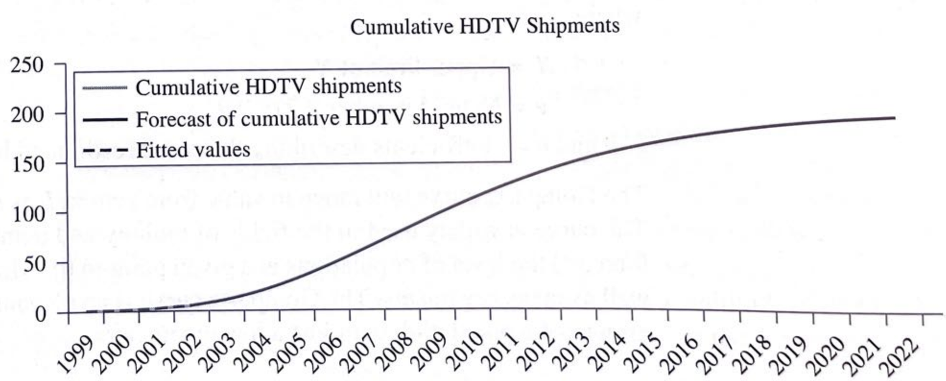

FIGURE 3.15
ForecastX™ Method Selection Dialog Box for a Gompertz Model

Source: John Galt Solutions

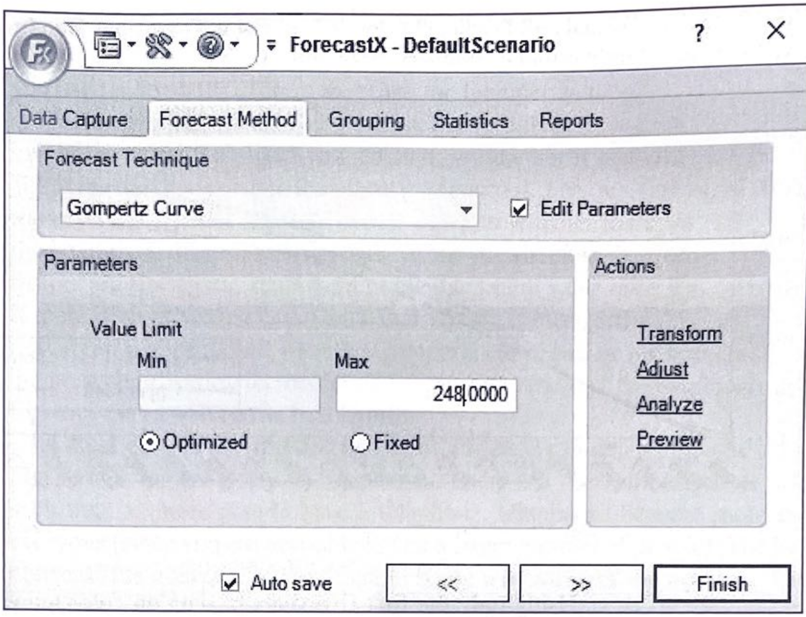

The Edit Parameters box is checked in Figure 3.15 and the maximum value of 248 is entered; this, along with the information on the first four years of actual shipments, allows the estimation of the best-fit Gompertz curve.

The Gompertz curve is best used in situations where it becomes more difficult to achieve an increase in the growth rate as the maximum value is approached. We will see that this is the exact opposite of the recommendation for the best situation in which to use the logistics function. Consider the adoption of color televisions shown in Table 3.8.

TABLE 3.8
Color Television Adoption in Percentages (c3t8)

Year	Percent Adoptions	Year	Percent Adoptions
Dec-65	***0***	Dec-78	79.37021852
Dec-66	***6.145835684***	Dec-79	82.72351665
Dec-67	***12.72965645***	Dec-80	85.59728784
Dec-68	***19.64872441***	Dec-81	88.03946805
Dec-69	***26.77512032***	Dec-82	90.10007711
Dec-70	33.96440431	Dec-83	91.82826105
Dec-71	41.06698245	Dec-84	93.27031978
Dec-72	47.94034951	Dec-85	94.46854445
Dec-73	54.46013052	Dec-86	95.46066873
Dec-74	60.52818203	Dec-87	96.27975333
Dec-75	66.07679819	Dec-88	96.95435387
Dec-76	71.06899248	Dec-89	97.50885698
Dec-77	75.49558333	Dec-90	97.96390133

FIGURE 3.16 Actual and Predicted Adoptions of Color Televisions Obtained by Using the First Five Years of Adoptions and a Gompertz Estimate (c3t8)

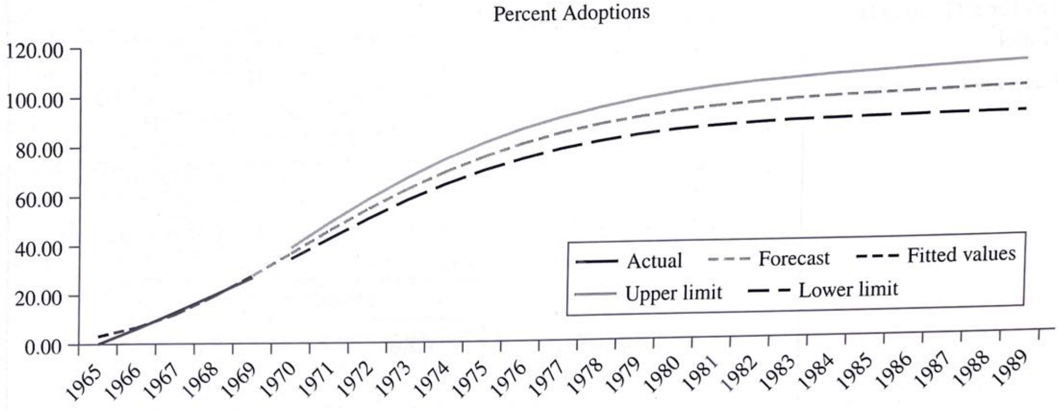

Using only the first five years of data on color television adoptions (and the assumption of a maximum of 100 percent), it is possible to very closely approximate the future growth pattern with a Gompertz function. Figure 3.16 shows the actual and predicted shipments obtained using only the first five data points from Table 3.8.

Had you been asked in late 1969 to forecast color television adoptions with only the first five years of annual data, you would have produced very accurate forecasts if you had used a Gompertz model. The assumption regarding the increased difficulty of obtaining the maximum value as it is approached probably describes color television adoptions quite well. The same assumption might also apply to HDTV adoptions and sales since the situation is similar.

Logistics Curve

The logistics curve is a second way of forecasting with sparse data and is also used frequently to forecast new-product sales. The logistics curve has the following form:

$$Y_t = \frac{L}{1 + ae^{-bt}}$$

where:

L = Upper limit of Y
e = Natural number 2.718282 . . .
a and b = coefficients describing the curve (estimated by ForecastX™)

Just as in the Gompertz function, there is an upper limit to Y called L, and e is the base of natural logarithms. The logistics curve is symmetric about its point of inflection (the upper half of the curve is a reflection of the lower half); the Gompertz curve is not necessarily symmetric about its points of inflection. Why would you use the logistics curve rather than the Gompertz curve?

The answer lies in whether, in a particular situation, it is easier to achieve the maximum value the closer you get to it, or whether it becomes more difficult to attain the maximum value the closer you get to it. The question of which function to use in a particular estimate comes down to whether there are factors assisting the attainment of the maximum value once you get close to it, or whether there are factors preventing the attainment of the maximum value once it is nearly attained. If there is an offsetting factor such that growth is more difficult to maintain as the maximum is approached, then the Gompertz curve will be the best choice. If there are no such offsetting factors hindering the attainment of the maximum value, the logistics curve will be the best choice.

A clear case of the appropriate use of a logistics function might be the prediction of U.S. households with telephones. There is a "network effect" at work here such that, as more people have a telephone, telephones become more useful to everyone (since you are now able to call a larger number of people). The larger the network, the greater the advantage to being a member of the network. The more recent case of the adoption of cellular telephones would likely progress in much the same manner as the original telephone adoption. The adoption data for cellular telephones in the United States is presented in Table 3.9.

TABLE 3.9
Percentage of Cellular Telephone Adoption in the United States (c3t9)

Year	Cellular Telephone Adoption Percentage
12/31/1986	*0*
12/31/1987	*0.989132021*
12/31/1988	*2.47063319*
12/31/1989	*4.661420125*
12/31/1990	*7.840611381*
12/31/1991	12.33023714
12/31/1992	18.43261935
12/31/1993	26.3098137
12/31/1994	35.82721276
12/31/1995	46.4488152
12/31/1996	57.30203601
12/31/1997	67.43707273
12/31/1998	76.13611792
12/31/1999	83.07839603
12/31/2000	88.3037847
12/31/2001	92.06622305
12/31/2002	94.68961176
12/31/2003	96.47806155
12/31/2004	97.6786669

FIGURE 3.17 Actual and Predicted Adoptions of Cellular Telephones in the United States (c3t9)

Year	Cellular Telephone Adoption Percentage	Logistics
31-12-1986	0	0.54
31-12-1987	0.989132021	1.09
31-12-1988	2.47063319	2.20
31-12-1989	4.661420125	4.37
31-12-1990	7.840611381	8.48
31-12-1991		15.84
31-12-1992		27.65
31-12-1993		43.70
31-12-1994		61.18
31-12-1995		76.19
31-12-1996		86.66
31-12-1997		92.96
31-12-1998		96.40
31-12-1999		98.20
31-12-2000		99.10
31-12-2001		99.56
31-12-2002		99.78
31-12-2003		99.89
31-12-2004		99.95
31-12-2005		99.97

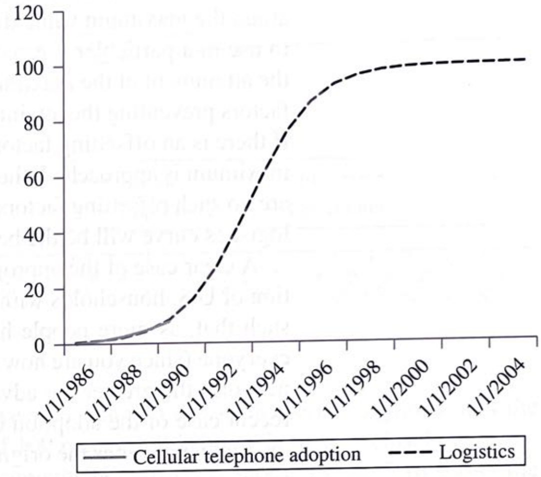

By fitting a logistics curve to the first five years of cellular telephone data, the results in Figure 3.17 are calculated.

It is not surprising that a logistics estimate of cellular telephone adoption works so well; as more individuals have cellular telephones, it becomes more advantageous to have one yourself. Thus, there is a factor assisting the attainment of the maximum value the closer you get to the maximum value (i.e., the network effect).

Note that there should be some theoretical reason for choosing a logistics function for your forecast estimate before estimating the model. In the case of cellular phones, the hypothesized existence of a network effect would lead a researcher to choose a logistics model. The ForecastX^{TM} Method Selection dialog box used to select the cellular telephone model appears in Figure 3.18.

Let's generalize our suggestions for employing the Gompertz and logistics models. Use a Gompertz model when you expect it to be more difficult to attain constant improvement as the maximum value is approached. On the other hand, select a logistics model when there are factors that help maintain improvements as the maximum value is approached. At times it will not be easy to predict which of the two models may work best; in those instances ForecastX^{TM} allows the choice of "New Product Forecasting" as a selection in the Method Selection dialog box.

FIGURE 3.18
ForecastX™ Method Selection Dialog Box for a Logistics Model

Source: John Galt Solutions

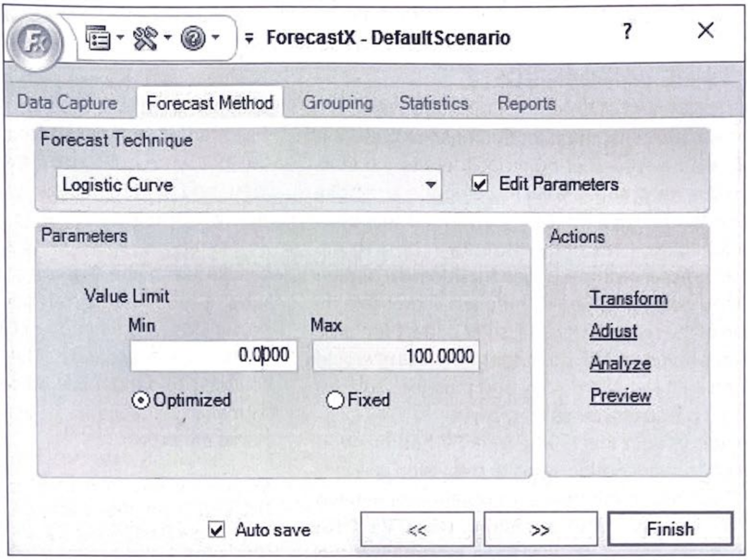

Choosing New Product Forecasting allows ForecastX™ to choose the optimal model from among three contenders: the logistics model, the Gompertz model, and a Probit curve.

Bass Model

Named after Professor Frank M. Bass, this model has been used for over 30 years to forecast the diffusion of innovations, to forecast the penetration of new products in a market, and in a variety of biological, medical, and scientific forecasts. This is a relatively simple model in which only three parameters are chosen by the researcher.

As they are used in ForecastX™, the three parameters are p, r, and qbar, where:

p = The innovation rate

r = The imitation rate (called q in the forecasting literature)

qbar = The cumulative value of all the historical values

The Bass model could be called a model of social contagion where the p (the innovation rate) refers to the probability of initial purchase of a new good independent of the influence of previous buyers (i.e., with no network effect considered). The r (the imitation rate) refers to the pressure of imitation on previous purchasers. The Bass model would appear most often in a graph like the S-curves we have been examining. As we indicated in Chapter 1, getting the estimates of the three parameters of the model is the difficult part. We can be helped significantly here by using past studies to suggest parameters that may place us in the ballpark for our own estimates.

The Bass Model in Action 1

CHRISTOPHE VAN DEN BULTE

The U.S. Department of Energy (DOE) in 1980 used the Bass model to forecast the adoption of solar batteries. The DOE used a survey of home builders to aid in its initial choices for *p* and *q* values. Using these empirically suggested values, the DOE concluded that solar battery technology was not sufficiently robust to encourage word-of-mouth propagation. Because of their finding, they postponed their proposed wide-scale introduction of the technology until solar battery technology had improved to the point that new users would be satisfied with the technology and thus the higher *q* value would predict faster sales growth.

A decade later in the 1990s, DirecTV had planned a launch of its subscription satellite television delivery service. Prudently, it attempted to obtain a prelaunch forecast for five years into the future. DirecTV's forecast was again based on the Bass model, and the *p* and *q* values were also obtained from a survey of prospective users; this information was combined with histories of similar services. The forecasts produced in 1992 were quite good from the company's point of view, and after the fact, the estimates compared favorably with the actual 1994 to 1999 experience.

Numerous other firms have reported productive results using the Bass model. RCA in the mid-1980s used a modified Bass model to forecast the sales of music CDs as a function of CD player sales. The model proved quite accurate. The Bass model is also used routinely to predict box office revenues for movies and to make decisions on how many screens to use for a particular movie.

Source: Van den Bulte, Christophe, "Want to Know How Diffusion Speed Varies across Countries and Products? Try Using a Bass Model," *PDMA Visions* 26, no. 4 (2002), pp. 12–15.

Christopher Van den Bulte of the Wharton School has constructed a database of 1,586 sets of *p* and *q* parameters from 113 separate recent articles.[5] Some suggestions from Van den Bulte's work appear in Table 3.10.

An interesting set of patterns emerges from this large number of *p* and *q* estimations. Recall that the parameter Van den Bulte refers to as *q* is the *r* parameter (the imitation rate) in ForecastX™ What estimates for *p* and *q* would be best for your product? Van den Bulte took as a baseline durable goods launched in the United States in 1976. The *p* factor measures the intrinsic tendency for an individual to adopt a new product, while the *q* measures the "word of mouth" or "social contagion" effect on purchases. Van den Bulte recommends that when a forecaster tries to set the values of *p* and *q* in a Bass model, you should use a range of values within his estimated confidence interval (given in Table 3.10). For countries with a collectivist mentality (like Japan) as opposed to an individualistic mentality (like the United States), a higher *q* value is better. People in collectivist cultures care more about what others think of them, according to Van den Bulte's study. In countries with higher purchasing power, the *p* tends to be higher. More disposable income makes it easier to adopt innovations. Finally, products that exhibit significant network effects or require heavy investment in complementary infrastructure (like television and the cellular telephone) will have higher values for *q*. Van den Bulte has summarized these results in a set of conclusions presented in Table 3.11.

[5] Christophe Van den Bulte, "Want to Know How Diffusion Speed Varies across Countries and Products? Try Using a Bass Model," *PDMA Visions* 26, no. 4 (2002), pp. 12–15.

TABLE 3.10
Van den Bulte's p and q Estimates from Selected Articles

p Estimates	Best Guess	90% Confidence Interval	
Baseline case:			
U.S. consumer, durable, launch in 1976	0.409	0.355 (t^0)	0.471
For other cases, multiply by the following factors:			
Cellular telephone	0.635	0.465	0.868
Nondurable product	0.931	0.713	1.216
Industrial	1.149	0.909	1.451
Noncommercial innovation	2.406	1.488	3.891
Western Europe	0.949	0.748	1.203
Asia	0.743	0.571	0.966
Other regions	0.699	0.429	1.137
For each year after 1976, multiply by	1.028	1.018	1.039
***q* Estimates (labeled *r* in ForecastX)**			
Baseline case:			
U.S. consumer, durable, launch in 1976	0.016	0.012	0.021
For other cases, multiply by the following factors:			
Cellular telephone	0.226	0.125	0.409
Nondurable product	0.689	0.415	1.143
Industrial	1.058	0.679	1.650
Noncommercial innovation	0.365	0.146	0.910
Western Europe	0.464	0.296	0.729
Asia	0.595	0.360	0.981
Other regions	0.796	0.315	2.008
For each year after 1976, multiply by	1.021	1.002	1.041

TABLE 3.11
Van den Bulte's Conclusions Regarding p and q Values

- There are systematic regional differences in diffusion patterns.
- The average coefficient of innovation p (speed of takeoff) in Europe and Asia is roughly half of that in the United States.
- The average coefficient of imitation q (speed of late growth) in Asia is roughly a quarter less than that in the United States and Europe.
- Also, economic differences explain national variations in speed better than cultural differences do.
- There are systematic product differences in diffusion patterns. For instance, takeoff is slower for nondurables and products with competing standards that require heavy investments in infrastructure, while late growth is faster for industrial products and products with competing standards, which require heavy investments in infrastructure.

Table 3.12 presents data for the adoption of telephone-answering devices in the United States.

Using only the first five observations in Table 3.12, it is possible to accurately represent the entire adoption cycle for telephone-answering devices. After some

TABLE 3.12
Adoption of Telephone-Answering Devices in the United States (c3t12)

Year	Adoption
Dec-84	0
Dec-85	3.030551
Dec-86	7.351138
Dec-87	13.29582
Dec-88	21.08724
Dec-89	30.67365
Dec-90	41.59211
Dec-91	52.98598
Dec-92	63.84035
Dec-93	73.31923
Dec-94	80.98819
Dec-95	86.81843
Dec-96	91.04448
Dec-97	94.00304
Dec-98	96.02418
Dec-99	97.38195
Dec-00	98.28381
Dec-01	98.87837
Dec-02	99.26839
Dec-03	99.52341
Dec-04	99.6898

trial and error, the researcher has selected a p value of 0.035 and an r value of 0.406. Note that the r value in ForecastX™ is the same as the q value used for explaining the imitation rate in the Bass model, as shown in Figure 3.19. The qbar value is 100 because we are working with percentages.

FIGURE 3.19
ForecastX™ Method Selection Dialog for the Bass Model

Source: John Galt Solutions

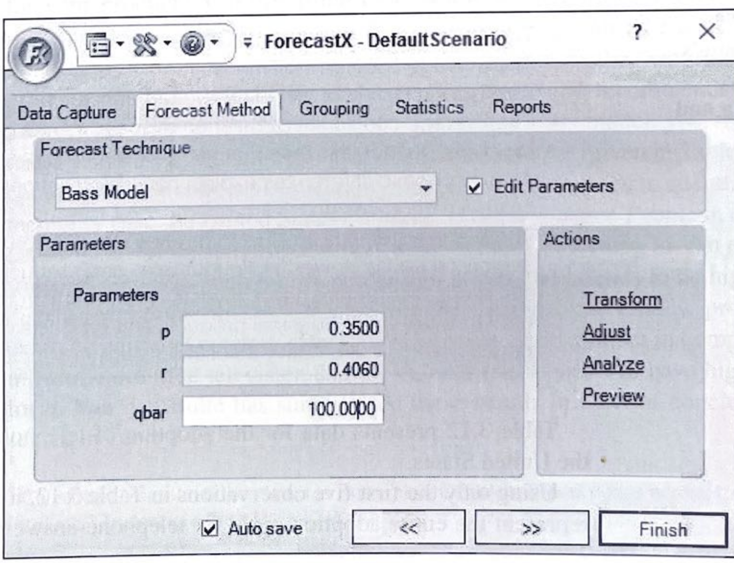

GURE 3.20
ass Model of
lephone-Answering
achine Adoptions
the United
ates (c3t12)

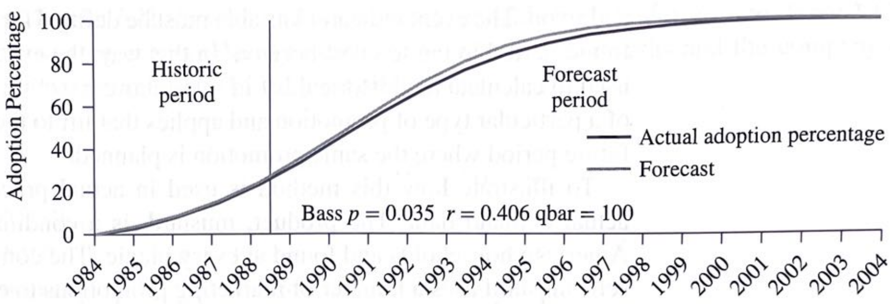

The resulting plot of actual and predicted values in Figure 3.20 shows a model that closely approximates the actual occurrence in the United States for answering-machine adoptions.

EVENT MODELING

Event modeling is a feature of some exponential smoothing programs such as ForecastX™. This feature allows the user to specify the time of one or more special events, such as irregular promotions and natural disasters, in the calibration data. For each type of special event, the effect is estimated and the data adjusted so that the events do not distort the trend and seasonal patterns of the time series.

When forecasting sales or demand in a highly promoted market, using this smoothing technique will significantly improve forecast accuracy. Consider the case of a manufacturer of a popular condiment (e.g., ketchup, mustard, steak sauce, and so on). This type of product tends to be highly seasonal and also tends to be aggressively promoted by marketers. Is there a method for modeling the effect of future promotions on the sales or demand for such a product?

When forecasting sales r demand in a highly romoted market, using e smoothing technique f event modeling will ignificantly improve orecast accuracy.

The answer to this dilemma is *event modeling*. By using the basic smoothing models already developed earlier in the chapter as a base, an event model may be generated to replicate the effects of various promotions and combinations of promotions.

The method of event modeling follows in the same pattern for the smoothing models already examined: after the systematic patterns are identified in the historical data, the exponential smoothing method uses smoothing equations for each component in the series to estimate and build up structural patterns. The event model adds a smoothing equation for each of the "events" identified as being important. The weights for each smoothing equation are represented by a parameter.

Event models are analogous to seasonal models: just as each month is assigned its own index for seasonality, so, too, each event type is assigned its own index for a specific promotional activity. For example, when monthly data are used, the seasonal index for a particular month is updated at regular intervals, each time that month recurs. However, event adjustments are created through the use of an indicator variable that assigns an integer for each event type to the period during which it recurs. Thus, one example of integer value assignment would be that 0 indicates a period where no event has occurred, 1 indicates a period where a free-standing insert (FSI) was circulated, 8 indicates a period where thematic advertising was used,

Event models are nalogous to seasonal models.

and so on. The event indicator variable must be defined for each historic period *and* future period in the forecast horizon. In this way, the event smoothing equation is used to calculate the historical lift in sales above baseline that occurred as a result of a particular type of promotion and applies that lift to the baseline forecast in the future period where the same promotion is planned.

To illustrate how this method is used in actual practice, we examine some actual demand data. The product, mustard, is a condiment commonly used in American households and found at every picnic. The company that produces and sells mustard uses a number of marketing promotions to enhance sales and maintain market share. Free-standing inserts are perhaps the most common of the promotions for this type of product; these are the familiar coupons found in Sunday newspapers or in online promotions and redeemable when the item is purchased. These FSIs are often used in conjunction with themed advertising campaigns, especially during particular seasons of the year. Our condiment manufacturer uses a separate event value, 7, to stand for the combination of FSIs and an advertising campaign. On-pack coupons are a separate type of coupon usually attached to the product packaging itself and redeemed at the cash register at checkout.

In addition to adjusting the price to the consumer through coupons, the mustard manufacturer also adjusts the price to the jobber by reducing case prices for a short period of time. When this takes place, it is common for jobbers to stock up on the reduced-price item and delay future purchases. Because of this, the manufacturer uses two event values called *load* and *deload* to signify periods of reduced prices and the periods immediately following such a promotion; these are actually two separate events.

The event values for this particular condiment manufacturer are listed in the following table.

Event Indices Legend:

0 = Nothing
1 = FSI
2 = Thematics
3 = Big Load
4 = After Load
5 = Deload
6 = Light Load
7 = FSI / Act Media
8 = Cross Cpn

FSI = free standing inserts. These are cents-off coupons distributed in newspapers
Thematics = Themed ad campaign
Big Load = Large trade promotion—often a deep drop in the case price for the retailer
Deload = Month after effect of a "load"
Act Media = Radio, television, print ad campaign
Cross Coupons = Cents-off coupons placed directly on the packaging of other goods

Figure 3.21 shows monthly historical demand of mustard over time. Table 3.13 shows the events related to each of these historical months and the company's planned promotions for the next six months.

FIGURE 3.21
Mustard Demand
(c3t13)

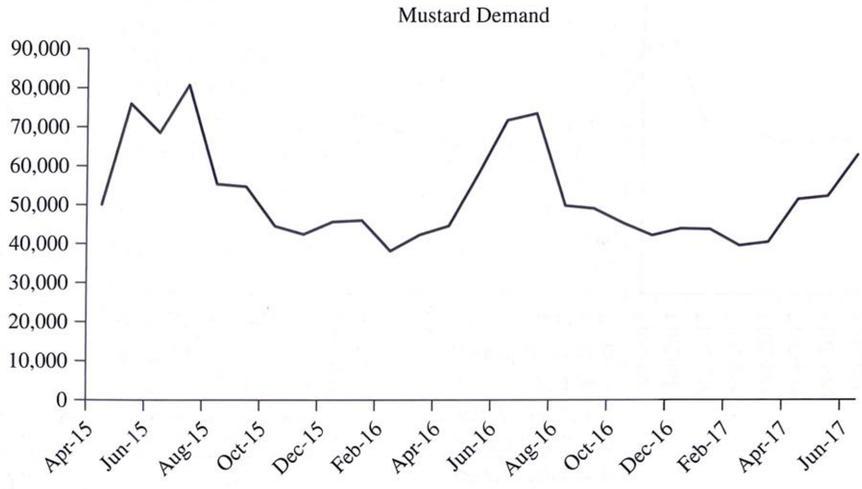

TABLE 3.13
An Event Model Example (c3t13)

Date	Mustard	Event Index
Apr-15	50,137	7
May-15	76,030	7
Jun-15	68,590	3
Jul-15	80,681	4
Aug-15	55,228	5
Sep-15	54,577	0
Oct-15	44,384	8
Nov-15	42,337	0
Dec-15	45,512	6
Jan-16	45,798	4
Feb-16	38,045	5
Mar-16	42,127	0
Apr-16	44,422	2
May-16	57,662	1
Jun-16	71,427	6
Jul-16	73,269	5
Aug-16	49,695	8
Sep-16	49,021	1
Oct-16	45,263	0
Nov-16	42,210	1
Dec-16	43,968	6
Jan-17	43,778	4

(continued on next page)

TABLE 3.13 (continued)

Date	Mustard	Event Index
Feb-17	39,524	0
Mar-17	40,476	0
Apr-17	51,167	2
May-17	51,916	1
Jun-17	62,274	6
Jul-17		4
Aug-17		5
Sep-17		0
Oct-17		2
Nov-17		1
Dec-17		6

Legend:

0 = Nothing
1 = FSI
2 = Thematics
3 = Big Load
4 = After Load
5 = Deload
6 = Light Load
7 = FSI / Act Media
8 = Cross Cpn

FSI = free standing inserts. These are off-cents coupons distributed in newspapers
Thematics = Themed add campaign
Big Load = Large trade promotion – often a deep drop in the case price for the retailer
Deload = Month after effect of a "load"
Act Media = Radio, television, print add campaign
Cross Coupons = Cents off coupons placed directly on the packaging of other goods

Using a Winters' smoothing model on these data picks up the implied seasonality and trend quite well; the calculated level, seasonal, and trend are 0.05, 0.88, and 0.26, respectively. This indicates that there is very little trend in the data but a high degree of seasonality. Actually, some of the apparent seasonality is not seasonality at all; instead, it is "induced seasonality" caused by the company's various promotions. The MAPE for the Winters' model without including events is 5.18 percent. This result is shown in Figure 3.22.

Using the Winters' smoothing model again, but with eight event smoothing factors added in, the level, seasonal, and trend factors are 0.2, 0.92, and 0.26. By examining these factors, we see that there is definitely little trend, but now the seasonality has also apparently changed. The seasonality has not disappeared; it is changed by the eight event indices. When the eight events are included, the MAPE improves to 2.62 percent. This result is shown in Figure 3.23.

3.22 Mustard Demand Forecast Using Winters' Exponential Smoothing. No Events Are MAPE = 5.18%

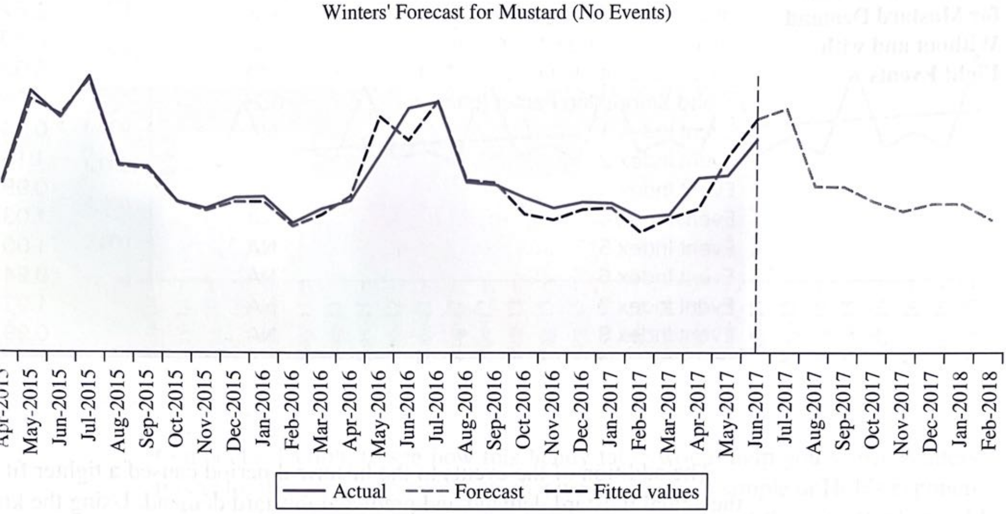

3.23 Mustard Demand Forecast Using Winters' Exponential Smoothing; Including ents. MAPE = 2.62%

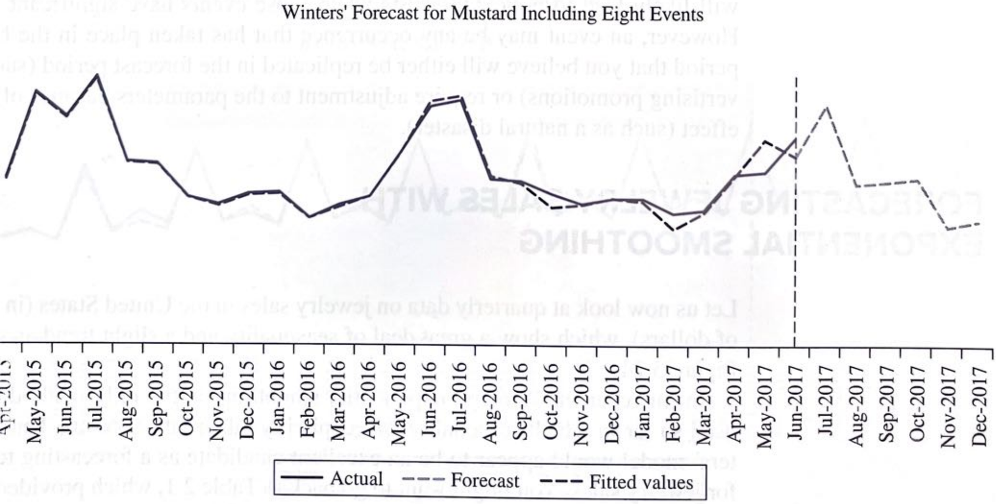

TABLE 3.14
Summary of Results of a Winters' Model for Mustard Demand Without and with Eight Events

	Winters' Model	Winters' Model with Event Indices
Historical MAPE	5.18%	2.62%
Level Smoothing Factor (alpha)	0.05	0.20
Seasonal Smoothing Factor (beta)	0.88	0.92
Trend Smoothing Factor (gamma)	0.26	0.26
Event Index 1	NA	0.94
Event Index 2	NA	1.16
Event Index 3	NA	0.99
Event Index 4	NA	1.03
Event Index 5	NA	1.00
Event Index 6	NA	0.94
Event Index 7	NA	1.03
Event Index 8	NA	0.99

The addition of the events to the historical period caused a tighter fit between the actual mustard demand and predicted mustard demand. Using the knowledge of the planned company promotions for the next six months allows the forecaster to calculate a much better picture of predicted demand than the Winters' model alone.

In this particular case, we used the Winters' model as a base because we believed the original data had both trend and seasonality. If the data had lacked trend or seasonality, we could have used simple smoothing as the base model. ForecastX^{TM} allows a number of models to be used as the underlying basis for event forecasting.

Ignoring events (usually promotions) that a company has scheduled in advance will likely lead to poorer forecasts when those events have significant impacts. However, an event may be any occurrence that has taken place in the historical period that you believe will either be replicated in the forecast period (such as advertising promotions) or require adjustment to the parameters because of its large effect (such as a natural disaster).

FORECASTING JEWELRY SALES WITH EXPONENTIAL SMOOTHING

Let us now look at quarterly data on jewelry sales in the United States (in millions of dollars), which show a great deal of seasonality and a slight trend as shown in Figure 3.24.

In this chapter, Winters' exponential smoothing is the only method we have used so far in which seasonality was explicitly taken into account. Thus, a Winters' model would appear to be an excellent candidate as a forecasting technique for jewelry sales. You might want to go back to Table 2.1, which provided a guid

FIGURE 3.24 **Jewelry Sales (M$)**

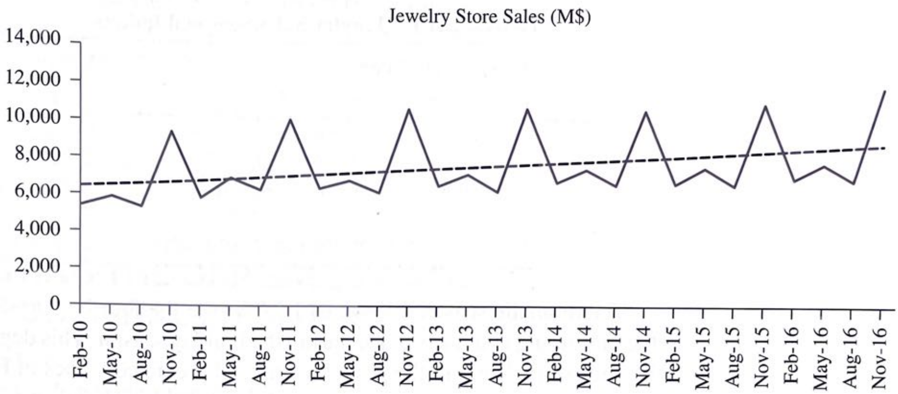

to model selection, to see how this handy table would help you select Winters' model for this series. You will note that we do not apply simple or Holt's exponential smoothing to the jewelry sales data because neither of those methods would be an appropriate model, given the guidelines in Table 2.1.

Applying the Winters' model, ForecastX™ finds the optimum values for the weights to be level, 0.38; seasonal, 0.49; trend, 0.02. The historic MAPE for this Winters' model is 2.25 percent. The seasonal value shows that we indeed do have a rather high degree of seasonality, and the trend value indicates trend is also present. The resulting forecast is shown graphically in Figure 3.25.

FIGURE 3.25 **Jewelry Sales with a Winters' Forecast. MAPE = 2.25%**

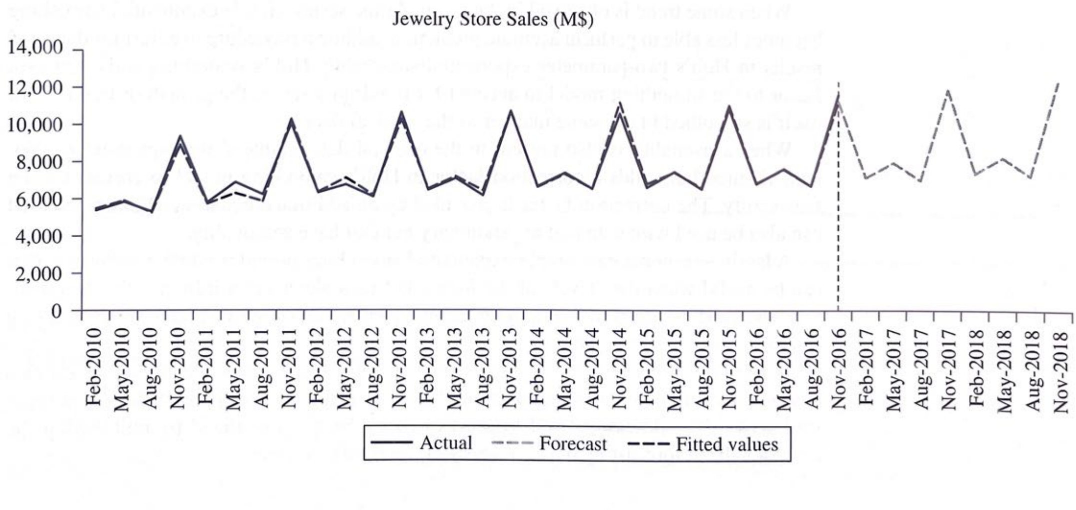

As shown in Table 3.15, the seasonal indices for this model are quite revealing:

TABLE 3.15 Jewelry Sales Seasonal Indices

Seasonal Indices	Value
Index 1 (Quarter 1 = February)	0.87
Index 2 (Quarter 2 = May)	0.93
Index 3 (Quarter 3 = August)	0.82
Index 4 (Quarter 4 = November)	1.38

Months are the middle month of each calendar quarter.

The dramatic seasonal index of 1.38 for the the fourth quarter is likely due to the gift giving that takes place during the holiday season. This degree of seasonality shows up clearly in the plot of actual and predicted values in Figure 3.24.

Summary

If the time series you are forecasting is a stationary one, the moving-average method of forecasting may accurately predict future values. The moving-average method calculates the average of the past observations, and this average becomes the forecast for the next period.

When recent-past observations are thought to contain more information than distant-past observations, some form of exponential smoothing may be appropriate. Exponential smoothing provides self-correcting forecasts that adjust so as to regulate the forecast values by changing them in the opposite direction from recent errors. It is a characteristic of smoothing models in general, however, that their forecasts lag behind movements in the original time-series data. Exponential smoothing requires the specification of a smoothing constant, which determines the relative weights accorded to recent as opposed to more distant historical observations.

A suggested method for choosing an optimal smoothing constant is to minimize the mean absolute percentage error (MAPE).

When some trend is observed in the original time series, simple exponential smoothing becomes less able to perform accurate prediction; adding a procedure to adjust for the trend results in Holt's two-parameter exponential smoothing. Holt's smoothing adds a growth factor to the smoothing model to account for trend; in a sense, the growth or trend factor itself is smoothed in the same manner as the original data.

When seasonality is also present in the original data, Winters' three-parameter exponential smoothing adds a correction factor to Holt's smoothing model to correct for the seasonality. The correction factor is provided by an additional equation. Winters' method can also be used with data that are stationary but that have seasonality.

Adaptive–response-rate single exponential smoothing provides another technique that can be useful when the "level" of the forecasted variable changes infrequently. Adaptive-response models adjust the smoothing factor for changing conditions rather than choosing a constant smoothing factor.

In addition to trying Winters' exponential smoothing for seasonal data, you might also deseasonalize the data and then use another forecasting tool to forecast the deseasonalized series. The deseasonalized forecast can then be reseasonalized by multiplying the deseasonalized forecast by the corresponding seasonal indices.

Event models are very useful for many business applications. The events may be planned and implemented by a business, or they may be caused by forces outside of the business. Examples of the latter type could be a strike, a natural disaster, or some action by another business.

Integrative Case:

The Gap

FORECASTING THE GAP SALES DATA WITH EXPONENTIAL SMOOTHING

The sales of The Gap stores for the 44 quarters covering April 2006 through January 2017 are once again shown below. From this graph, it is clear that The Gap sales are quite seasonal and are increasing over time. Recall that the 2004 data are used as a holdout period.

(c3Gap Sales Data)

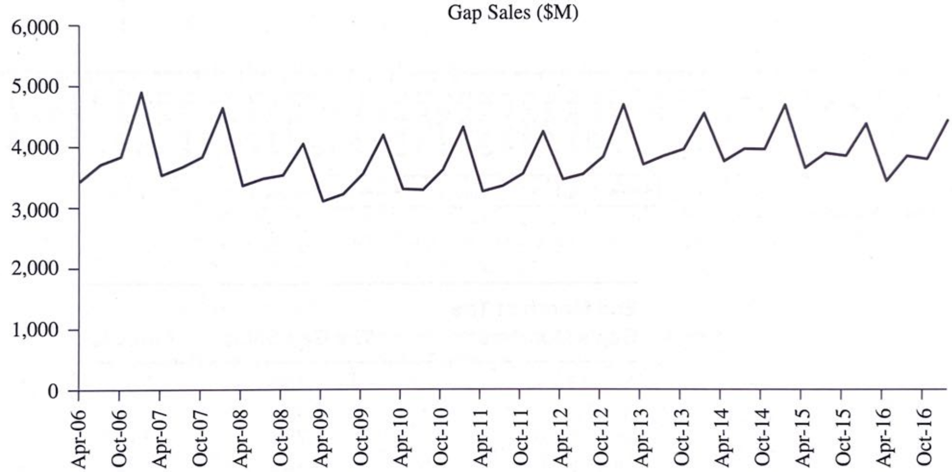

Gap Sales for Chapter 3 Case.

Case Questions

1. Using The Gap data for April 2006 through January 2016, what exponential smoothing model do you think would be the most appropriate if you want to develop a quarterly forecast for April 2016 through January 2017 (Gap's 2016 fiscal year) sales? Explain why.
2. Make a one-year (four-quarter) forecast for The Gap sales using the method you selected. What is the MAPE for this period of historic data?

3. To evaluate your forecast accuracy for the four quarters of the 2016 fiscal year, calculate the MAPE for those four months that were not used in developing the forecast. Compare the historic period MAPE with the forecast period MAPE.
4. What are the seasonal indices for the The Gap sales, and what do they tell you about this company's sales pattern?

Solutions to Case Questions

1. Of the exponential smoothing models discussed in the text, the one that is most appropriate for The Gap sales data is Winters' exponential smoothing. This model takes both trend and seasonality into account. Allowing ForecastX™ to determine the optimal smoothing weights, we obtain level = 0.57, seasonal = 1.00, and trend = 0.00.
2. The forecast graph is shown below. The MAPE using the historic period is 2.60 percent,

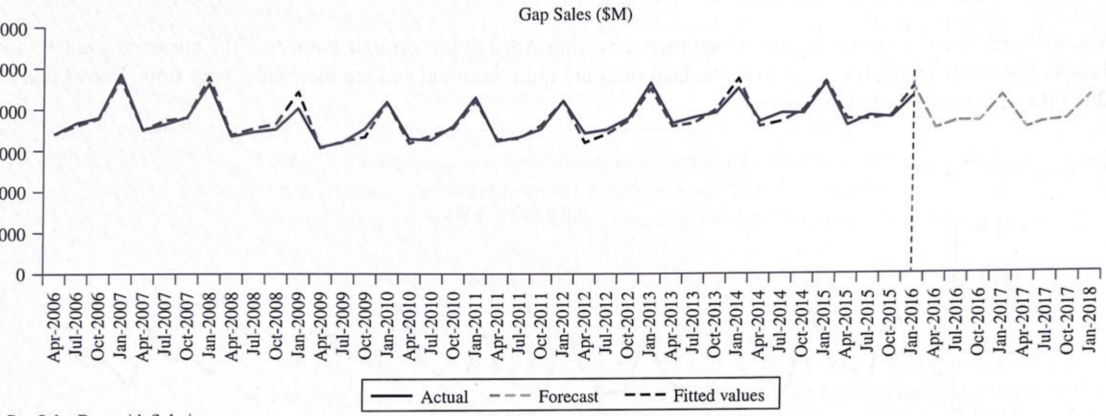

c3 Gap Sales Data with Solutions

End Month of The Gap's Quarters	The Gap Sales	Forecast
April-16	3,438	3,573
July-16	3,851	3,763
October-16	3,798	3,747
January 17	4,429	4,385

3. Using The Gap sales data and the forecast values (both are shown in the above table), the MAPE for the 2016 fiscal year is calculated as shown in the following table:

Date	Actual Sales	Forecast	Error	Absolute Error	Absolute % Error
Apr-16	3,438	3,573	−135	135	3.93
Jul-16	3,851	3,763	88	88	2.29
Oct-16	3,798	3,747	51	51	1.34
Jan-17	4,429	4,385	44	44	0.99
					MAPE = 2.14%

In this case, it turns out that the forecast period MAPE (2.14 percent) is slightly lower than the historic MAPE (2.60 percent). Typically, a forecast period MAPE will be larger than the historic MAPE.

4. The seasonal indices are:

Seasonal Indexes	Value
Index 1April	0.88
Index 2 July	0.94
Index 3 October	0.97
Index 4 January	1.21

These indices suggest The Gap has strong holiday season sales since the fourth quarter (ending in January) would include the months when holiday shopping peaks.

USING FORECASTX™ TO MAKE EXPONENTIAL SMOOTHING FORECASTS

What follows is a brief discussion of how to use ForecastX™ for preparing an exponential smoothing forecast. This also serves as a further introduction to the ease of use of ForecastX™. The illustration used here is for a forecast of The Gap data that has trend and seasonality.

First, put your data into an Excel spreadsheet in column format, such as the The Gap data shown in the table below. Once you have your data in this format, while in Excel put your cursor in any cell with sales data. For example, put the cursor in cell B4 in the Excel file. See the Excel file c3 Gap Sales Data with Solutions.

Date	Gap Sales ($M)	Date	Gap Sales ($M)
Apr-06	3441	Jul-09	3,245
Jul-06	3,714	Oct-09	3,589
Oct-06	**3,851** Example cell in which to place your cursor (cell B4).	Jan-10	4,236
		Apr-10	3,329
		Jul-10	3,317
Jan-07	4,919	Oct-10	3,654
Apr-07	3,549	Jan-11	4,364
Jul-07	3,685	Apr-11	3,295
Oct-07	3,854	Jul-11	3,386
Jan-08	4,675	Oct-11	3,585
Apr-08	3,384	Jan-12	4,283
Jul-08	3,499	Apr-12	3,487
Oct-08	3,561	Jul-12	3,575
Jan-09	4,082	Oct-12	3,864
Apr-09	3,127	Jan-13	4,725

(continued on next page)

(continued)

Date	Gap Sales ($M)	Date	Gap Sales ($M)
Apr-13	3,729	Apr-15	3,657
Jul-13	3,868	Jul-15	3,898
Oct-13	3,976	Oct-15	3,857
Jan-14	4,575	Jan-16	4,385
Apr-14	3,774	Apr-16	3,438
Jul-14	3,981	Jul-16	3,851
Oct-14	3,972	Oct-16	3,798
Jan-15	4,708	Jan-17	4,429

Then select Add-Ins from the top of your Excel sheet. In Add-Ins, select "Forecast" as shown below:

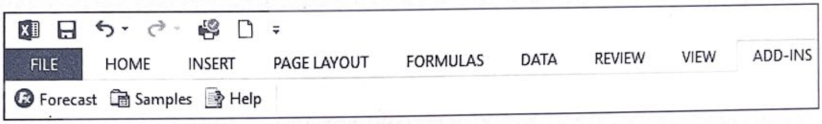

This will start ForecastX^{TM}. The following dialog box will appear.

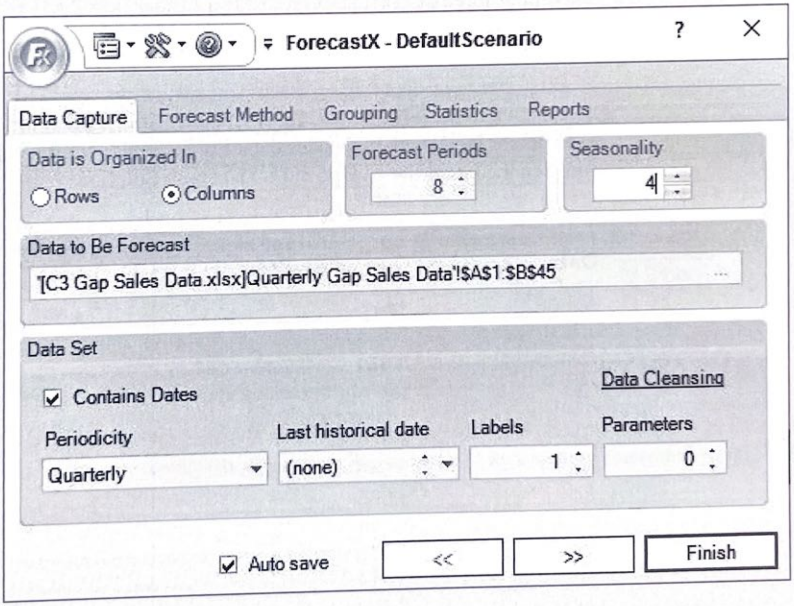

Source: John Galt Solutions

Check the items in this box to be sure ForecastX has correctly identified your data for periodicity, number of lines with labels (usually 1), whether your data are in rows or columns (typically columns), and the number of periods you want to forecast (8 in this example). Here we have set the seasonality to 4 since we know The Gap data have a four-period seasonality. Note that "Auto save" is also checked. This saves all selections, which makes it easier when repeating the same process for more than on Excel sheet. Once you have verified the selections, click the **Forecast Method** tab at the top. The following screen appears.

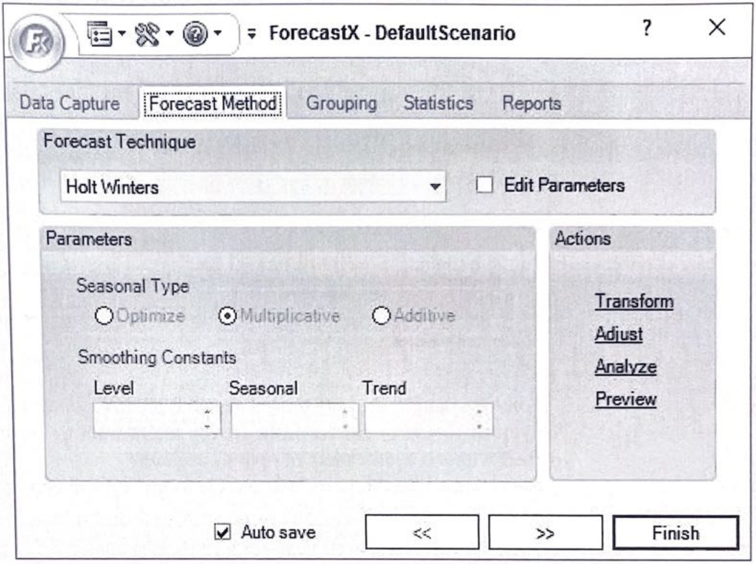

Source: John Galt Solutions

Click the down arrow in the **Forecasting Technique** window and select **Holt Winters,** which is what ForecastX™ calls Winters'exponential smoothing. This would be an appropriate method for data such as The Gap series. You can enter your own values for the constants, or you can leave those spaces blank and let ForecastX™ select the best set of values. Letting the software pick optimal values for the constants is most common.

We will not be using the "Grouping" tab, so now click on the "Statistics" tab. The following will appear:

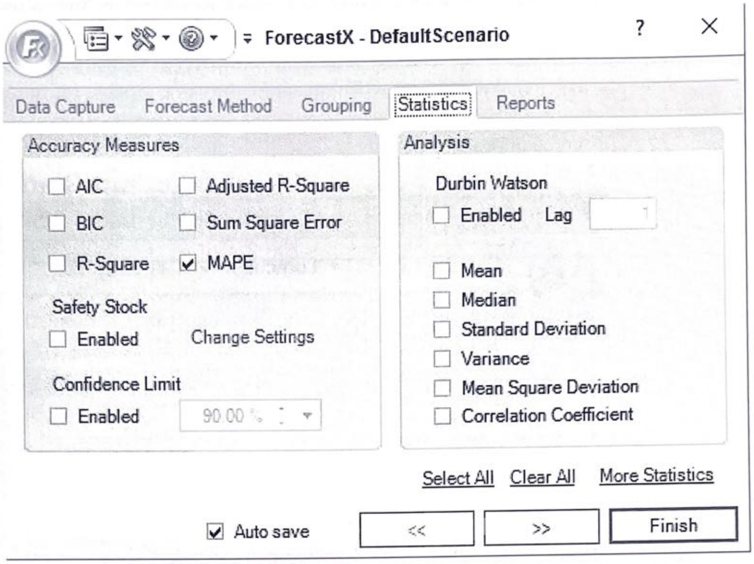

Source: John Galt Solutions

For this example, the only thing we want is the MAPE, which you see is checked above. Many other measures are available, so you might want to explore this box to familiarize yourself with the wide variety of options available.

Nest click on the "Reports" tab and the following will appear:

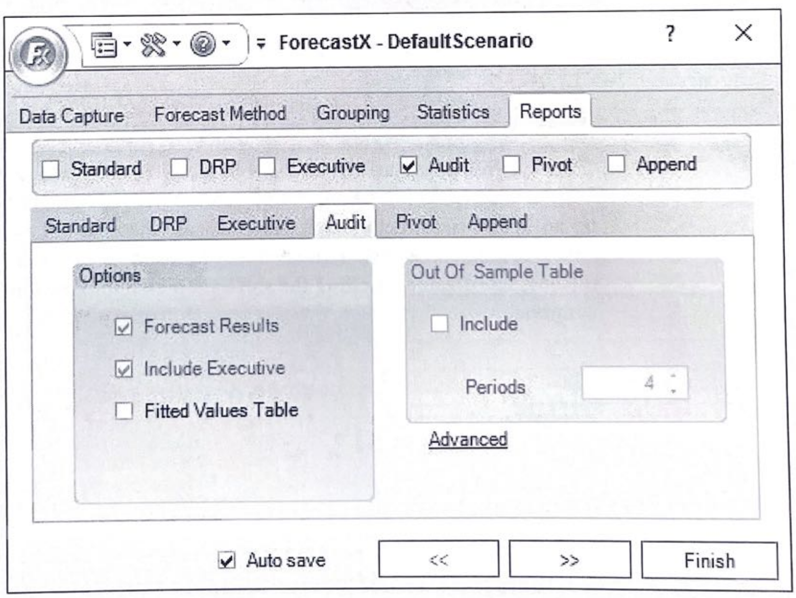

Source: John Galt Solutions

For this example, we only want what is called the "Audit" report, so that is what is checked. You will want to explore the other options. The "Standard" report is often useful.

Near the lower right of this box you see the word *Advanced* in dark blue. You may find it useful to click on *Advanced*. The following appears:

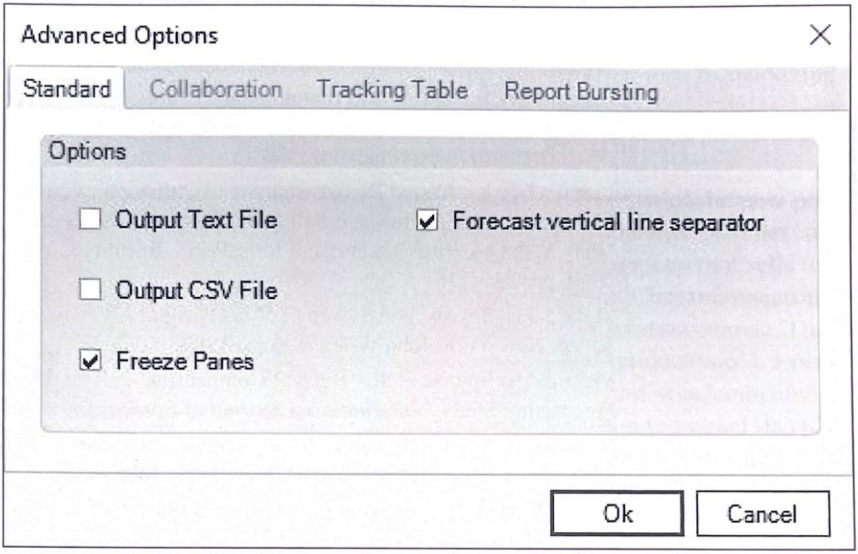

Source: John Galt Solutions

The two options most students find useful are the "Freeze Panes" and the "Forecast vertical line separator" options. You see those are checked above. Once you click on Ok, you will return to the "Reports" tab. Now click on "Finish." Your results will be calculated and displayed in a new Excel book. You may first need to click "View Output" in the following box.

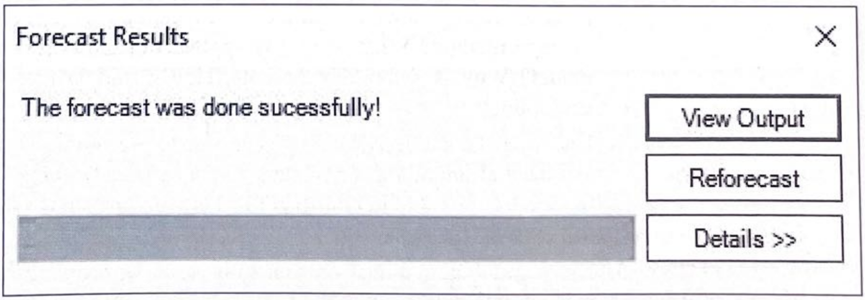

Source: John Galt Solutions

Suggested Readings

Armstrong, J. Scott, ed. *Principles of Forecasting: A Handbook for Researchers and Practitioners*. Boston: Kluwer Academic Publishers, 2001.

Bass, Frank M. "A New Product Growth Model for Consumer Durables." *Management Science* 15 (January 1969), pp. 215–27. (Web site: http://www.basseconomics.com.)

Gardner, Everette S. "Exponential Smoothing: The State of the Art." *Journal of Forecasting* 4, no. 1 (1985), pp. 1–28.

Holt, C. C. "Forecasting Seasonal and Trends by Exponentially Weighted Moving Averages." Office of Naval Research, Memorandum No. 52, 1957.

Lilien, Gary L.; and Arvind Rangaswamy. *Marketing Engineering: Computer-Assisted Marketing Analysis and Planning*. 2nd ed. Upper Saddle River, NJ: Prentice-Hall, 2003. (Web site: http://www.mktgeng.com.)

Lilien, Gary L.; Arvind Rangaswamy; and Christophe Van den Bulte, "Diffusion Models: Managerial Applications and Software." In *New-Product Diffusion Models,* edited by Vijay Mahajan, Eitan Muller, and Jerry Wind. Boston: Kluwer Academic Publishers, 2000, pp. 295–336.

Makridakis, Spyros; and Steven C. Wheelwright. *Forecasting Methods for Management.* 5th ed. New York: John Wiley & Sons, 1989.

Makridakis, Spyros, et al. "The M2-Competition: A Real-Time Judgmentally Based Forecasting Study." *International Journal of Forecasting* 9, no. 1 (April 1993), pp. 5–22.

Makridakis, Spyros; Steven C. Wheelwright; and Victor E. McGee. *Forecasting: Methods and Applications.* 2nd ed. New York: John Wiley & Sons, 1983.

Pegels, Carl C. "Exponential Forecasting: Some New Variations." *Management Science* 15, no. 12 (January 1969), pp. 311–15.

West, Douglas C. "Number of Sales Forecast Methods and Marketing Management." *Journal of Forecasting* 13, no. 4 (August 1994), pp. 395–407.

Winters, P. R. "Forecasting Sales by Exponentially Weighted Moving Averages." *Management Science* 6 (1960), pp. 324–42.

Exercises

1. Assume you were to use α values of 0.1, 0.5, and 0.9 in a separate simple exponential smoothing models. How would these different α values weight past observations of the variable to be forecast? How would you know which of these α values provided the best forecasting model? If the $\alpha = 0.9$ value provided the best forecast for your data, would this imply that you should do anything else? Does exponential smoothing place more or less weight on the most recent data when compared with the moving-average method? What weight is applied to each observation in a moving-average model? Why is smoothing (simple, Holt's, and Winters') also called *exponential* smoothing?
2. Under what conditions would you choose to use simple exponential smoothing, Holt's exponential smoothing, or Winters' exponential smoothing? Are these the only smoothing models possible to construct? If there are other possible models, suggest one that might be useful.
3. Exponential smoothing is meant to be used with time-series data when the data are made up of some or all of the basic components of level, trend, seasonality, and error. If the data series only fluctuates about an average with no trend and no seasonality, which form of smoothing would you employ? If the data include all of these

components, which form of smoothing would you employ? How should the smoothing constants be chosen?

4. The smoothing constant chosen in simple exponential smoothing determines the weight to be placed on different terms of time-series data. If the smoothing factor is high rather than low, is more or less weight placed on recent observations? If α is .3, what weight is applied to the observation four periods ago?
5. Consider the following rates offered on certificates of deposit at a large metropolitan bank during a recent year (c3p5):

Month	Rate (%)
Jan-16	1.025
Feb-16	2.047
Mar-16	2.28
Apr-16	2.65
May-16	2.714
Jun-16	2.963
Jul-16	1.575
Aug-16	2.612
Sep-16	2.985
Oct-16	2.298
Nov-16	1.454
Dec-16	2.461

Use a three-month average to forecast the rate for the following January.

6. The following inventory pattern has been observed in the Zahm Corporation over 12 months (c3p6):

Month	Inventory
Apr-16	1,544
May-16	1,913
Jun-16	2,028
Jul-16	1,178
Aug-16	1,554
Sep-16	1,910
Oct-16	1,208
Nov-16	2,467
Dec-16	2,101
Jan-17	1,662
Feb-17	2,432
Mar-17	2,443

Use both three-month and five-month moving-average models to forecast the inventory for the next January. Use mean absolute percentage error (MAPE) to evaluate these two forecasts.

7. Consider the following data on full-service restaurant sales. Calculate both the three-month and five-month moving averages for these data, and compare the forecasts by calculating the mean absolute percentage errors and MAPEs. The data are in millions of dollars. (c3p7)

Date	Sales (000,000)	Date	Sales (000,000)
Jan-02	6910	Jun-05	8549
Feb-02	6959	Jul-05	8902
Mar-02	7268	Aug-05	9035
Apr-02	7023	Sep-05	8271
May-02	7555	Oct-05	8328
Jun-02	7021	Nov-05	7987
Jul-02	7297	Dec-05	8383
Aug-02	7558	Jan-06	7532
Sep-02	6945	Feb-06	7943
Oct-02	7464	Mar-06	8685
Nov-02	7138	Apr-06	8502
Dec-02	7355	May-06	8977
Jan-03	6854	Jun-06	8716
Feb-03	6699	Jul-06	8978
Mar-03	7324	Aug-06	9548
Apr-03	7514	Sep-06	8675
May-03	7898	Oct-06	9032
Jun-03	7814	Nov-06	9005
Jul-03	8049	Dec-06	8921
Aug-03	8322	Jan-07	8688
Sep-03	7730	Feb-07	8640
Oct-03	8049	Mar-07	9592
Nov-03	7449	Apr-07	9332
Dec-03	7774	May-07	9976
Jan-04	6998	Jun-07	9460
Feb-04	7275	Jul-07	10071
Mar-04	8177	Aug-07	10517
Apr-04	8143	Sep-07	9539
May-04	8364	Oct-07	9850
Jun-04	8292	Nov-07	9227
Jul-04	8689	Dec-07	9699
Aug-04	8661	Jan-08	9147
Sep-04	8080	Feb-08	9114
Oct-04	8264	Mar-08	9972
Nov-04	7822	Apr-08	9825
Dec-04	8352	May-08	10423
Jan-05	7507	Jun-08	10203
Feb-05	7341	Jul-08	10458
Mar-05	8243	Aug-08	10541
Apr-05	8269	Sep-08	9844
May-05	8615	Oct-08	10455

(continued on next page)

(continued)

Date	Sales (000,000)	Date	Sales (000,000)
Nov-08	9715	Jun-12	12685
Dec-08	10338	Jul-12	12873
Jan-09	9583	Aug-12	13357
Feb-09	9515	Sep-12	11743
Mar-09	10385	Oct-12	12129
Apr-09	10571	Nov-12	12003
May-09	10792	Dec-12	12794
Jun-09	10553	Jan-13	11811
Jul-09	11083	Feb-13	11523
Aug-09	10939	Mar-13	12957
Sep-09	10297	Apr-13	12423
Oct-09	11056	May-13	13741
Nov-09	10229	Jun-13	13250
Dec-09	10703	Jul-13	13673
Jan-10	10092	Aug-13	14329
Feb-10	10532	Sep-13	12465
Mar-10	11464	Oct-13	13026
Apr-10	11240	Nov-13	12606
May-10	11393	Dec-13	13281
Jun-10	11332	Jan-14	12953
Jul-10	11752	Feb-14	12926
Aug-10	11581	Mar-14	13709
Sep-10	11257	Apr-14	13324
Oct-10	11447	May-14	14042
Nov-10	10742	Jun-14	13669
Dec-10	11372	Jul-14	14572
Jan-11	10726	Aug-14	14149
Feb-11	10691	Sep-14	13268
Mar-11	11919	Oct-14	13918
Apr-11	11312	Nov-14	12992
May-11	12002	Dec-14	14312
Jun-11	12191	Jan-15	13202
Jul-11	12374	Feb-15	13260
Aug-11	12797	Mar-15	14359
Sep-11	11292	Apr-15	14368
Oct-11	11523	May-15	14687
Nov-11	11259	Jun-15	14445
Dec-11	12596	Jul-15	15142
Jan-12	11520	Aug-15	14905
Feb-12	11414	Sep-15	13982
Mar-12	12696	Oct-15	14575
Apr-12	12140	Nov-15	13838
May-12	12857	Dec-15	15478

8. Forecasters at Siegfried Corporation are using simple exponential smoothing to forecast the sales of its major product. They are trying to decide what smoothing constant will give the best results. They have tried a number of smoothing constants with the following results:

Smoothing Constant	MAPE
0.10	12.5%
0.15	9.7%
0.20	13.6%
0.25	14.1%

Which smoothing constant appears best from these results? Why? Could you perhaps get even better results, given these outcomes? How would you go about improving the MAPE?

9. The number of tons of brake assemblies received at an auto parts distribution center last month was 670. The forecast tonnage was 720 for last month. The company uses a simple exponential smoothing model with a smoothing constant of 0.6 to develop its forecasts. What will be the company's forecast for the next month?

10. The number of service calls received at LaFortune Electric during four months is shown in the following table (c3p10):

Month	Number of Service Calls
April	19
May	31
June	27
July	29

Forecast the number of service calls in August by using a simple exponential smoothing model with a smoothing constant of 0.1. (Assume the forecast for April was 21.)

11. *a.* Plot the data presented in Exercise 7 to examine the possible existence of trend and seasonality in the data. (c3p11)

    b. Prepare three separate exponential smoothing models to forecast the full-service restaurant sales data using the monthly data.

    1. A simple smoothing model
    2. Holt's model
    3. Winters' model

    c. Examine the accuracy of each model by calculating the mean absolute percentage error for each during the historical period. Explain carefully what characteristics of the original data led one of these models to have the lowest MAPE.

12. The data in the table below represent warehouse club and superstore sales in the eastern and central United States on a monthly basis. The data are in millions of dollars. (c3p12)

Date	Sales	Date	Sales
Jan-02	2,580	Jan-06	4,758
Feb-02	2,616	Feb-06	4,914
Mar-02	2,838	Mar-06	5,431
Apr-02	2,985	Apr-06	5,474
May-02	3,258	May-06	6,124
Jun-02	3,107	Jun-06	6,027
Jul-02	3,097	Jul-06	5,914
Aug-02	3,288	Aug-06	6,244
Sep-02	3,077	Sep-06	5,808
Oct-02	3,429	Oct-06	6,373
Nov-02	4,011	Nov-06	6,994
Dec-02	5,739	Dec-06	9,018
Jan-03	2,877	Jan-07	5,694
Feb-03	2,885	Feb-07	5,431
Mar-03	3,259	Mar-07	6,240
Apr-03	3,454	Apr-07	6,101
May-03	3,771	May-07	6,849
Jun-03	3,667	Jun-07	6,694
Jul-03	3,743	Jul-07	6,815
Aug-03	3,792	Aug-07	6,948
Sep-03	3,699	Sep-07	6,450
Oct-03	4,082	Oct-07	7,190
Nov-03	4,727	Nov-07	7,738
Dec-03	6,672	Dec-07	9,769
Jan-04	3,560	Jan-08	6,665
Feb-04	3,575	Feb-08	6,400
Mar-04	4,220	Mar-08	7,277
Apr-04	4,282	Apr-08	7,584
May-04	4,594	May-08	8,169
Jun-04	4,691	Jun-08	8,179
Jul-04	4,629	Jul-08	8,118
Aug-04	4,795	Aug-08	8,284
Sep-04	4,632	Sep-08	7,962
Oct-04	5,067	Oct-08	8,636
Nov-04	5,746	Nov-08	9,433
Dec-04	7,965	Dec-08	11,786
Jan-05	4,317	Jan-09	8,082
Feb-05	4,118	Feb-09	7,761
Mar-05	4,855	Mar-09	8,994
Apr-05	4,999	Apr-09	8,803
May-05	5,343	May-09	9,712
Jun-05	5,392	Jun-09	9,843
Jul-05	5,274	Jul-09	9,769
Aug-05	5,435	Aug-09	9,944
Sep-05	5,217	Sep-09	9,582
Oct-05	5,460	Oct-09	10,209
Nov-05	6,288	Nov-09	11,115
Dec-05	8,403	Dec-09	14,995

(continued on next page)

(continued)

Date	Sales	Date	Sales
Jan-10	9,183	Nov-13	20,336
Feb-10	9,478	Dec-13	24,665
Mar-10	10,751	Jan-14	17,686
Apr-10	10,518	Feb-14	17,908
May-10	11,349	Mar-14	18,691
Jun-10	11,728	Apr-14	19,030
Jul-10	11,590	May-14	20,623
Aug-10	11,871	Jun-14	19,596
Sep-10	11,336	Jul-14	20,122
Oct-10	11,986	Aug-14	20,029
Nov-10	13,130	Sep-14	18,669
Dec-10	16,694	Oct-14	20,518
Jan-11	11,195	Nov-14	21,967
Feb-11	10,919	Dec-14	27,584
Mar-11	12,389	Jan-15	19,315
Apr-11	12,619	Feb-15	19,186
May-11	13,489	Mar-15	21,211
Jun-11	13,620	Apr-15	20,985
Jul-11	13,438	May-15	22,385
Aug-11	14,084	Jun-15	22,223
Sep-11	13,172	Jul-15	22,602
Oct-11	14,040	Aug-15	22,456
Nov-11	15,759	Sep-15	21,418
Dec-11	19,992	Oct-15	23,092
Jan-12	13,162	Nov-15	24,598
Feb-12	13,394	Dec-15	30,706
Mar-12	15,285	Jan-16	21,692
Apr-12	14,467	Feb-16	21,699
May-12	16,086	Mar-16	23,402
Jun-12	16,027	Apr-16	24,046
Jul-12	15,622	May-16	24,881
Aug-12	16,360	Jun-16	24,602
Sep-12	14,714	Jul-16	24,631
Oct-12	15,894	Aug-16	24,831
Nov-12	18,152	Sep-16	23,603
Dec-12	22,089	Oct-16	24,608
Jan-13	15,161	Nov-16	26,705
Feb-13	15,342	Dec-16	34,023
Mar-13	16,997	Jan-17	23,837
Apr-13	16,623	Feb-17	23,438
May-13	18,064	Mar-17	26,305
Jun-13	17,605	Apr-17	25,429
Jul-13	17,746	May-17	27,152
Aug-13	18,907	Jun-17	27,218
Sep-13	16,735	Jul-17	26,722
Oct-13	18,146		

a. Prepare a time-series plot of the data, and visually inspect that plot to determine the characteristics you see in this series.

b. Use an exponential smoothing model to develop a forecast of sales for the next 12 months, and explain why you selected that model. Plot the actual and forecast values. Determine the MAPE for your model during the historical period.

13. The data in the table below are for retail sales in book stores by quarter. (c3p13)

U.S. Retail Book Sales (in Millions of Dollars, NSA)

Date	Sales	Date	Sales
Mar-02	1,866	Mar-09	3,480
Jun-02	1,666	Jun-09	2,943
Sep-02	2,351	Sep-09	3,654
Dec-02	2,455	Dec-09	4,108
Mar-03	2,169	Mar-10	3,628
Jun-03	1,815	Jun-10	3,203
Sep-03	2,498	Sep-10	4,051
Dec-03	2,637	Dec-10	4,010
Mar-04	2,326	Mar-11	3,719
Jun-04	2,020	Jun-11	3,084
Sep-04	2,858	Sep-11	4,234
Dec-04	2,915	Dec-11	4,073
Mar-05	2,725	Mar-12	3,983
Jun-05	2,283	Jun-12	3,132
Sep-05	3,134	Sep-12	4,328
Dec-05	3,066	Dec-12	4,007
Mar-06	2,876	Mar-13	3,969
Jun-06	2,445	Jun-13	3,257
Sep-06	3,190	Sep-13	4,824
Dec-06	3,407	Dec-13	4,129
Mar-07	3,197	Mar-14	4,298
Jun-07	2,575	Jun-14	3,312
Sep-07	3,290	Sep-14	4,811
Dec-07	3,693	Dec-14	4,336
Mar-08	3,273	Mar-15	4,261
Jun-08	2,713	Jun-15	3,278
Sep-08	3,514	Sep-15	4,991
Dec-08	3,794	Dec-15	4,447

a. Plot these data and examine the plot. Does this view of the data suggest a particular smoothing model? Do the data appear to be seasonal? Explain.

b. Use an exponential smoothing method to forecast the next four quarters. Plot the actual and forecast values.

14. Monthly data from March 2014 through September 2017 are provided below for the number of lunches served in public schools. You are charged with making a 12-month forecast of the meals to be served. Begin by plotting the data and examining it for the patterns of trend and seasonality. Choose an appropriate model for the data and forecast for the next 12 months. (c3p14)

Month	Meals Served	Month	Meals Served
Mar-14	108,371,749	Jan-16	98,887,496
Apr-14	99,199,094	Feb-16	96,477,065
May-14	92,195,689	Mar-16	114,094,756
Jun-14	81,447,374	Apr-16	96,093,092
Jul-14	72,792,981	May-16	107,527,897
Aug-14	78,931,911	Jun-16	87,135,336
Sep-14	89,982,843	Jul-16	72,397,374
Oct-14	96,761,533	Aug-16	88,657,480
Nov-14	92,772,827	Sep-16	94,566,627
Dec-14	83,103,478	Oct-16	106,889,806
Jan-15	93,109,115	Nov-16	97,638,605
Feb-15	93,267,674	Dec-16	83,280,944
Mar-15	105,290,897	Jan-17	102,522,133
Apr-15	103,625,467	Feb-17	95,537,211
May-15	100,549,323	Mar-17	111,462,237
Jun-15	85,155,854	Apr-17	103,542,365
Jul-15	71,406,448	May-17	111,242,080
Aug-15	85,623,392	Jun-17	85,765,747
Sep-15	94,828,432	Jul-17	78,943,762
Oct-15	97,917,922	Aug-17	89,965,185
Nov-15	95,753,418	Sep-17	92,934,809
Dec-15	83,145,194		

15. Describe what is meant by the term *moving average*? When would a moving average be an appropriate forecast method?
16. How are simple moving averages models different from exponential smoothing models?
17. Why is the term *exponential* used when describing exponential smoothing forecast models?
18. For what type of data pattern would a simple exponential smoothing model be good as a forecast method?
19. When is a Holt's exponential smoothing model most appropriate?
20. What data pattern would suggest the use of a Winters' exponential smoothing model?
21. What are some methods that might be useful to forecast "new products" for which there are few historical observations?
22. What is an "event model?" Give some examples of when such a model might be useful.