Forecasting

Chapter 3

Introduction

Current factors and conditions

Past experience in a similar situation

Accounting. New product/process cost estimates, profit projections, cash management.

Finance. Equipment/equipment replacement needs, timing and amount of funding/borrowing needs.

Human resources. Hiring activities, including recruitment, interviewing, training, layoff planning, including outplacement, counseling.

Marketing. Pricing and promotion, e-business strategies, global competition strategies.

MIS. New/revised information systems, Internet services.

Operations. Schedules, capacity planning, work assignments and workloads, inventory planning, make-or-buy decisions, outsourcing, project management.

Product/service design. Revision of current features, design of new products or services.

to help managers plan the system

to help managers plan the use of the system

Features Common to All Forecasts

Forecasting techniques generally assume that the same underlying causal system that existed in the past will continue to exist in the future.

Forecasts are rarely perfect; actual results usually differ from predicted values. No one can predict precisely how an often large number of related factors will impinge upon the variable in question; this, along with the presence of randomness, precludes a perfect forecast. Allowances should be made for forecast errors.

Forecasts for groups of items tend to be more accurate than forecasts for individual items because forecasting errors among items in a group usually have a canceling effect. Opportunities for grouping may arise if parts or raw materials are used for multiple products or if a product or service is demanded by a number of independent sources.

Forecast accuracy decreases as the time period covered by the forecast—the time horizon—increases. Generally speaking, short-range forecasts must contend with fewer uncertainties than longer-range forecasts, so they tend to be more accurate.

Elements of A Good Forecast

The forecast should be timely. Usually, a certain amount of time is needed to respond to the information contained in a forecast. The forecasting horizon must cover the time necessary to implement possible changes.

The forecast should be accurate, and the degree of accuracy should be stated. This will enable users to plan for possible errors and will provide a basis for comparing alternative forecasts.

The forecast should be reliable; it should work consistently.

The forecast should be expressed in meaningful units. The choice of units depends on user needs.

The forecast should be in writing. Although this will not guarantee that all concerned are using the same information, it will at least increase the likelihood of it.

The forecasting technique should be simple to understand and use. Users often lack confidence in forecasts based on sophisticated techniques; they do not understand either the circumstances in which the techniques are appropriate or the limitations of the techniques.

The forecast should be cost-effective: The benefits should outweigh the costs.

Steps in the Forecasting Process

1. Determine the purpose of the forecast. How will it be used and when will it be needed? This step will provide an indication of the level of detail required in the forecast, the amount of resources (personnel, computer time, dollars) that can be justified, and the level of accuracy necessary.
2. Establish a time horizon. The forecast must indicate a time interval, keeping in mind that accuracy decreases as the time horizon increases.
3. Select a forecasting technique.
4. Obtain, clean, and analyze appropriate data. Obtaining the data can involve significant effort. Once obtained, the data may need to be "cleaned" to get rid of outliers and obviously incorrect data before analysis.
5. Make the forecast.
6. Monitor the forecast. A forecast has to be monitored to determine whether it is performing in a satisfactory manner. If it is not, reexamine the method, assumptions, validity of data, and so on; modify as needed; and prepare a revised forecast.

Approaches to Forecasting

Qualitative methods consist mainly of subjective inputs, which often defy precise numerical description. Qualitative techniques permit inclusion of soft information (e.g., human factors, personal opinions, hunches) in the forecasting process. Those factors are often omitted or downplayed when quantitative techniques are used because they are difficult or impossible to quantify.

Quantitative methods involve either the projection of historical data or the development of associative models that attempt to utilize causal (explanatory) variables to make a forecast. Quantitative techniques consist mainly of analyzing objective, or hard, data. They usually avoid personal biases that sometimes contaminate qualitative methods. In practice, either or both approaches might be used to develop a forecast.

Judgmental forecasts rely on analysis of subjective inputs obtained from various sources, such as consumer surveys, the sales staff, managers and executives, and panels of experts. Quite frequently, these sources provide insights that are not otherwise available.

Time-series forecasts simply attempt to project past experience into the future. These techniques use historical data with the assumption that the future will be like the past. Some models merely attempt to smooth out random variations in historical data; others attempt to identify specific patterns in the data and project or extrapolate those patterns into the future, without trying to identify causes of the patterns.

Associative models Forecasting technique that uses explanatory variables to predict future demand. use equations that consist of one or more explanatory variables that can be used to predict demand. For example, demand for paint might be related to variables such as the price per gallon and the amount spent on advertising, as well as to specific characteristics of the paint (e.g., drying time, ease of cleanup).

Forecasts Based on Judgment and Opinion

Executive opinions

Salesforce opinions

Consumer surveys

Other approaches (Delphi method, ...)

Forecast Based on Time Series Data

Trend refers to a long-term upward or downward movement in the data. Population shifts, changing incomes, and cultural changes often account for such movements.

Seasonality refers to short-term, fairly regular variations generally related to factors such as the calendar or time of day. Restaurants, supermarkets, and theaters experience weekly and even daily “seasonal” variations.

Cycles are Wavelike variations lasting more than one year. are wavelike variations of more than one year’s duration. These are often related to a variety of economic, political, and even agricultural conditions.

Irregular variations are due to unusual circumstances such as severe weather conditions, strikes, or a major change in a product or service. They do not reflect typical behavior, and their inclusion in the series can distort the overall picture. Whenever possible, these should be identified and removed from the data.

Random variations are residual variations after all other behaviors are accounted for. are residual variations that remain after all other behaviors have been accounted for.

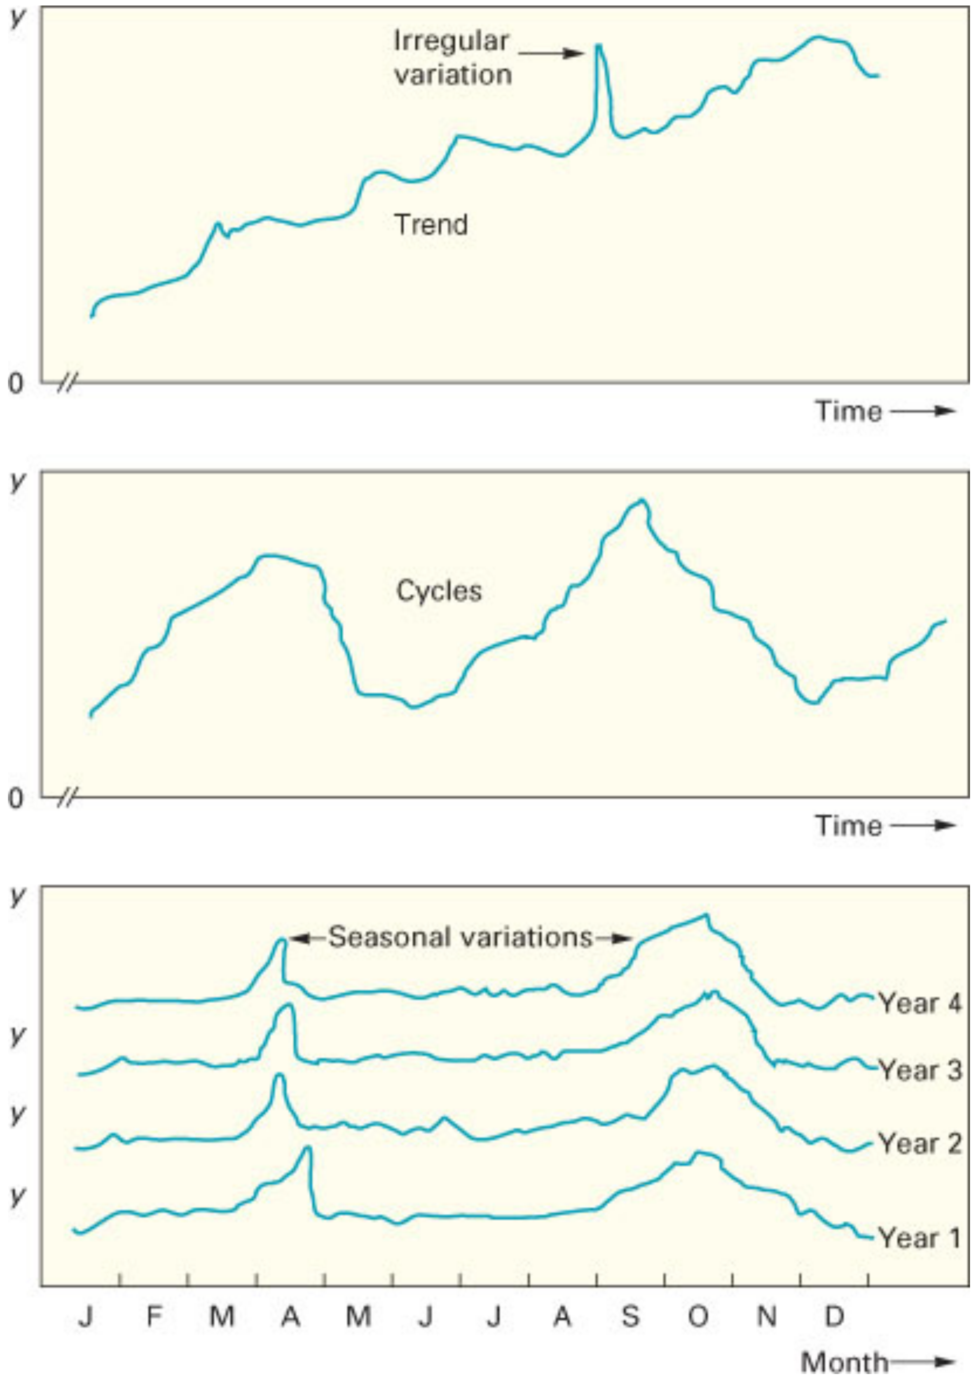
y
Irregular variation
Trend
0
Time
y
Cycles
0
Time
y
Seasonal variations
Year 4
y
Year 3
y
Year 2
y
Year 1
J F M A M J J A S O N D
Month

Naive Methods

Period	Actual	Change from Previous Value	Forecast
$t - 2$	50		
$t - 1$	53	+3	
(next) t			$53 + 3 = 56$

Techniques for Averaging

Historical data typically contain a certain amount of random variation, or *white noise,* that tends to obscure systematic movements in the data. This randomness arises from the combined influence of many—perhaps a great many—relatively unimportant factors, and it cannot be reliably predicted. Averaging techniques smooth variations in the data. Ideally, it would be desirable to completely remove any randomness from the data and leave only “real” variations, such as changes in the demand. As a practical matter, however, it is usually impossible to distinguish between these two kinds of variations, so the best one can hope for is that the small variations are random and the large variations are “real.”

Moving Average

$$F_t = \text{MA}_n = \frac{\sum_{i=1}^{n} A_{t-i}}{n} = \frac{A_{t-n} + \cdots + A_{t-2} + A_{t-1}}{n}$$

F_t = Forecast for time period t
MA_n = n period moving average
$A_t - 1$ = Actual value in period $t - 1$
n = Number of periods (data points) in the moving average

Example

Period	Demand	
1	42	
2	40	
3	43	
4	40	the 3 most recent demands
5	41	

$$F_6 = \frac{43 + 40 + 41}{3} = 41.33$$

If actual demand in period 6 turns out to be 38, the moving average forecast for period 7 would be

$$F_7 = \frac{40 + 41 + 38}{3} = 39.67$$

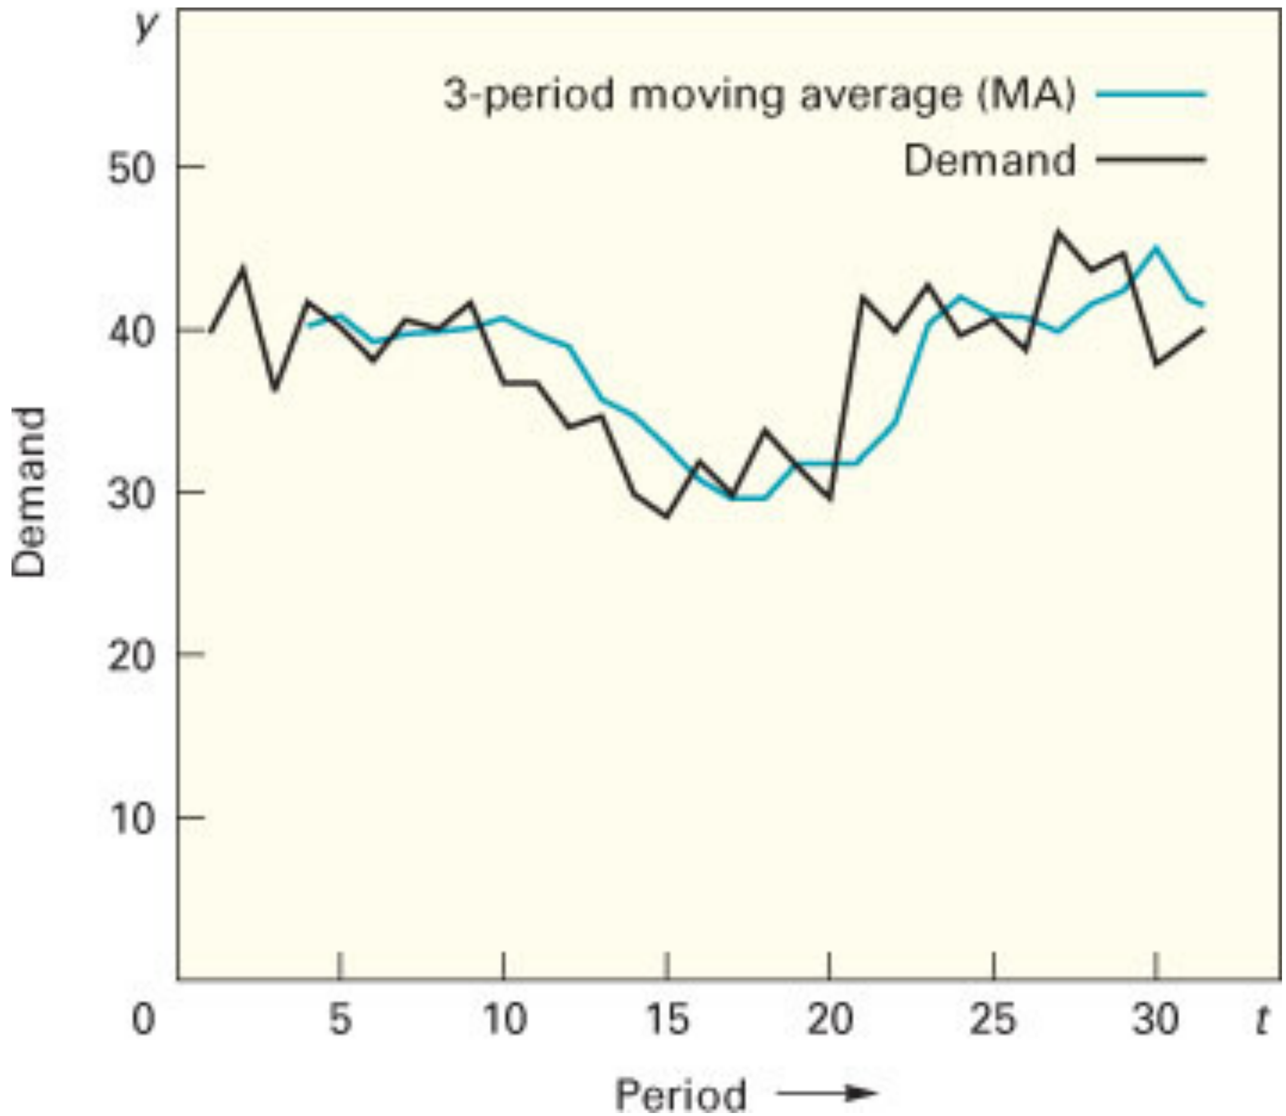

y
3-period moving average (MA)
Demand
50
40
30
20
10
0
5
10
15
20
25
30
t
Demand
Period

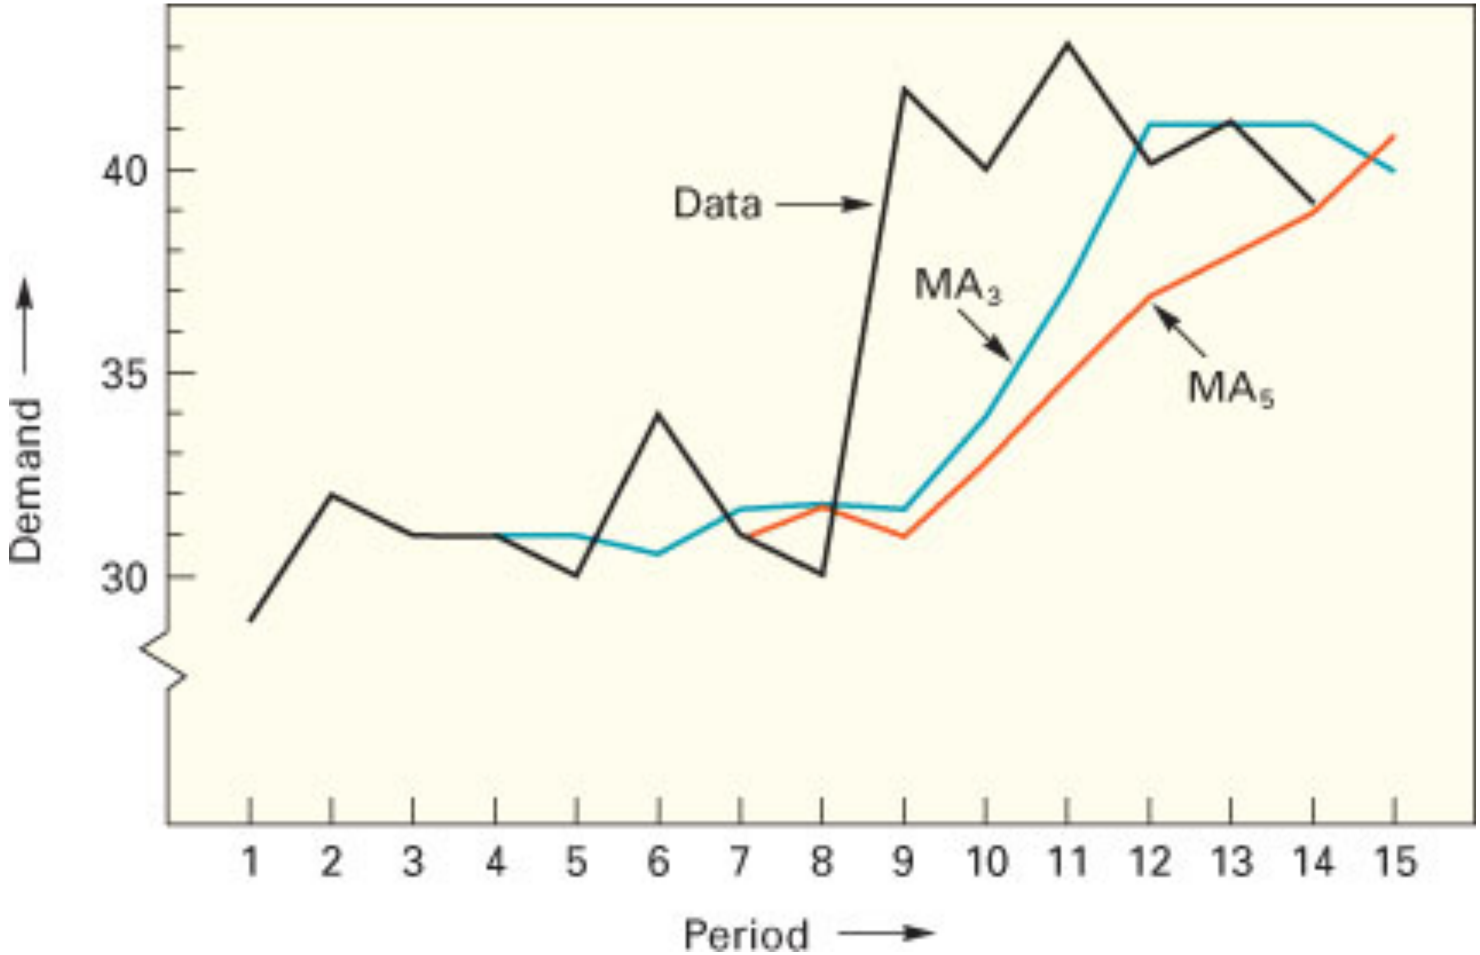

40
35
30
Demand
Data
MA3
MA5
1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
Period

Weighted Moving Average

Given the following demand data,

Compute a weighted average forecast using a weight of .40 for the most recent period, .30 for the next most recent, .20 for the next, and .10 for the next.

If the actual demand for period 6 is 39, forecast demand for period 7 using the same weights as in part a.

Period	Demand
1	42
2	40
3	43
4	40
5	41

Solution:

F_6 = .10(40) + .20(43) + .30(40) + 40(41) = 41.0
F_7 = .10(43) + .20(40) + .30(41) + .40(39) = 40.2

Exponential Smoothing

$$\text{Next forecast} = \text{Previous forecast} + \alpha(\text{Actual} - \text{Previous forecast})$$

$$F_t = F_{t-1} + \alpha(A_{t-1} - F_{t-1})$$

F_t = Forecast for period t
F_{t-1} = Forecast for the previous period (i.e., period t = 1)
α = Smoothing constant
A_{t-1} = Actual demand or sales for the previous period

The smoothing constant α represents a percentage of the forecast error. Each new forecast is equal to the previous forecast plus a percentage of the previous error. For example, suppose the previous forecast was 42 units, actual demand was 40 units, and α = .10. The new forecast would be computed as follows:

$$F_t = 42 + .10(40 - 42) = 41.8$$

Then, if the actual demand turns out to be 43, the next forecast would be

$$F_t = 41.8 + .10(43 - 41.8) = 41.92$$

Example

The following table illustrates two series of forecasts for a data set, and the resulting (Actual – Forecast) = Error, for each period. One forecast uses α = .10 and one uses α = .40. The following figure plots the actual data and both sets of forecasts.

		$\alpha = .10$		$\alpha = .40$	
Period (t)	Actual Demand	Forecast	Error	Forecast	Error
1	42	—	—	—	—
2	40	42	–2	42	–2
3	43	41.8	1.2	41.2	1.8
4	40	41.92	–1.92	41.92	–1.92
5	41	41.73	–0.73	41.15	–0.15
6	39	41.66	–2.66	41.09	–2.09
7	46	41.39	4.61	40.25	5.75
8	44	41.85	2.15	42.55	1.45
9	45	42.07	2.93	43.13	1.87
10	38	42.35	–4.35	43.88	–5.88
11	40	41.92	–1.92	41.53	–1.53
12		41.73		40.92	

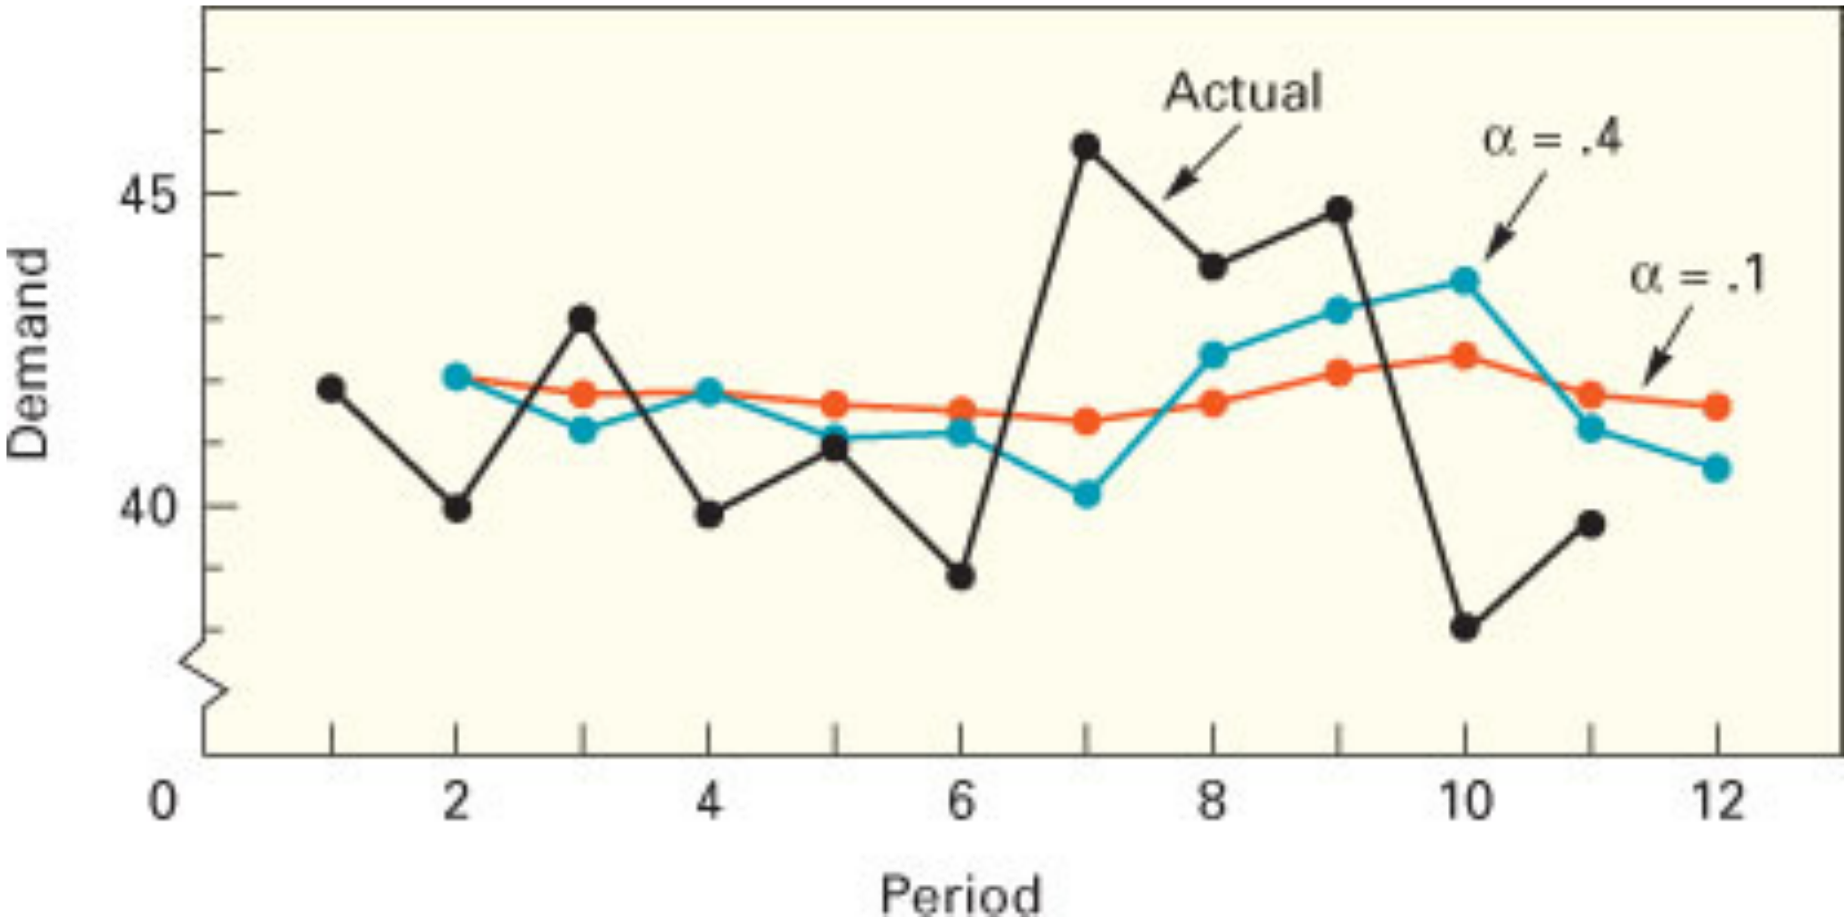
Actual
α = .4
α = .1
Demand
45
40
0
2
4
6
8
10
12
Period

Techniques For Trends

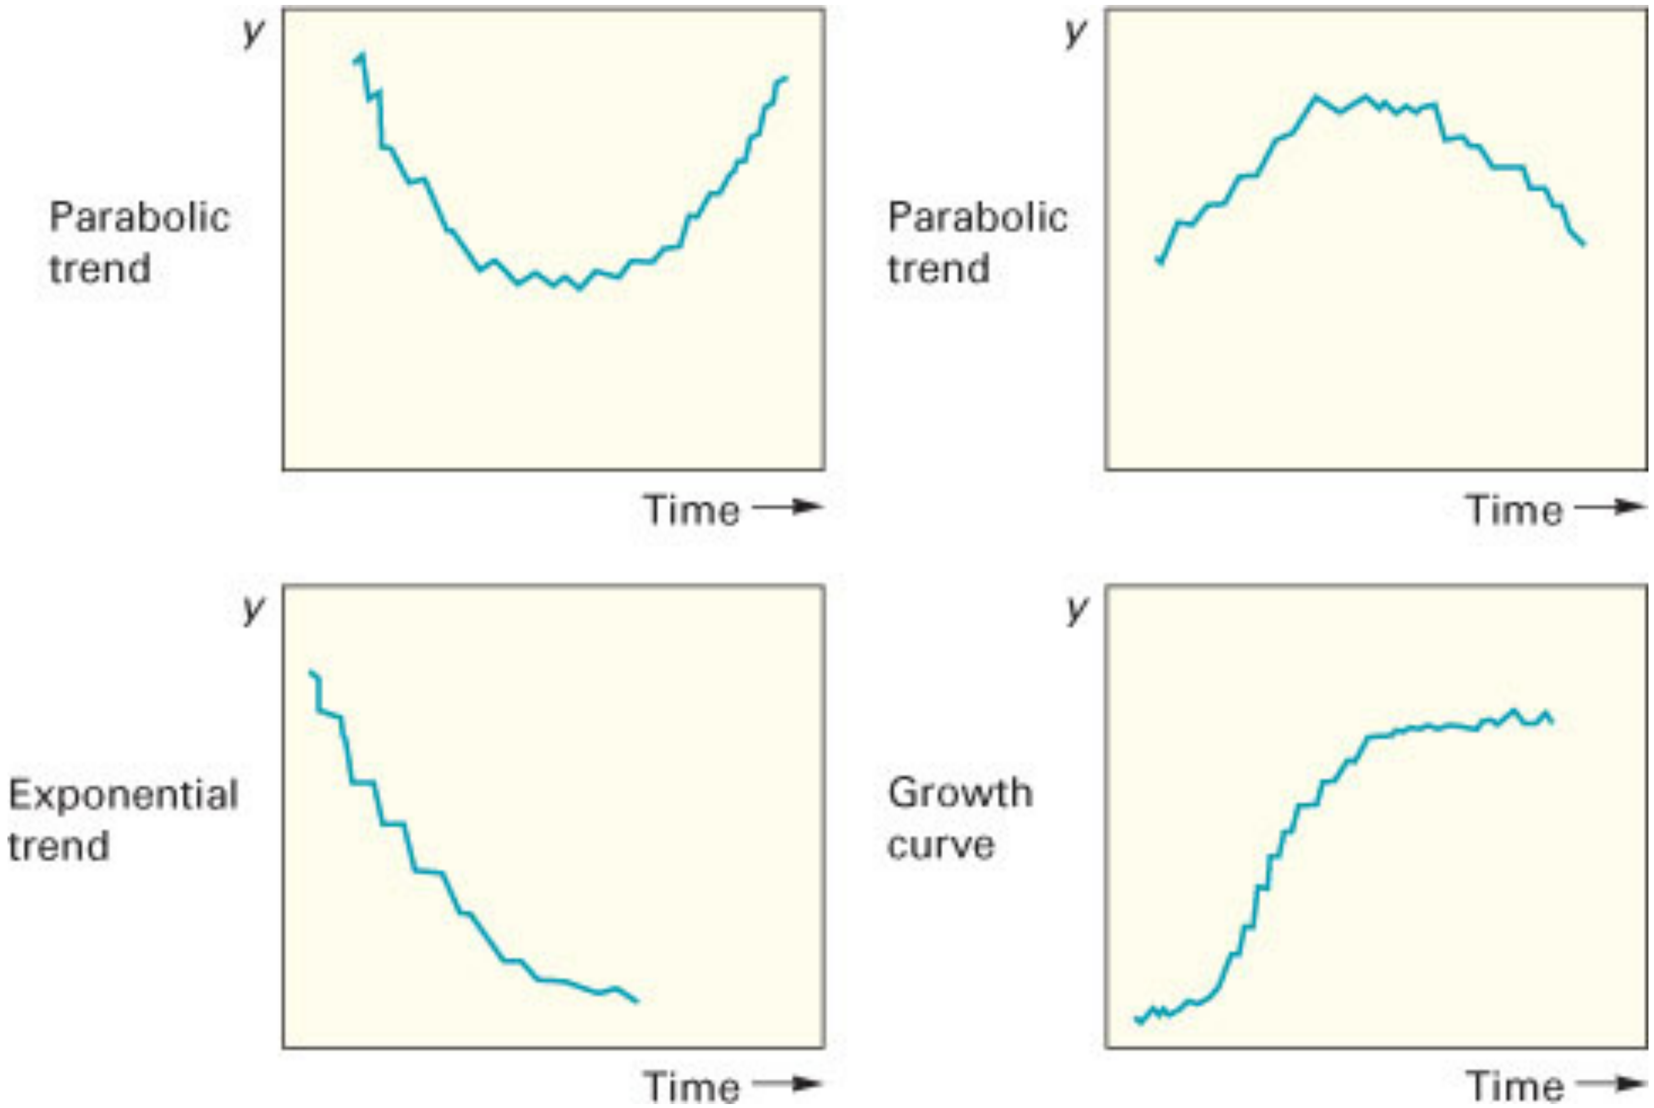

Trend Equation

$$F_t = a + bt$$

F_t = Forecast for period t
a = Value of F_t at t = 0
b = Slope of the line
t = Specified number of time periods from t = 0

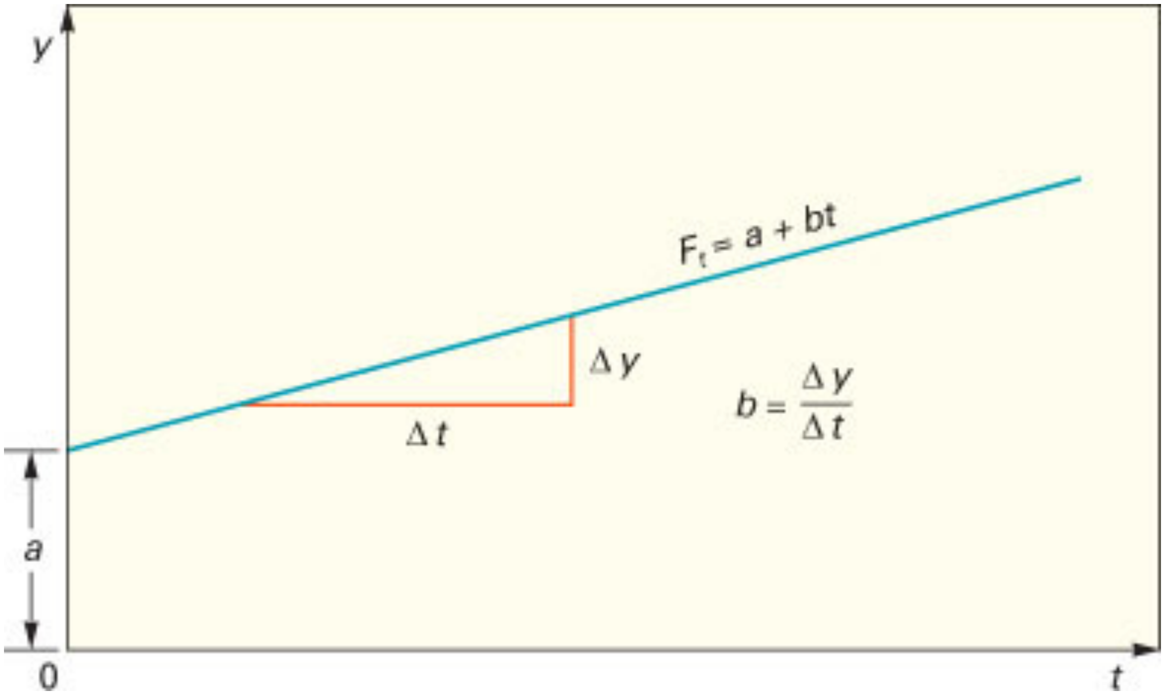

Example

For example, consider the trend equation $F_t = 45 + 5t$. The value of F_t when $t = 0$ is 45, and the slope of the line is 5, which means that, on the average, the value of F_t will increase by five units for each time period. If $t = 10$, the forecast, Ft, is $45 + 5(10) = 95$ units. The equation can be plotted by finding two points on the line. One can be found by substituting some value of t into the equation (e.g., $t = 10$) and then solving for F_t. The other point is a (i.e., F_t at $t = 0$). Plotting those two points and drawing a line through them yields a graph of the linear trend line.

The coefficients of the line, a and b, can be computed from historical data using the following two equations:

$$b = \frac{n\Sigma ty - \Sigma t \Sigma y}{n\Sigma t^2 - (\Sigma t)^2}$$

$$a = \frac{\Sigma y - b\Sigma t}{n} \text{ or } \bar{y} - b\bar{t}$$

n = Number of periods
y = Value of the time series

n	$\sum t$	$\sum t^2$
1 . . .	1	1
2 . . .	3	5
3 . . .	6	14
4 . . .	10	30
5 . . .	15	55
6 . . .	21	91
7 . . .	28	140
8 . . .	36	204
9 . . .	45	285
10 . . .	55	385
11 . . .	66	506
12 . . .	78	650
13 . . .	91	819
14 . . .	105	1,015
15 . . .	120	1,240
16 . . .	136	1,496
17 . . .	153	1,785
18 . . .	171	2,109
19 . . .	190	2,470
20 . . .	210	2,870

Example

Cell phone sales for a California-based firm over the last 10 weeks are shown in the table below. Plot the data, and visually check to see if a linear trend line would be appropriate. Then determine the equation of the trend line, and predict sales for weeks 11 and 12.

Week	Unit Sales
1	700
2	724
3	720
4	728
5	740
6	742
7	758
8	750
9	770
10	775

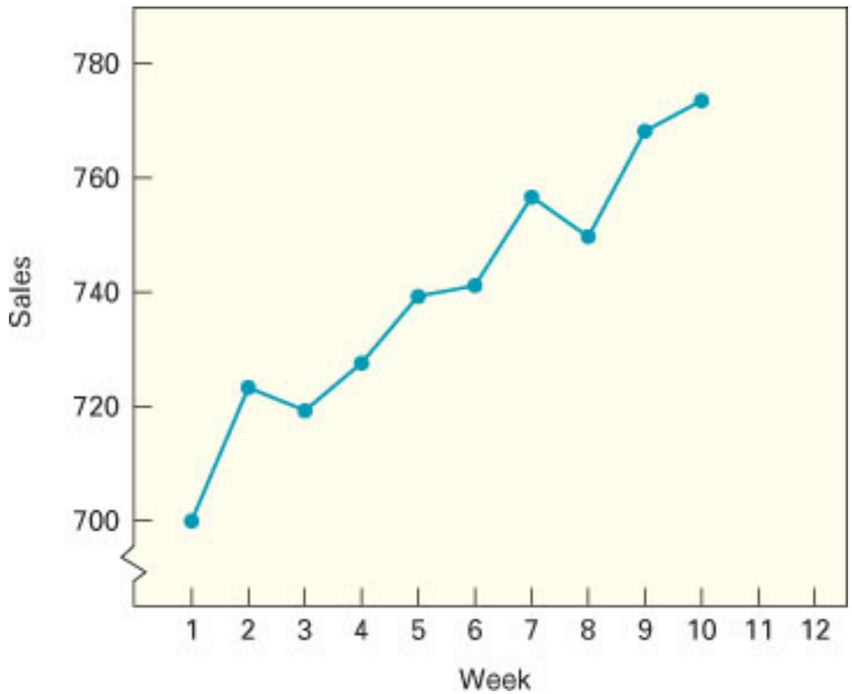

Week (t)	y	ty
1	700	700
2	724	1,448
3	720	2,160
4	728	2,912
5	740	3,700
6	742	4,452
7	758	5,306
8	750	6,000
9	770	6,930
10	775	7,750
	7,407	41,358

b. From table, for n = 10, Σt = 55 and Σt^2 = 385. You can compute the coefficients of the trend line:

$$b = \frac{10(41{,}358) - 55(7{,}407)}{10(385) - 55(55)} = \frac{6{,}195}{825} = 7.51$$

$$a = \frac{7{,}407 - 7.51(55)}{10} = 699.40$$

The trend line is $F_t = 699.40 + 7.51t$, where $t = 0$ for period 0.

c. Substituting values of *t* into this equation, the forecasts for the next two periods (i.e., *t* = 11 and *t* = 12) are:

$$F_{11} = 699.40 + 7.51(11) = 782.01$$

$$F_{12} = 699.40 + 7.51(12) = 789.52$$

d.

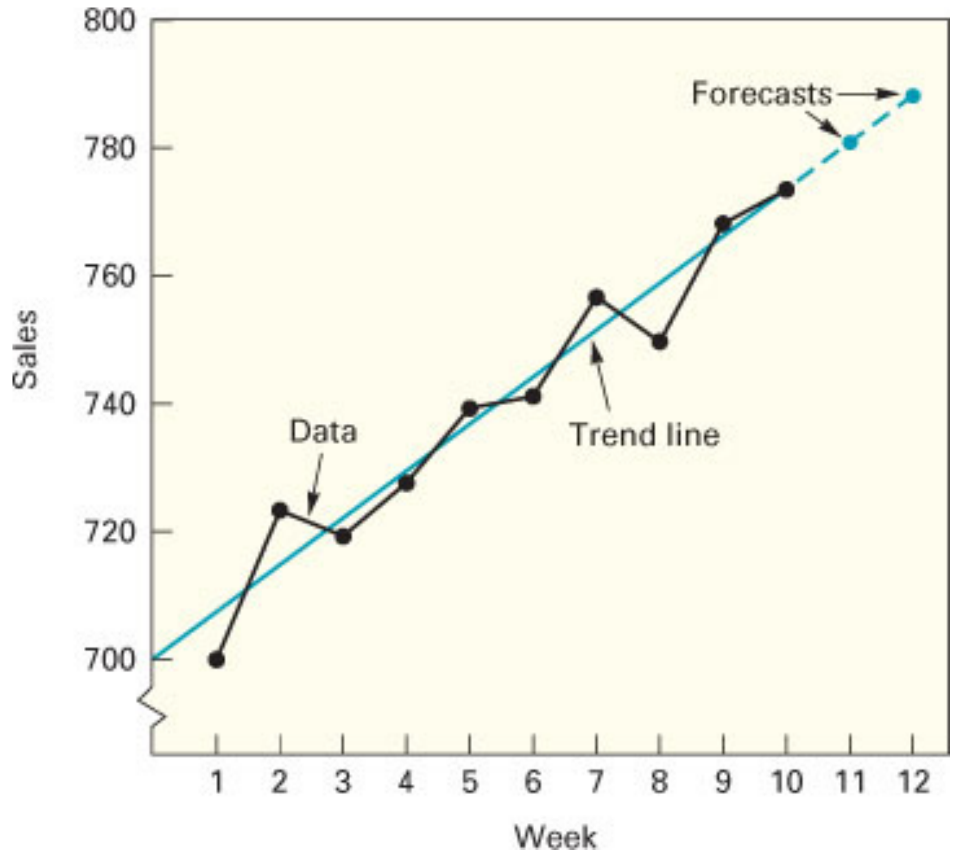

Trend-Adjusted Exponential Smoothing

$$\mathrm{TAF}_{t+1} = S_t + T_t$$

S_t = Previous forecast plus smoothed error
T_t = Current trend estimate

$$S_t = \mathrm{TAF}_t + \alpha(A_t - \mathrm{TAF}_t)$$

$$T_t = T_{t-1} + \beta(\mathrm{TAF}_t - \mathrm{TAF}_{t-1} - T_{t-1})$$

where α and β are smoothing constants. In order to use this method, one must select values of α and β (usually through trial and error) and make a starting forecast and an estimate of trend.

Example

Using the cell phone data from the previous example (where it was concluded that the data exhibited a linear trend), use trend-adjusted exponential smoothing to obtain forecasts for periods 6 through 11, with α =.40 and β = .30.

The initial estimate of trend is based on the net change of 28 for the *three changes* from period 1 to period 4, for an average of 9.33. The Excel spreadsheet is shown in table. Notice that an initial estimate of trend is estimated from the first four values, and that the starting forecast (period 5) is developed using the previous (period 4) value of 728 plus the initial trend estimate:

Starting forecast = 728 + 9.33 = 737.33

Trend Adjusted Exponential Smoothing

<Back

Clear

Model Initialization:

Period =	5		
Forecast =	737.33	α =	0.4
Trend =	9.3	$\Delta\alpha$ =	0.1
MAD =	5.39	β =	0.3
MSE =	69.41	$\Delta\beta$ =	0.1

Period	Actual	Forecast	Error
1	700		
2	724		
3	720		
4	728		
5	740		
6	742	747.728	−5.728
7	758	755.0872	2.9128
8	750	765.21536	−15.21536
9	770	768.44179	1.558208
10	775	776.55181	−1.551808
11		783.6048	
12			

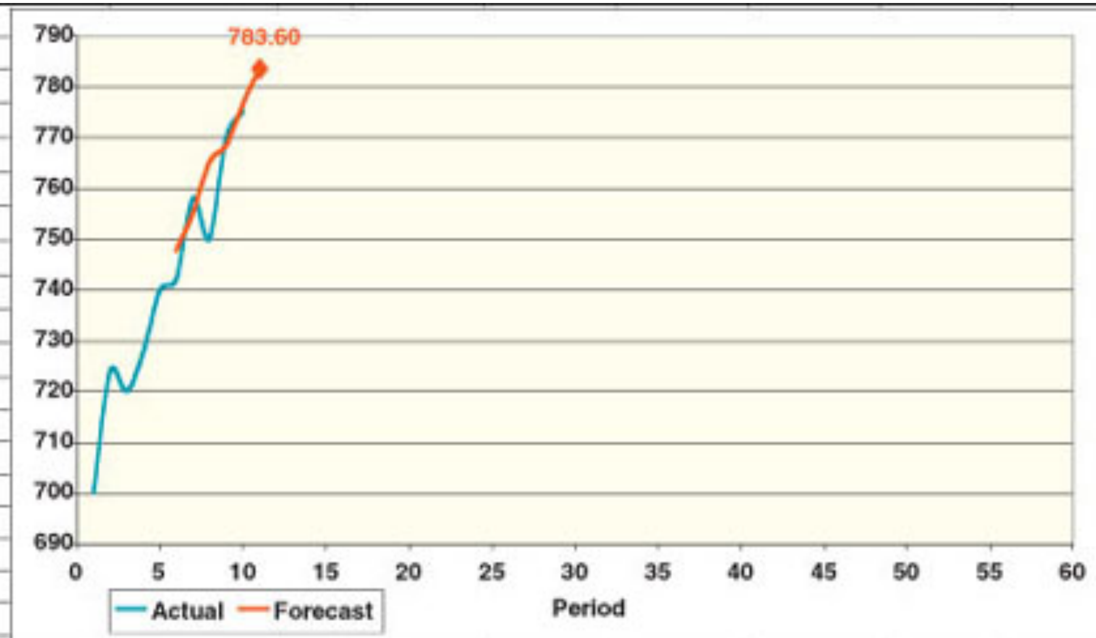

		Actual:	
Period:	11	Forecast:	783.60

Techniques for Seasonality

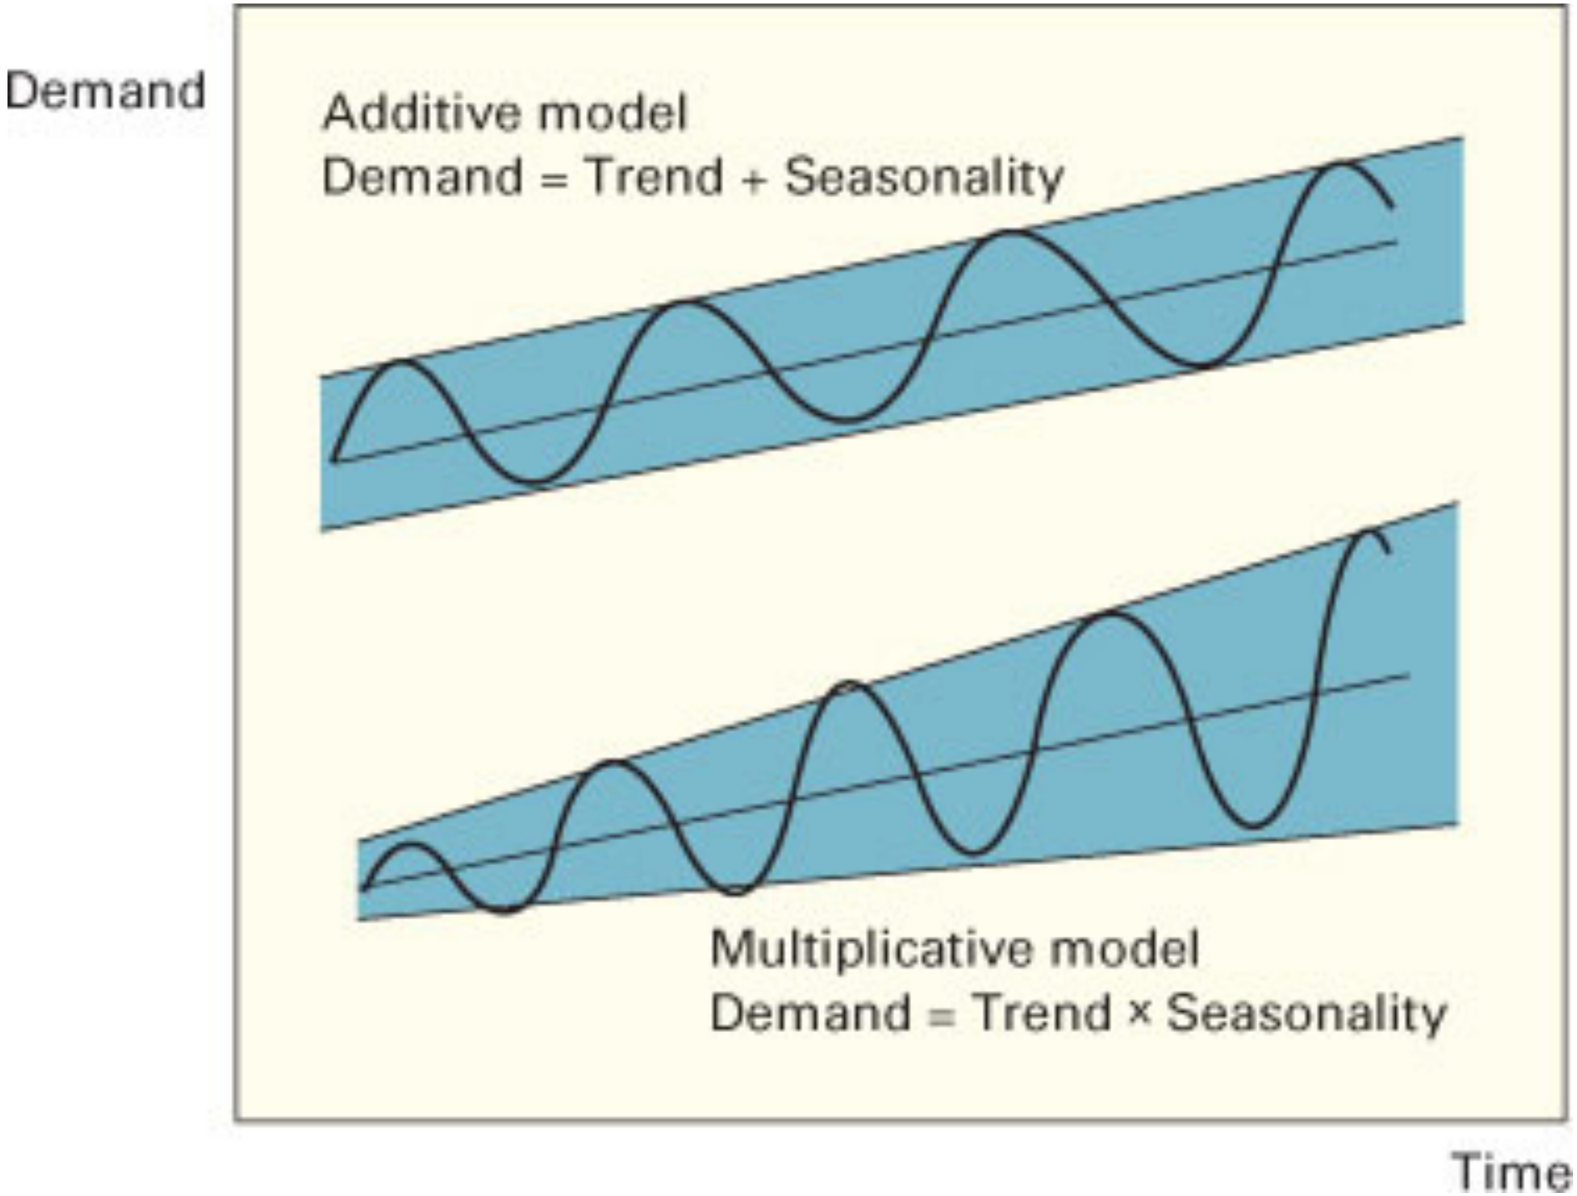

Seasonal variation: Regularly repeating movements in series values that can be tied to recurring events.

Seasonal relative: Percentage of average or trend.

Using Seasonal Relatives

1. Obtain trend estimates for desired periods using a trend equation.
2. Add seasonality to the trend estimates by multiplying (assuming a multiplicative model is appropriate) these trend estimates by the corresponding seasonal relative (e.g., multiply the November trend estimate by the November seasonal relative, multiply the December trend estimate by the December seasonal relative, and so on).

Example

A furniture manufacturer wants to predict quarterly demand for a certain loveseat for periods 15 and 16, which happen to be the third and fourth quarters of a particular year. The series consists of both trend and seasonality. The trend portion of demand is projected using the equation $F_t = 124 + 7.5t$. Quarter relatives are $Q1 = 1.20$, $Q2 = 1.10$, $Q3 = 0.75$, and $Q4 = 0.95$.

a. Use this information to deseasonalize sales for quarters 1 through 8.

b. Use this information to predict demand for periods 15 and 16.

a.

Period	Quarter	Sales	÷	Quarter Relative	=	Deseasonalized Sales
1	1	132	÷	1.20	=	110.00
2	2	140	÷	1.10	=	127.27
3	3	146	÷	0.75	=	194.67
4	4	153	÷	0.95	=	161.05
5	1	160	÷	1.20	=	133.33
6	2	168	÷	1.10	=	152.73
7	3	176	÷	0.75	=	234.67
8	4	185	÷	0.95	=	194.74

b. The trend values at t = 15 and t = 16 are:

$$F_{15} = 124 + 7.5(15) = 236.5$$
$$F_{16} = 124 + 7.5(16) = 244.0$$

Multiplying the trend value by the appropriate quarter relative yields a forecast that includes both trend and seasonality. Given that t = 15 is a second quarter and t = 16 is a third quarter, the forecasts are

Period 15: 236.5(0.75) = 177.38
Period 16: 244.0(0.95) = 231.80

Computing Seasonal Relatives

A commonly used method for representing the trend portion of a time series involves a centered moving average. Computations and the resulting values are the same as those for a moving average forecast. However, the values are not projected as in a forecast; instead, they are positioned in the middle of the periods used to compute the moving average. The implication is that the average is most representative of that point in the series. For example, assume the following time-series data:

Period	Demand	Three-Period Centered Average	
1	40		
2	46	42.67	$Average = \frac{40 + 46 + 42}{3} = 42.67$
3	42		

The three-period average is 42.67. As a centered average, it is positioned at period 2; the average is most representative of the series at that point.

The ratio of demand at period 2 to this centered average at period 2 is an estimate of the seasonal relative at that point. Because the ratio is 46/42.67 = 1.08 the series is about 8 percent above average at that point.

Example

Day	Volume	Moving Total	Centered MA_7	Volume/MA
Tues	67			
Wed	75			
Thur	82			
Fri	98		71.86	98/71.86 = 1.36 (Friday)
Sat	90		70.86	90/70.86 = 1.27
Sun	36		70.57	36/70.57 = 0.51
Mon	55	503 ÷ 7 =	71.00	55/71.00 = 0.77
Tues	60	496 ÷ 7 =	71.14	60/71.14 = 0.84 (Tuesday)
Wed	73	494 etc.	70.57	73/70.57 = 1.03
Thur	85	497	71.14	85/71.14 = 1.19
Fri	99	498	70.71	99/70.71 = 1.40 (Friday)
Sat	86	494	71.29	86/71.29 = 1.21
Sun	40	498	71.71	40/71.71 = 0.56
Mon	52	495	72.00	52/72.00 = 0.72
Tues	64	499	71.57	64/71.57 = 0.89 (Tuesday)
Wed	76	502	71.86	76/71.86 = 1.06
Thur	87	504	72.43	87/72.43 = 1.20
Fri	96	501	72.14	96/72.14 = 1.33 (Friday)
Sat	88	503		
Sun	44	507		
Mon	50	505		

The estimated Friday relative is (1.36 + 1.40 + 1.33)/3 = 1.36. Relatives for other days can be computed in a similar manner. For example, the estimated Tuesday relative is (0.84 + 0.89)/2 = 0.87.

The number of periods needed in a centered moving average is equal to the number of “seasons” involved. For example, with monthly data, a 12-period moving average is needed. When the number of periods is even, one additional step is needed because the middle of an even set falls between two periods. The additional step requires taking a centered two-period moving average of the even-numbered centered moving average, which results in averages that “line up” with data points and, hence, permit determination of seasonal ratios.

A centered moving average is used to obtain representative values because by virtue of its centered position—it "looks forward" and "looks backward"—it is able to closely follow data movements whether they involve trends, cycles, or random variability alone.

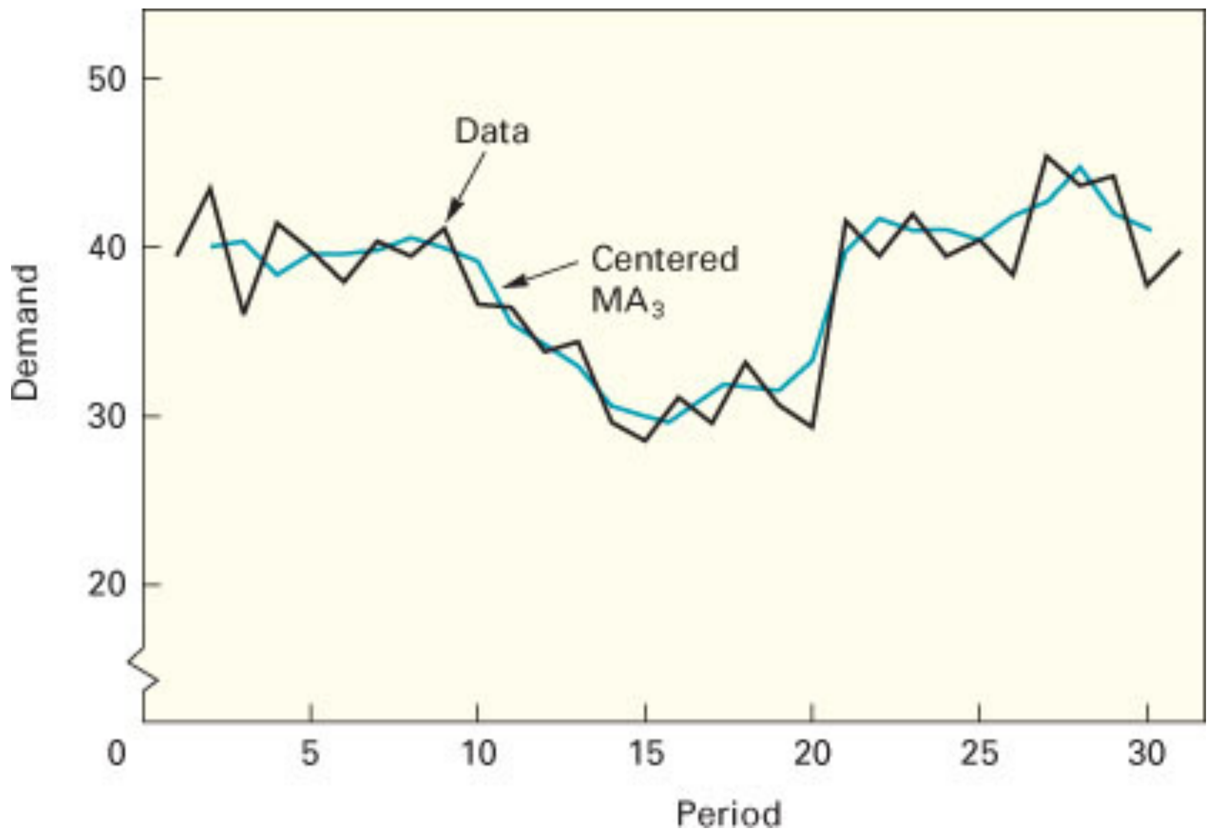

Techniques for Cycles

Cycles are up-and-down movements similar to seasonal variations but of longer duration—say, two to six years between peaks. When cycles occur in time-series data, their frequent irregularity makes it difficult or impossible to project them from past data because turning points are difficult to identify. A short moving average or a naive approach may be of some value, although both will produce forecasts that lag cyclical movements by one or several periods.

The most commonly used approach is explanatory: Search for another variable that relates to, and leads, the variable of interest. For example, the number of housing starts (i.e., permits to build houses) in a given month often is an indicator of demand a few months later for products and services directly tied to construction of new homes (landscaping; sales of washers and dryers, carpeting, and furniture; new demands for shopping, transportation, schools). Thus, if an organization is able to establish a high correlation with such a leading variable (i.e., changes in the variable precede changes in the variable of interest), it can develop an equation that describes the relationship, enabling forecasts to be made. It is important that a persistent relationship exists between the two variables. Moreover, the higher the correlation, the better the chances that the forecast will be on target.

Associative Forcasting Techniques

Pretictor variables: Variables that can be used to predict values of the variable of interest.

Regression: Technique for fitting a line to a set of points.

Least squares line: Minimizes the sum of the squared vertical deviations around the line.

Simple Linear Regression

$$y_c = a + bx$$

y_c = Predicted (dependent) variable
x = Predicted (independent) variable
b = Slope of the line
a = Value of y_c when x = 0 (i.e., the height of the line at the y intercept)

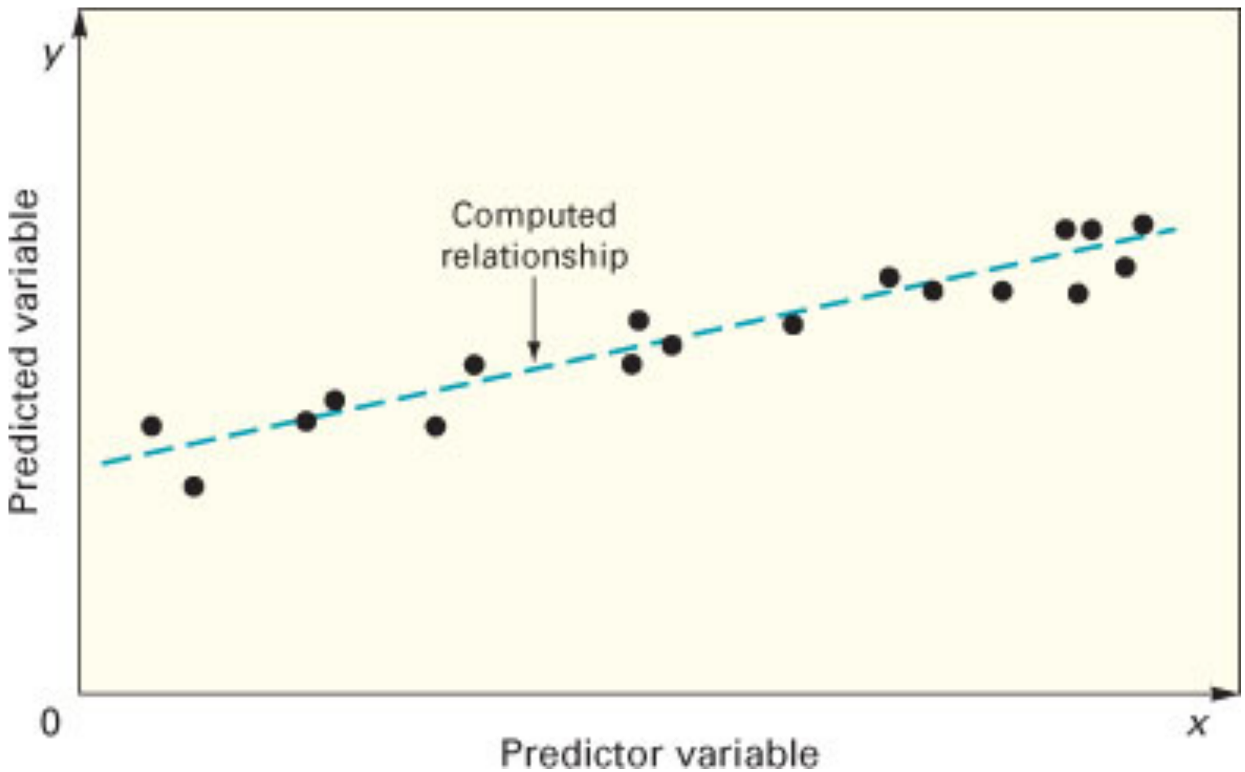

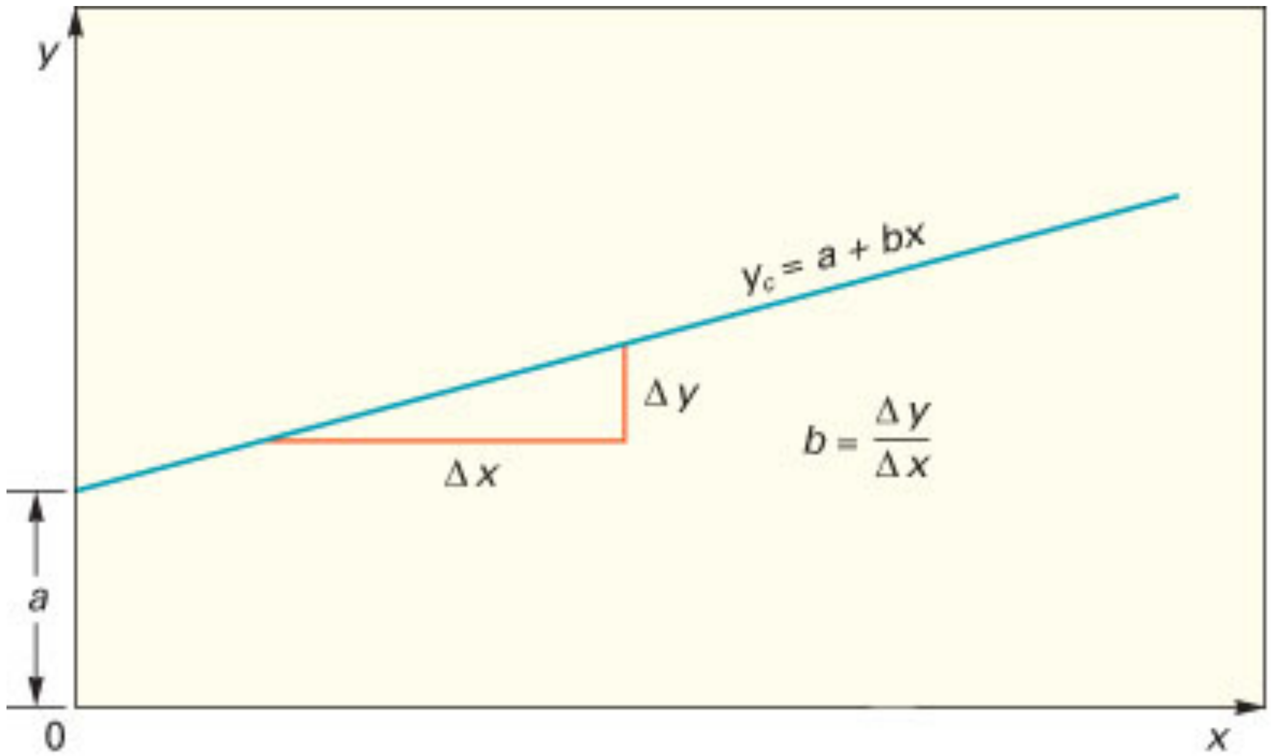

The line intersects the y axis where $y = a$. The slope of the line = b.

The coefficients *a* and *b* of the line are based on the following two equations:

$$b = \frac{n(\Sigma xy) - (\Sigma x)(\Sigma y)}{n(\Sigma x^2) - (\Sigma x)^2}$$

$$a = \frac{\Sigma y - b\Sigma x}{n} \text{ or } \bar{y} - b\bar{x}$$

n = Number of paired observations

Example

Healthy Hamburgers has a chain of 12 stores in northern Illinois. Sales figures and profits for the stores are given in the following table. Obtain a regression line for the data, and predict profit for a store assuming sales of $10 million.

Unit Sales, x (in $ millions)	Profits, y (in $ millions)
$ 7	$0.15
2	0.10
6	0.13
4	0.15
14	0.25
15	0.27
16	0.24
12	0.20
14	0.27
20	0.44
15	0.34
7	0.17

$$y_c = 0.0506 + 0.0159x$$

For sales of x = 10 (i.e., 10 million), estimated profit is;

$$y_c = 0.0506 + 0.0159(10) = 0.2099, \text{ or } \$209,900.$$

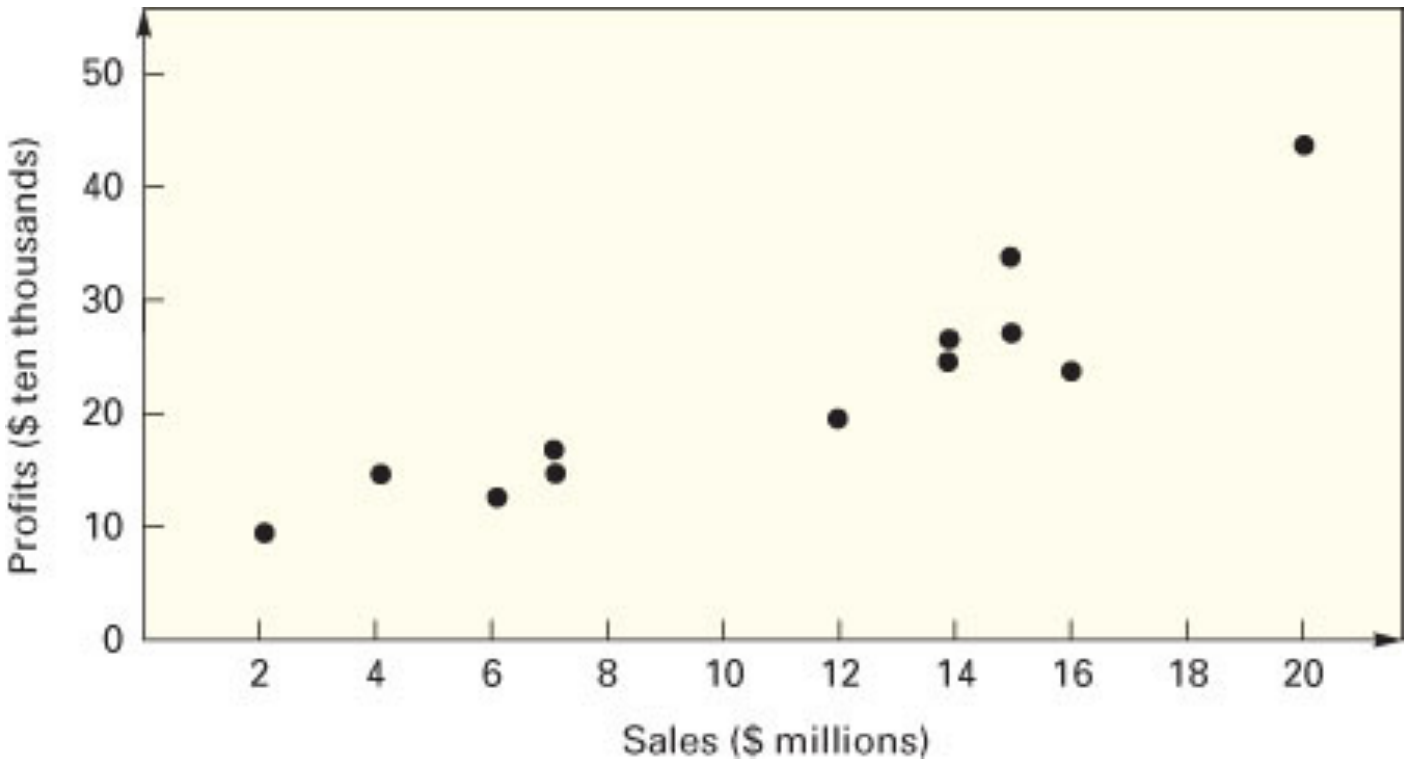

Simple Linear Regression

<Back

Clear

Slope =	0.0159	r =	0.9166657
Intercept =	0.0506008	r^2 =	0.840276

x	y	Forecast	Error
7	0.15	0.1621124	–0.0121124
2	0.1	0.0824612	0.0175388
6	0.13	0.1461822	–0.0161822
4	0.15	0.1143217	0.0356783
14	0.25	0.273624	–0.023624
15	0.27	0.2895543	–0.0195543
16	0.24	0.3054845	–0.0654845
12	0.2	0.2417636	–0.0417636
14	0.27	0.273624	–0.003624
20	0.44	0.3692054	0.0707946
15	0.34	0.2895543	0.0504457
7	0.17	0.1621124	0.0078876

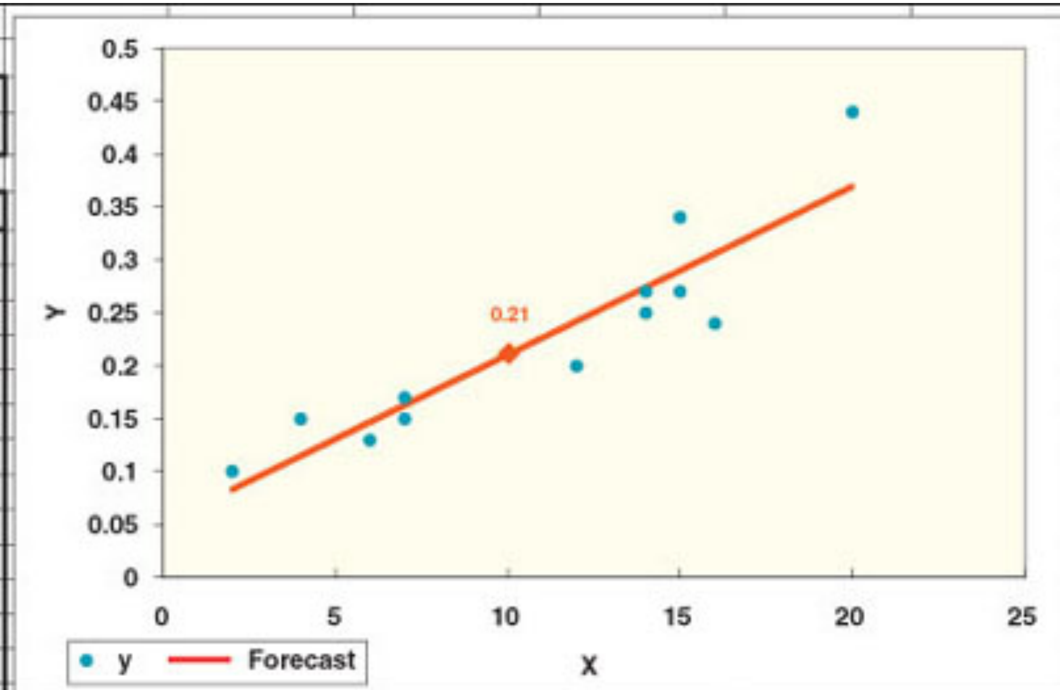

x = 10

Δx = 1

Forecast: 0.2099031

Correlation measures the strength and direction of relationship between two variables. Correlation can range from –1.00 to +1.00. A correlation of +1.00 indicates that changes in one variable are always matched by changes in the other; a correlation of –1.00 indicates that increases in one variable are matched by decreases in the other; and a correlation close to zero indicates little linear relationship between two variables. The correlation between two variables can be computed using the equation

$$r = \frac{n(\Sigma xy) - (\Sigma x)(\Sigma y)}{\sqrt{n(\Sigma x^2) - (\Sigma x)^2}\,\sqrt{n(\Sigma y^2) - (\Sigma y)^2}}$$

Assumptions

- Variations around the line are random. If they are random, no patterns such as cycles or trends should be apparent when the line and data are plotted.
- Deviations around the line should be normally distributed. A concentration of values close to the line with a small proportion of larger deviations supports the assumption of normality.
- Predictions are being made only within the range of observed values.

Best Results from Reggression

- Always plot the data to verify that a linear relationship is appropriate.

- The data may be time-dependent. Check this by plotting the dependent variable versus time; if patterns appear, use analysis of time series instead of regression, or use time as an independent variable as part of a multiple regression analysis.

- A small correlation may imply that other variables are important.

Weaknesses of Reggression

- Simple linear regression applies only to linear relationships with one independent variable.
- One needs a considerable amount of data to establish the relationship—in practice, 20 or more observations.
- All observations are weighted equally.

Example

Sales of new houses and three-month lagged unemployment are shown in the following table. Determine if unemployment levels can be used to predict demand for new houses and, if so, derive a predictive equation.

Period	1	2	3	4	5	6	7	8	9	10	11
Units sold	20	41	17	35	25	31	38	50	15	19	14
Unemployment % (three-month lag)	7.2	4.0	7.3	5.5	6.8	6.0	5.4	3.6	8.4	7.0	9.0

Plot the data to see if a linear model seems reasonable. In this case, a linear model seems appropriate for the range of the data.

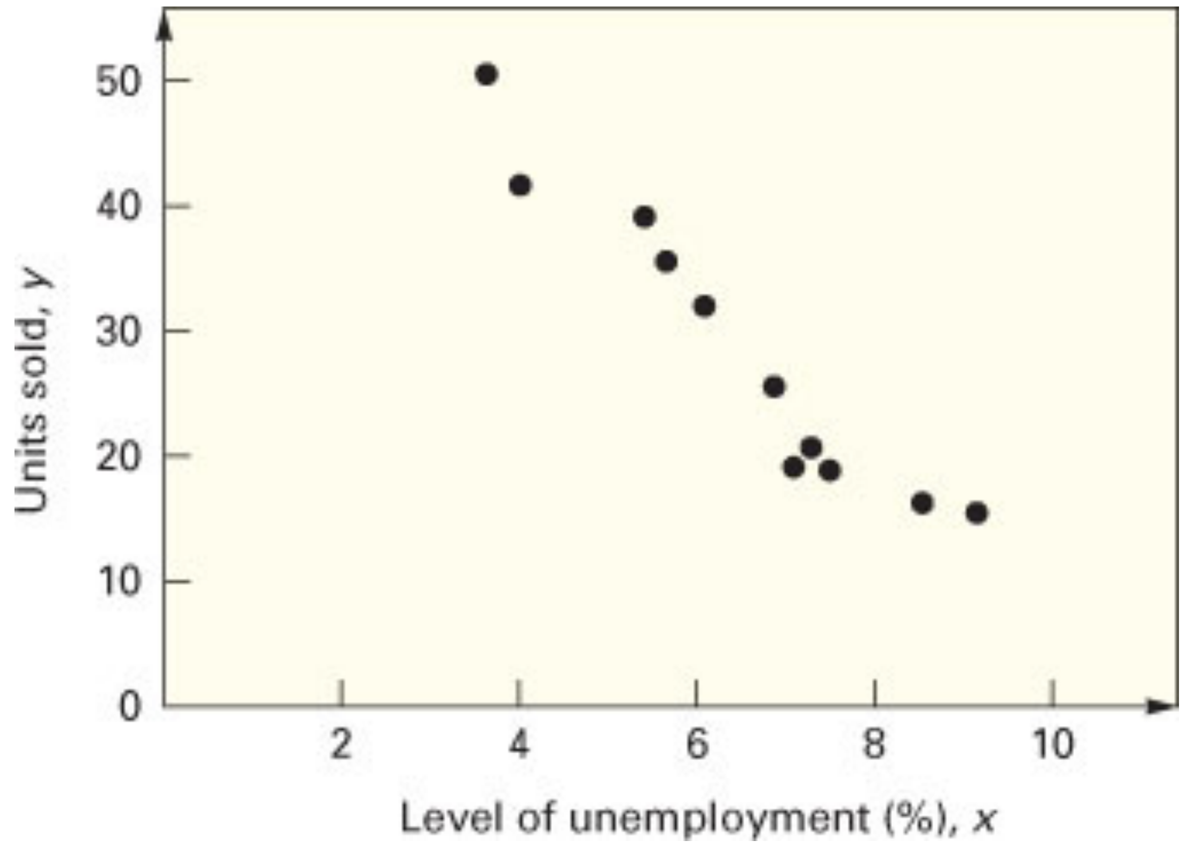

Check the correlation coefficient to confirm that it is not close to zero, and then obtain the regression equation:

$$r = -.966$$

This is a fairly high negative correlation. The regression equation is

$$y = 71.85 - 6.91x$$

Curvilinear Reggression and Multiple Reggression

Accuracy and Control of Forecasts

Forecast error: is the difference between the value that occurs and the value that was predicted for a given time period.

Error = Actual – Forecast:

$$e_t = A_t - F_t$$

Forecast Accuracy

Three commonly used measures for summarizing historical errors are;

1. The mean absolute deviation (MAD)
2. The mean squared error (MSE)
3. The mean absolute percent error (MAPE).

The formulas used to compute MAD, MSE, and MAPE are as follows:

$$\text{MAD} = \frac{\Sigma|\text{Actual}_t - \text{Forecast}_t|}{n}$$

$$\text{MSE} = \frac{\Sigma(\text{Actual}_t - \text{Forecast}_t)^2}{n - 1}$$

$$\text{MAPE} = \frac{\Sigma \frac{|\text{Actual}_t - \text{Forecast}_t|}{\text{Actual}_t} \times 100}{n}$$

Example

Compute MAD, MSE, and MAPE for the following data, showing actual and predicted numbers of accounts serviced.

Period	Actual	Forecast	(A − F) Error	|Error|	Error^2	[|Error| ÷ Actual] × 100
1	217	215	2	2	4	.92%
2	213	216	−3	3	9	1.41
3	216	215	1	1	1	.46
4	210	214	−4	4	16	1.90
5	213	211	2	2	4	.94
6	219	214	5	5	25	2.28
7	216	217	−1	1	1	.46
8	212	216	−4	4	16	1.89
			−2	22	76	10.26%

$$\text{MAD} = \frac{\Sigma|e|}{n} = \frac{22}{8} = 2.75$$

$$\text{MSE} = \frac{\Sigma e^2}{n-1} = \frac{76}{8-1} = 10.86$$

$$\text{MAPE} = \frac{\Sigma\left[\dfrac{|e|}{\text{Actual}} \times 100\right]}{n} = \frac{10.26\%}{8} = 1.28\%$$

Controlling the Forecast

- The model may be inadequate due to (a) the omission of an important variable, (b) a change or shift in the variable that the model cannot deal with (e.g., sudden appearance of a trend or cycle), or (c) the appearance of a new variable (e.g., new competitor).

- Irregular variations may occur due to severe weather or other natural phenomena, temporary shortages or breakdowns, catastrophes, or similar events.

- The forecasting technique may be used incorrectly, or the results misinterpreted.

- There are always random variations in the data. Randomness is the inherent variation that remains in the data after all causes of variation have been accounted for.

Control Chart

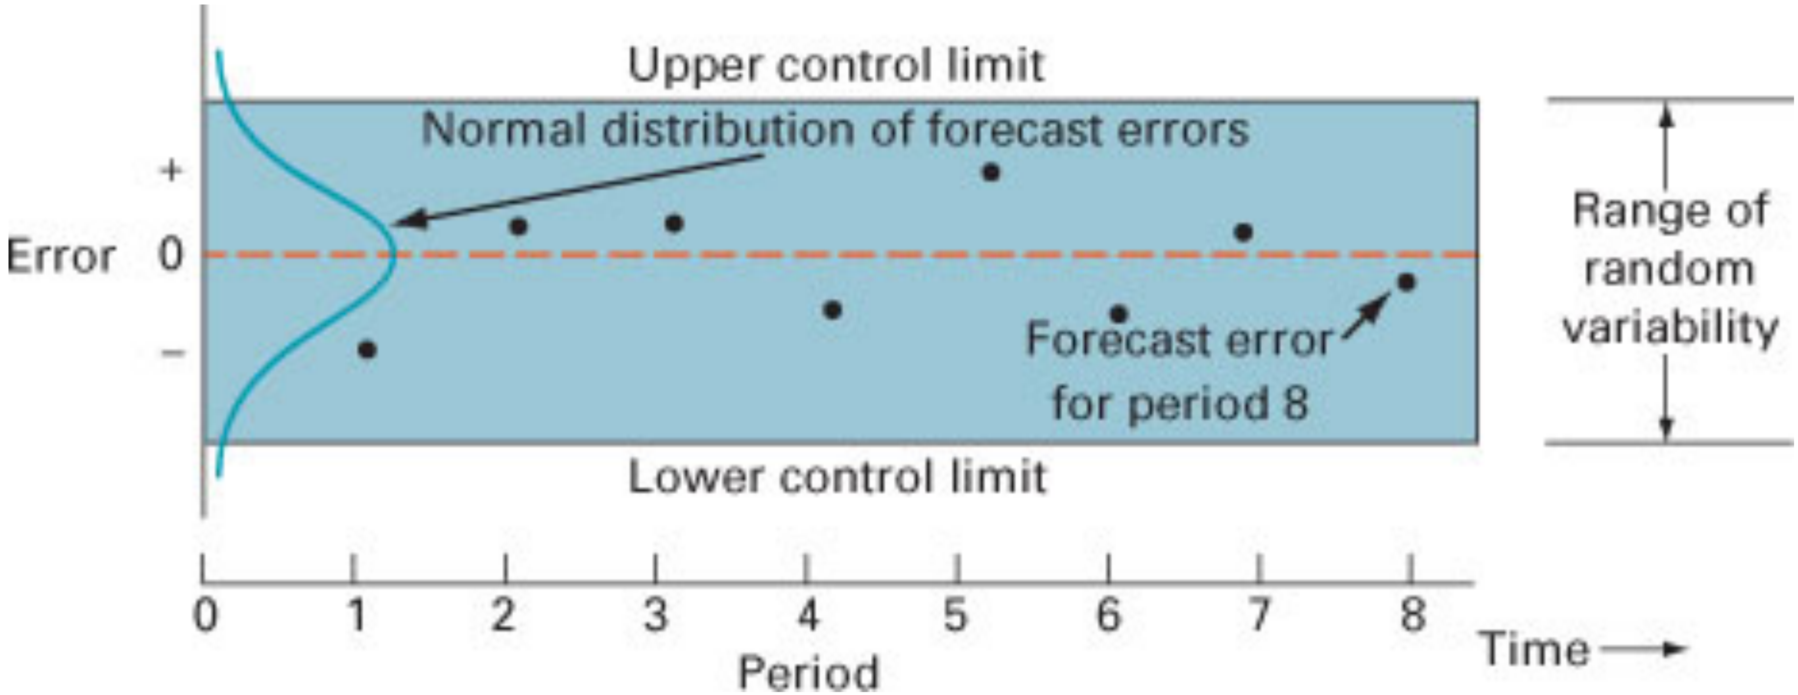

Examples of Nonrandomness

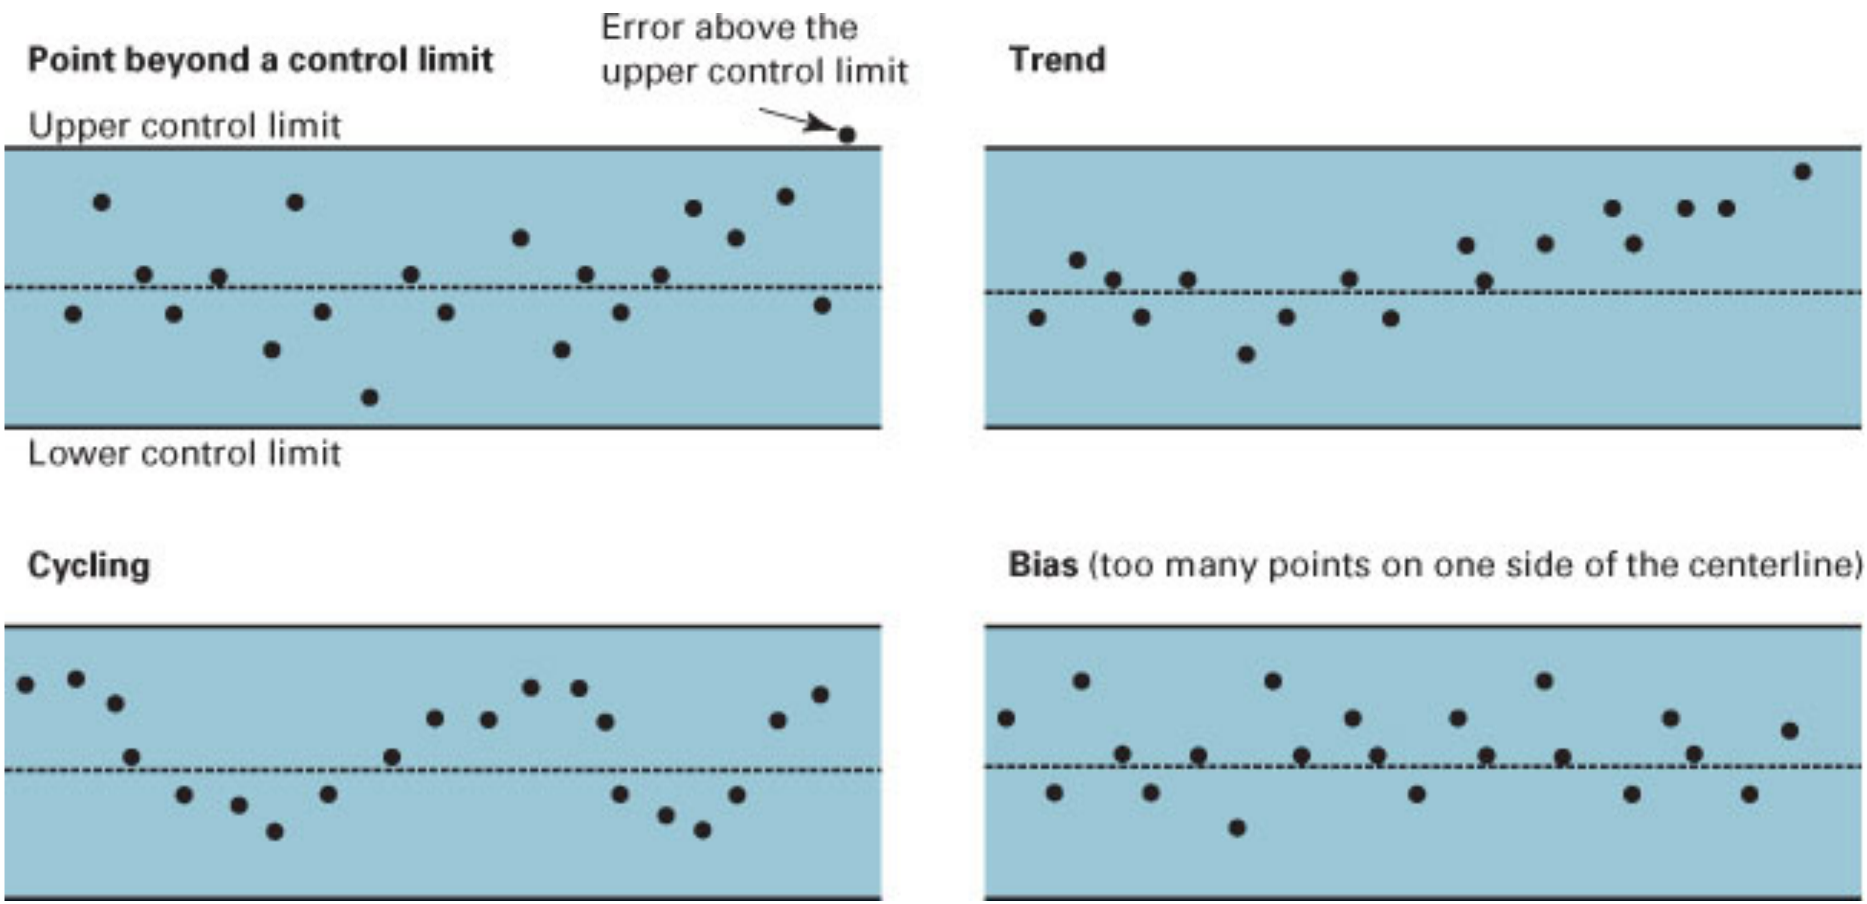

Control charts are based on the assumption that when errors are random, they will be distributed according to a normal distribution around a mean of zero. Recall that for a normal distribution, approximately 95.5 percent of the values (errors in this case) can be expected to fall within limits of 0 ± 2*s* (i.e., 0 ± 2 standard deviations), and approximately 99.7 percent of the values can be expected to fall within ± 3*s* of zero. With that in mind, the following formulas can be used to obtain the upper control limit (UCL) and the lower control limit (LCL):

$$s = \sqrt{\text{MSE}}$$

$$\text{UCL: } 0 + z\sqrt{\text{MSE}}$$
$$\text{LCL: } 0 - z\sqrt{\text{MSE}}$$

z = the number of standard deviations from the mean.

Example

Compute 2s control limits for forecast errors when the MSE is 2.0.

$$s = \sqrt{\text{MSE}} = 1.41$$

$$\text{UCL} = 0 + 2(1.41) = +2.82$$

$$\text{LCL} = 0 - 2(1.41) = -2.82$$

An older, less informative technique that is sometimes employed to monitor forecast errors is the tracking signal. It relates the cumulative forecast error to the average absolute error (i.e., MAD). The intent is to detect any bias in errors over time (i.e., a tendency for a sequence of errors to be positive or negative). The tracking signal is computed period by period using the following formula:

$$\text{Tracking signal}_t = \frac{\Sigma(\text{Actual}_t - \text{Forecast}_t)}{\text{MAD}_t}$$

Values can be positive or negative. A value of zero would be ideal; limits of ± 4 or ± 5 are often used for a range of acceptable values of the tracking signal. If a value outside the acceptable range occurs, that would be taken as a signal that there is bias in the forecast, and that corrective action is needed.

After an initial value of MAD had been determined, MAD can be updated using exponential smoothing:

$$\text{MAD}_t = \text{MAD}_{t-1} + \alpha(|\text{Actual} - \text{Forecast}|_t - \text{MAD}_{t-1})$$

Example

Monthly attendance at financial planning seminars for the past 24 months, and forecasts and errors for those months, are shown in the following table. Determine if the forecast is working using these approaches:

1. A tracking signal, beginning with month 10, updating MAD with exponential smoothing. Use limits of ±4 and $\alpha = .2$.
2. A control chart with 2s limits. Use data from the first eight months to develop the control chart, then evaluate the remaining data with the control chart.

Month	*A* (Attendance)	*F* (Forecast)	*A − F* (Error)	|*e*|	Cumulative |*e*|
1	47	43	4	4	4
2	51	44	7	7	11
3	54	50	4	4	15
4	55	51	4	4	19
5	49	54	−5	5	24
6	46	48	−2	2	26
7	38	46	−8	8	34
8	32	44	−12	12	46
9	25	35	−10	10	56
10	24	26	−2	2	58
11	30	25	5	5	
12	35	32	3	3	
13	44	34	10	10	
14	57	50	7	7	
15	60	51	9	9	
16	55	54	1	1	
17	51	55	−4	4	
18	48	51	−3	3	
19	42	50	−8	8	
20	30	43	−13	13	
21	28	38	−10	10	
22	25	27	−2	2	
23	35	27	8	8	
24	38	32	6	6	
			−11		

1. The sum of absolute errors through the 10th month is 58. Hence, the initial MAD is 58/10 = 5.8. The subsequent MADs are updated using the formula

$$\text{MAD}_{\text{new}} = \text{MAD}_{\text{old}} + \alpha(|e| - \text{MAD}_{\text{old}})$$

The tracking signal for any month is

$$\frac{\text{Cumulative error at that month}}{\text{Updated MAD at that month}}$$

t (Month)	$\lvert e\rvert$	$MAD_t = MAD_{t-1} + .2(\lvert e\rvert - MAD_{t-1})$	Cumulative Error	Tracking Signal = Cumulative $Error_t \div MAD_t$
10			−20	−20/5.800 = −3.45
11	5	5.640 = 5.8 + .2(5 − 5.8)	−15	−15/5.640 = −2.66
12	3	5.112 = 5.640 + .2(3 − 5.64)	−12	−12/5.112 = −2.35
13	10	6.090 = 5.112 + .2(10 − 5.112)	−2	−2/6.090 = −0.33
14	7	6.272 = 6.090 + .2(7 − 6.090)	5	5/6.272 = 0.80
15	9	6.818 = 6.272 + .2(9 − 6.272)	14	14/6.818 = 2.05
16	1	5.654 = 6.818 + .2(1 − 6.818)	15	15/5.654 = 2.65
17	4	5.323 = 5.654 + .2(4 − 5.654)	11	11/5.323 = 2.07
18	3	4.858 = 5.323 + .2(3 − 5.323)	8	8/4.858 = 1.65
19	8	5.486 = 4.858 + .2(8 − 4.858)	0	0/5.486 = 0.00
20	13	6.989 = 5.486 + .2(13 − 5.486)	−13	−13/6.989 = −1.86
21	10	7.591 = 6.989 + .2(10 − 6.989)	−23	−23/7.591 = −3.03
22	2	6.473 = 7.591 + .2(2 − 7.591)	−25	−25/6.473 = −3.86
23	8	6.778 = 6.473 + .2(8 − 6.473)	−17	−17/6.778 = −2.51
24	6	6.622 = 6.778 + .2(6 − 6.778)	−11	−11/6.622 = −1.66

2. a. Make sure that the average error is approximately zero, because a large average would suggest a biased forecast.

$$\text{Average error} = \frac{\Sigma \text{ errors}}{n} = \frac{-11}{24} = -0.46$$

b. Compute the standard deviation:

$$s = \sqrt{\text{MSE}} = \sqrt{\frac{\Sigma e^2}{n-1}}$$

$$= \sqrt{\frac{4^2 + 7^2 + 4^2 + 4^2 + (-5)^2 + (-2)^2 + (-8)^2 + (-12)^2}{8-1}} = 6.91$$

c. Determine 2s control limits:

$$0 \pm 2\text{s} = 0 \pm 2(6.91) = -13.82 \text{ to } +13.82$$

d.

i. Check that all errors are within the limits. (They are.)

ii. Plot the data (see the following graph), and check for nonrandom patterns. Note the strings of positive and negative errors. This suggests nonrandomness (and that an improved forecast is possible). The tracking signal did not reveal this.

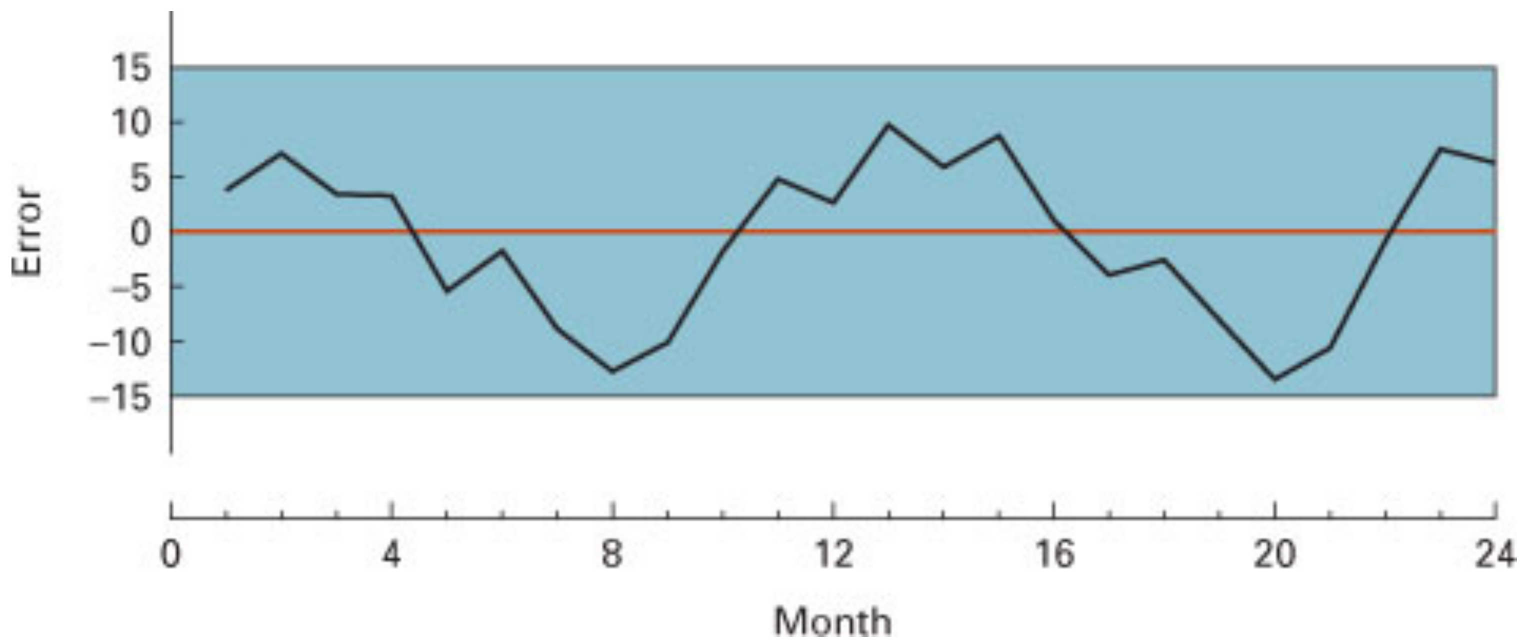

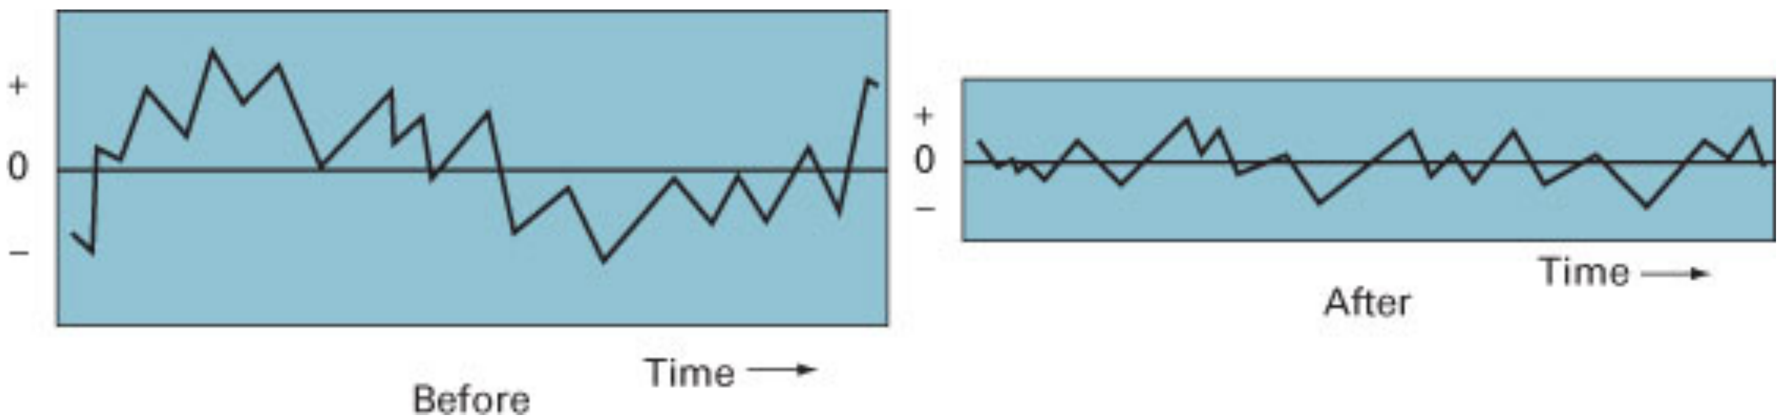
+
0
–
Time
Before
+
0
–
Time
After

Choosing a Forecasting Technique

Forecasting Method	Amount of Historical Data	Data Pattern	Forecast Horizon	Preparation Time	Personnel Background
Moving average	2 to 30 observations	Data should be stationary	Short	Short	Little sophistication
Simple exponential smoothing	5 to 10 observations	Data should be stationary	Short	Short	Little sophistication
Trend-adjusted exponential smoothing	10 to 15 observations	Trend	Short to medium	Short	Moderate sophistication
Trend models	10 to 20; for seasonality at least 5 per season	Trend	Short to medium	Short	Moderate sophistication
Seasonal	Enough to see 2 peaks and troughs	Handles cyclical and seasonal patterns	Short to medium	Short to moderate	Little sophistication
Causal regression models	10 observations per independent variable	Can handle complex patterns	Short, medium, or long	Long development time, short time for implementation	Considerable sophistication

Forecast Factors

Factor	Short Range	Intermediate Range	Long Range
1. Frequency	Often	Occasional	Infrequent
2. Level of aggregation	Item	Product family	Total output Type of product/ service
3. Type of model	Smoothing . . . . Projection Regression	Projection . . . . Seasonal Regression	Managerial judgment
4. Degree of management involvement	Low	Moderate	High
5. Cost per forecast	Low	Moderate	High