(continued)

Date	Sales	Date	Sales
Jan-15	47,906	Jan-16	65,711
Feb-15	53,570	Feb-16	68,005
Mar-15	69,189	Mar-16	78,029
Apr-15	64,346	Apr-16	92,764
May-15	77,267	May-16	97,175
Jun-15	75,787	Jun-16	86,255
Jul-15	74,052	Jul-16	90,496
Aug-15	79,756	Aug-16	87,602
Sep-15	73,292	Sep-16	83,577
Oct-15	77,207	Oct-16	92,610
Nov-15	68,423	Nov-16	73,949
Dec-15	67,274	Dec-16	77,711

(c5p12)

a. Prepare a time-series plot of the sales data. Does there appear to be a regular pattern of movement in the data that may be seasonal? Ronnie Mills, the product manager for this product line, believes that her brief review of sales data for the four-year period indicates that sales are slowest in November, December, January, and February than in other months. Do you agree?

b. Since production is closely related to orders for current shipment, Ronnie would like to have a monthly sales forecast that incorporates monthly fluctuations. She has asked you to develop a trend model that includes a time index and dummy variables for all but the above mentioned four months. Do these results support Ronnie's observations? Explain.

c. Ronnie believes that there has been some increase in the rate of sales growth. To test this and to include such a possibility in the forecasting effort, she has asked that you add the square of the time index (T) to your model (call this new term $T2$). Is there any evidence of increase of sales growth? Compare the results of this model with those found in part (b).

d. Use the model in part (c) to forecast sales for 2017. Calculate the mean absolute percentage error (MAPE) for the first six months of 2017. Actual sales for those six months were:

Jan-2017	87327
Feb-2017	84772
Mar-2017	112499
Apr-2017	102633
May-2017	112996
Jun-2017	119807

Chapter **Five Appendix**

Combining Forecasts (Ensemble Models)

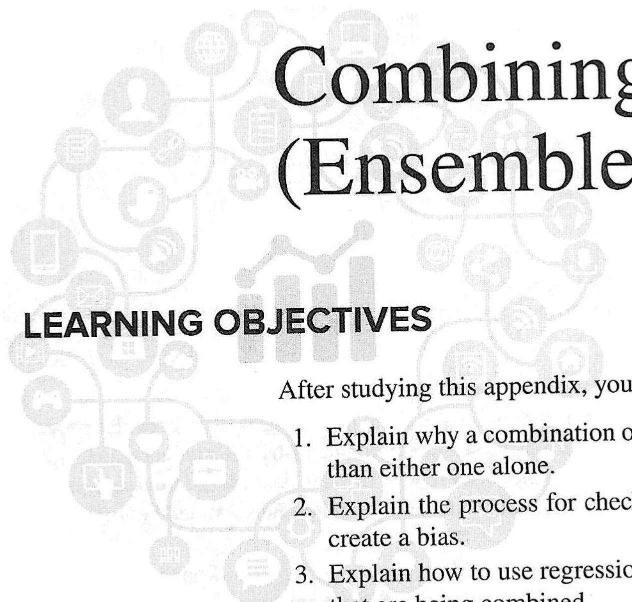

©VLADGRIN/Getty Images

LEARNING OBJECTIVES

After studying this appendix, you should be able to:

1. Explain why a combination of two forecasts (called ensembles) may be better than either one alone.
2. Explain the process for checking to see if a combination of forecasts would create a bias.
3. Explain how to use regression analysis to select the weights for the forecasts that are being combined.
4. Set up the data table that should be used when combining forecasts.
5. Use ForecastX™ to combine forecasts.

INTRODUCTION

The use of combinations of forecasts has been the subject of a great deal of research in forecasting. An indication of the importance of this concept is the fact that the prestigious *International Journal of Forecasting* had a special section, composed of seven articles, entitled "Combining Forecasts" in the year-end issue of the volume for 1989. In December 1992, an article in the same journal provided strong evidence on the importance of combining forecasts to improve accuracy. It was found that 83 percent of expert forecasters believe that combining forecasts will produce more accurate forecasts than could be obtained from the individual methods!

The idea of combining business forecasting models was originally proposed by Bates and Granger. Since the publication of their article, this strategy has received immense support in almost every empirical test of combined forecasts versus individual uncombined forecasts.

Throughout this book, we have emphasized the use of the mean absolute percentage error (MAPE) as a measure of the effectiveness of ***one*** particular forecasting model. In this appendix, instead of choosing the best model from among two or more alternatives, we are going to combine the forecasts from these different models to obtain *forecast improvement.* It may actually be unwise to simply determine which of a number of forecasting methods yields the most accurate predictions. A more reasoned approach, according to the empirical evidence, is to combine the forecasts already made in order to obtain a combined forecast that is more accurate than any of the separate predictions.

Any time a particular forecast is ignored because it is not the "best" forecast produced, it is likely that valuable independent information contained in the discarded forecast has been lost. The information lost is likely to be due to of one or both of the following:

1. Some variables included in the discarded forecast may not be included in the "best" forecast.
2. The discarded forecast may make use of a type of relationship ignored by the "best" forecast.

In the first of these cases, it is quite likely that various forecasts are based on different information. Thus, ignoring any one of these forecasts would necessarily exclude the explanatory power unique to the information included in the discarded model.

In the second situation, it is often the case that different assumptions are made in different models about the form of the relationship between the variables. Each of the different forms of relationship, however, may have some explanatory value. Choosing only the "best" of the relationships could exclude functional information. An example would be **accounting for seasonality** in a **multiplicative** manner with **Winters'** exponential smoothing and in an **additive** manner using seasonal dummy variables in **regression**.

BIAS

To be useful, forecasts we wish to combine must be unbiased. ***To be unbiased, each of the forecasts cannot consistently overestimate or underestimate the actual value.*** If we combined an unbiased forecast with one that consistently overestimated the true value, we would always wind up with a biased estimate. Combining forecasts will not a eliminate systematic bias in a forecast.

The idea of bias related to a forecast can be explained in graphical manner. Consider the graphs in Figure 5A.1. In each of the three graphs, actual values of the variable to be forecast are on the vertical axis and the forecast values are on the horizontal axis. In the top graph, the 45-degree line represents a situation in which there is a perfect forecast. While a perfect forecast is very unlikely, the purpose of this line is to show that such a forecast would not have a bias. Note two important observations about the line in the upper graph: 1) The slope of such a 45-degree line is equal to one; and 2) the 45-degree line has an intercept equal to zero. The importance of these observations will become clear as we discuss bias in the context of combining forecasts,

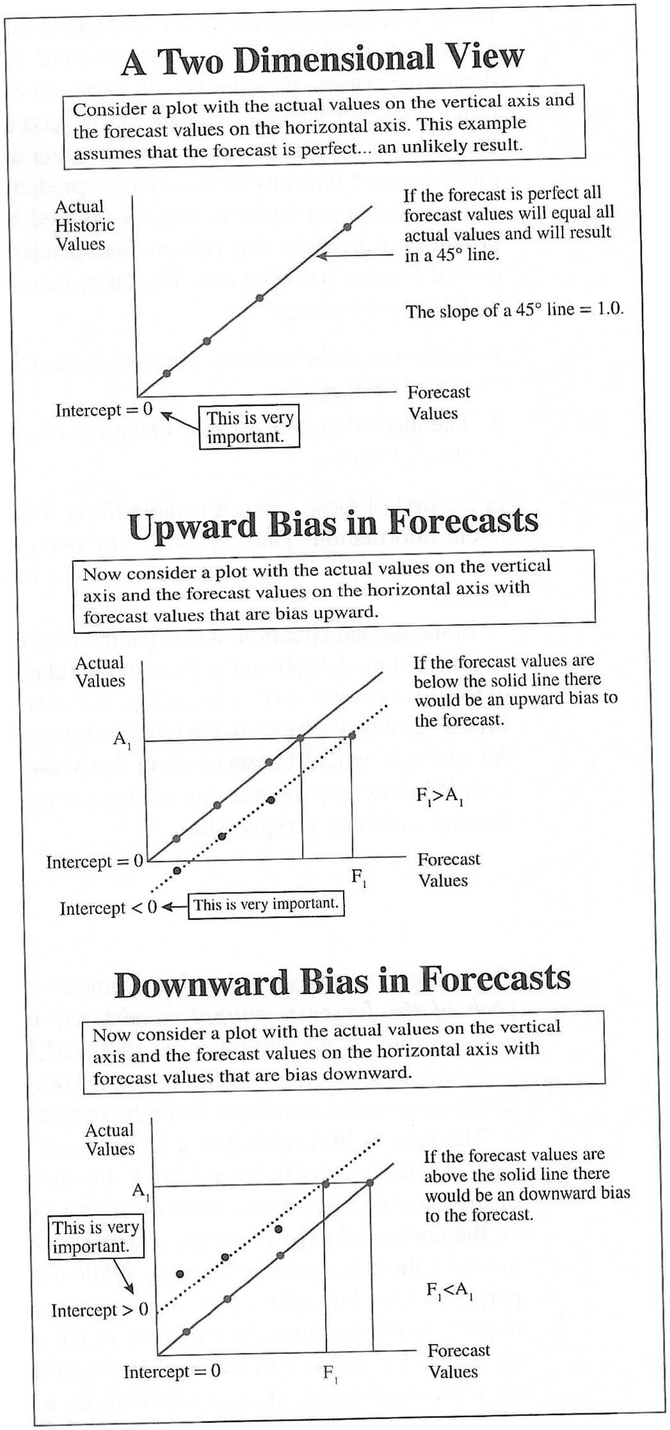

FIGURE 5A.1
One way to look at the concept of forecast bias.
In these graphs, the solid line represents a forecast with no bias.The two dotted lines show forecasts that have a bias. The top graph shows a forecast with no bias. The middle graph illustrates a situation in which a forecast has an upward bias. In the lower graph, a downward bias is shown.

Bias can arise from a number of sources, but perhaps the most common source is the forecaster's preconceived notions. Predictions of forecasters not only reflect what they believe to be the truth but also what they would *like* the truth to be. This statement is best demonstrated by the results obtained by Hayes in a survey of voters two weeks before the Roosevelt-Hoover election. Hayes found that of the people who intended to vote for Hoover, 84 percent thought that he would win the election. Of the people who intended to vote for Roosevelt, however, only 6 percent thought that Hoover would win. Apparently, those who intended to vote for a particular candidate are biased in the sense that they also believe that their favorite will actually win the election.[1]

A forecaster should spend some time examining multiple forecasting models in the hope of combining some or all of these models into a combined forecast that is superior to any of the individual forecasts.

Professional forecasters may suffer from the same bias as voters—they may look for forecasting models that confirm their own preconceived ideas. To eliminate bias, a forecaster will have to examine models that may contradict his or her current beliefs. What this means is that you must do something that runs counter to your intuition in order to examine models you may feel are incorrect; you must examine forecasting models that you may believe to be inferior to your "favorite" model. This prescription is more difficult to implement than it sounds. Much of a forecaster's time is spent in confirming existing beliefs of how the world works. However, we are suggesting that a forecaster should spend some time examining multiple forecasting models in the hope of combining some or all of these models into a combined forecast that is superior to any of the individual forecasts.

WHAT KINDS OF FORECASTS CAN BE COMBINED?

The example of combining forecasts we used in the previous section is one of the simpler combinations a researcher could try. In actual practice, it would be more common to find a forecaster using very different types of models in order to construct a combination forecast.

Recall that the premise in constructing combined forecasts is:

1. That the different forecasting models *extract different predictive factors* from essentially the same data, or
2. That the different models offer different predictions because they *use different variables.*

We should expect that combinations of forecasts that use very different models are likely to be effective in reducing forecast error.

We should expect that combinations of forecasts that use very different models are likely to be effective in reducing forecast error.

Consider Figure 5A.2, which conceptually presents a 10-year forecast of air travel in the United States. The judgmental method represents a mail survey of experts outside the airline industry. The extrapolation method could be a form of exponential smoothing. The segmentation method surveys airline travelers in different segments of the market and then combines the results to obtain a total picture of the industry. The econometric method refers to a causal regression model.

[1] S. P. Hayes, Jr., "The Predictive Ability of Voters," *Journal of Social Psychology* 7 (1936), pp. 183–91.

FIGURE 5A.2
Combining forecasts from different methods

Note: The w's, the relative weights on various forecasts, should sum to 1.0.

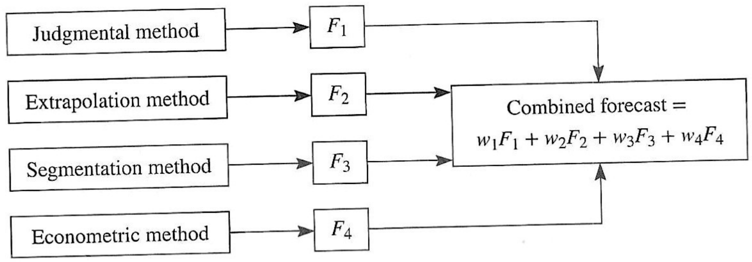

All four methods could be employed and their predictions weighted by the values w_1 to w_4 in order to calculate the combined forecast. Such a diverse combined forecast would benefit from both the use of the different techniques *and* the use of different sources of data. If each of the methods employed were also constructed and estimated by a different forecaster, another source of possible bias may also have been minimized; this provides a safeguard by making it difficult to cheat.

CONSIDERATIONS IN CHOOSING THE WEIGHTS FOR COMBINED FORECASTS

Combined forecasts are used in place of individual forecasts in order to reduce forecast error, and the results of the combined methods are quite often impressive. Armstrong has reported results from reanalyzing eight separate studies that provided sufficient information to test the combined forecasting method against individual forecast models.[2] In each case, Armstrong used equal weights for the individual forecasts, following his belief that weight should be chosen *ex ante*. The combinations of two forecasts reduced error (measured as mean absolute percentage error) by a significant 6.6 percent. In no single case did the accuracy ever suffer.

Even though the use of equal weights for each of the individual forecasts offers the advantage of simplicity and precludes the forecaster's own bias in the selection of weighting factors, there may be a good reason for weighting one individual forecast more than another. Equal weights do not take into account the relative accuracy of the individual forecasting models that are combined. Bates and Granger were the first to indicate that, by weighting the more accurate of the methods more heavily, the overall forecast could be improved.[3]

It is the diversity of information included in the individual models that allows a combined forecast model to assemble the pieces to form a more powerful forecasting model than any one of the parts.

In general, a combined forecast will have a smaller error, as measured by MAPE, *unless individual forecasting models are almost equally good and their forecast errors are highly correlated.* If, however, the forecast model with the lower MAPE is more heavily weighted, the combined forecast should improve even further.

[2] See J. Scott Armstrong, *Long-Range Forecasting from Crystal Ball to Computer*. 2nd ed. (New York: John Wiley & Sons, 1985).

[3] See J. M. Bates and C. W. J. Granger. "The Combination of Forecasts." *Operational Research Quarterly* 20, no. 4 (1969), p. 452.

It is the diversity of information included in the individual models that allows a combined forecast model to assemble the pieces to form a more powerful forecasting model than any one of the parts.

ONE TECHNIQUE FOR SELECTING WEIGHTS WHEN COMBINING FORECASTS

One technique that is used to combine forecasts in order to improve accuracy uses regression concepts you have now learned. This technique involves the use of a regression analysis in determining the weights. Charles Nelson suggests that if we are trying to weight a portfolio of forecasts in order to minimize the forecast error, an optimal linear composite forecast would be:[4]

$$F^* = b_1F(1) + b_2F(2)$$

where:

F^* = Optimal combined forecast
$F(1)$ = First individual forecast
$F(2)$ = Second individual forecast
b_1 = Weight allocated to the first forecast
b_2 = Weight allocated to the second forecast

The actual values of b_1 and b_2 would be calculated by running a regression with the past actual values as the dependent variable and the forecasted values for each individual model as the independent variables. Note that this is not exactly the type of regression we have run before in the text; this regression has no intercept term, and so the equation must be calculated in a manner different from that we have used earlier.

Using this method, if the two (or more) individual forecasts are free of systematic bias, the values of b_1 and b_2 will sum roughly to 1. The t-ratios for the regression will essentially answer the question whether individual forecast 1 adds any explanatory power to what is already present in forecast 2 and similarly for forecast 2. If the b_1 value passes the t-test at some reasonable confidence level, we can be assured that the first individual model, $F(1)$, did add explanatory power when combined with the second model, $F(2)$, using the weights calculated by the regression.

To apply this method and to determine the best values for b_1 and b_2, a two-step regression process is used. First, you perform a standard multiple regression of the actual values (dependent variable) on the values predicted from the individual forecasting methods (independent variables in this regression). We can express this as:

$$A = a + b_1F(1) + b_2F(2)$$

[4] Charles R. Nelson, "A Benchmark for the Accuracy of Econometric Forecasts of GNP," *Business Economics* 19, no. 3 (April 1984), pp. 52–58.

The value of the intercept (a) should be (not statistically different from) zero if there is no bias in the combined forecast. A standard t-test can be used to test whether the intercept is significantly different from zero.[5] Note that a two-tailed test would be appropriate here because you want to know if the intercept is different from zero in either a positive or negative direction.

Assuming that you conclude that $a = 0$, you then redo the regression, forcing the regression through the origin. Most regression programs provide an option that allows this to be done quite easily. The result of regressing the actual values on the two forecast series, without an intercept, yields the desired result to determine the best weights to be used in combining the forecasts. We have:

$$F^* = b_1F(1) + b_2F(2)$$

Using these values of b_1 and b_2, along with the $F(1)$ and $F(2)$ forecast series, the optimal combined forecast, F^*, is easily determined.

As indicated, the values of b_1 and b_2 should sum roughly to 1. On occasion, one of these weights may be negative, in which case interpretation is tenuous. Some forecasters use such a model even if b_1 or b_2 is negative, as long as the MAPE for F^* is lower than for $F(1)$ or $F(2)$ alone. However, we advise using this method only when both weights are positive. It should be noted that this method can be extended to include more than two forecast series in the combination process. Remember, however, that each method should have unique information content.

An Application of the Regression Method for Combining Forecasts

To illustrate the widely used regression method of combining forecasts, we will apply it to the problem of forecasting the sales of women's clothing by a retail chain. Figure 5A.3 shows the data in graphic form.

You have learned in Chapter 3 how to develop a Winters' model for seasonal series such as this. In Chapter 5, you learned to build a regression model to make a forecast. Therefore, here we will only present the results of those two forecasts. Winters' is an obvious choice since it accounts for both trend and seasonality in the data. In the causal regression model, dummy variables were used to account for the seasonality. Thus, in this case, we have one extrapolation model (Winters') that deals with seasonality in a multiplicative way and a causal model that accounts for seasonality in an additive manner. The two methods bring different perspectives to the table in terms of forecasting.

Based on 2010 through 2015 quarterly data, the MAPEs for the Winters' and regression models alone are:

	Winter's	Regression
Historical Period MAPEs	2.04%	4.04%

[5] This is one of the few cases in which we are interested in testing to see whether the intercept is different from zero. Normally, we do this test only for the slope terms.

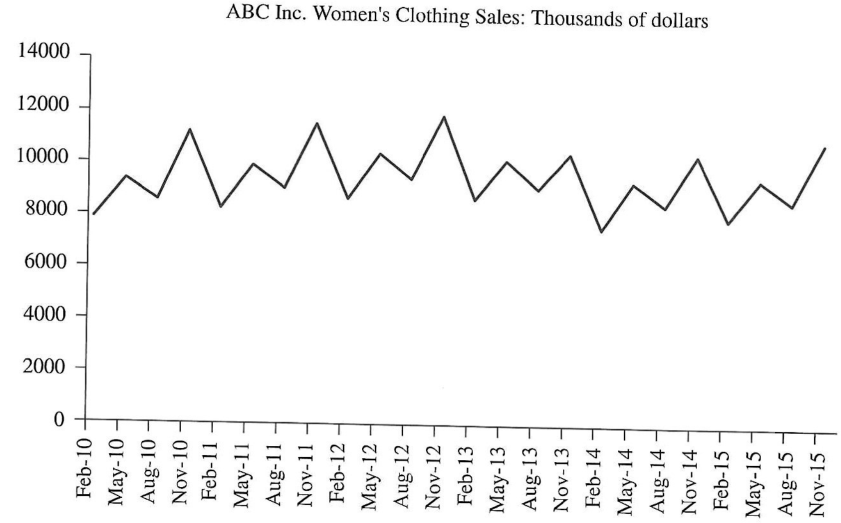

FIGURE 5A.3
Six years of quarterly sales of women's clothing by a retailer. In these sales data, there is clear seasonality and a very slight upward trend. (c5A3)

Based on these historical period MAPEs, the Winters' model clearly appears to be a better forecasting model than the regression model. Both MAPEs are relatively low.

If we want to combine these two models, the first thing we must do is to check to see if this combination would create a bias. We first regress the actual sales as a function of the regression and Winters' predictions. Doing so yields the following result:

Dependent Variable Is Sales of Women's Clothing by a Retailer

	Coefficients	Std Error	*t* Stat	P-value
Intercept	145.425	501.486	0.290	**0.775**
Regression Predictions	0.241	0.140	1.724	0.099
Winters' Predictions	0.745	0.127	5.858	0.000

Go back and look at Figure 5A.1. You see that if there is no bias, the fit between actual and predicted values results in a line with a zero intercept and a slope of 1. Here we have a three-dimensional situation (actual sales, Winters' predictions, and regression predictions), but the concepts presented in Figure 5A.1 still hold.

The only thing we are interested in at this point is whether the intercept is significantly different from zero. To evaluate bias, we look at a *t*-test for the intercept. This is a test we have not previously done because in most regression applications, we do not care about whether the intercept is positive, negative, or zero. In the current application, we look for statistical evidence that combining the Winters' and regression forecasts will not create a bias. The hypothesis test is:

$$H_o: \beta = 0$$
$$H_1: \beta \neq 0$$

We can evaluate this by looking at either how the t-statistic compares with the table value at [n – (k+1)] degrees of freedom. For the current example, this is [(24-(1+1)] = 22. Using a 95 percent confidence level, the table value of t is 2.074. We see above that the calculated t is only 0.290, which is less than 2.074. We could also compare the P-value with the desired significance level or 0.05 (remember, the significance level is 1 minus the confidence level).We see that the P-value is greater than 0.05. Either way one looks at this the statistical decision is to fail to reject the null hypothesis (fail to reject H_0: $\beta = 0$). Thus, we believe that the true intercept is essentially zero.

The next step is to do the regression forcing the intercept to be zero. This can be done in most statistical software. At the end of this appendix, you will see how to do it in ForecastX™. In Excel when doing regression, the dialog box looks as follows:

In the red oval, you see that "Constant is Zero" is checked.

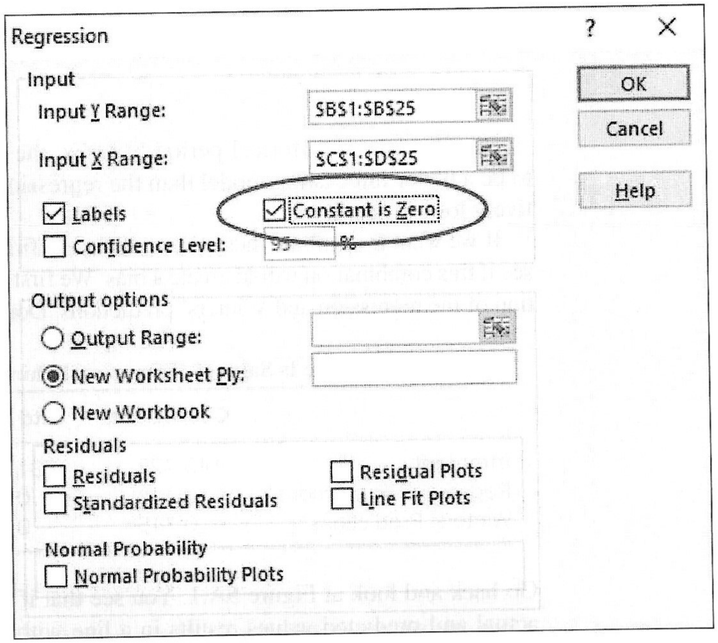

Source: John Galt Solutions

Regressing sales on the Regression predictions and Winters' predictions with a zero constant gives us the following result:

	Coefficients	Std Error	*t* Stat	P-value
Intercept	0	#N/A	#N/A	#N/A
Regression Predictions	**0.258**	0.124	2.076	0.0498
Winters' Predictions	**0.743**	0.124	5.976	0.000
Sum of Weights	1.001			

Now we have the optimal weights to assign to the regression and Winters' models. Recall that the MAPE for the regression model is 4.04 percent and the MAPE for the Winters' model is 2.04 percent. Thus, you would expect the Winters' model to get a higher weight in the combined model. The results above tell us exactly that. The Winters' model gets a weight of 74.3 percent, and the regression model gets a weight of 25.8 percent. The sum of these weights is about 1.001, which is essentially the 1.0 we would have expected.

Summarizing the Steps for Combining Forecasts

First, consider how the data should be set up in Excel. The Excel file, in general, should be set up as shown to the left in Figure 5A.4. The number of rows of historical data will almost always greatly exceed the number of data rows for the forecast values. On the right side of Figure 5A.4, you see the data sheet for our current example.

Second, regress the actual values of the variable to be forecast on the two (or more) forecast results for the historic period. In our example, that result is:

Dependent Variable Is Sales of Women's Clothing by a Retailer

	Coefficients	Std Error	*t* Stat	P-value
Intercept	145.425	501.486	0.290	0.775
Regression Predictions	0.241	0.140	1.724	0.099
Winters' Predictions	0.745	0.127	5.858	0.000

Since the P-value is greater than 0.05, we conclude that the intercept is essentially zero and thus combining the methods will not create a bias. If the P-value is less than 0.05, a bias would be created, so those forecasts should not be combined. In which case, stop here.

Assuming no bias, proceed to do the same regression but this time force the constant to be zero. In the current example, the result is:

	Coefficients	Std Error	*t* Stat	P-value
Intercept	0	#N/A	#N/A	#N/A
Regression Predictions	**0.258**	0.124	2.076	0.0498
Winters' Predictions	**0.743**	0.124	5.976	0.000
Sum of Weights	1.002			

This provides the weights to be applied to each of the two (or more) forecast methods. You will see that this process is easily done within ForecastX™.

FIGURE 5A.4 **The format for data file to combine forecasts. The actual and fitted values are at the top of the data file. In the forecast horizon, there are only values for the two forecasts. The left side is the general format, while the right side is the actual file for the current example.** (c5A3)

This is the general format for setting up the data file to combine forcasts.

Date	Actual Data	Model 1 Forecast	Model 2 Forecast
	Historic	Fitted	Fitted
	Values	Values	Values
		for Model 1	for Model 2
	?		
	?		
	?		
	?	Forecast	Forecast
	?	Values	Values
	?	for Model 1	for Model 2
	?		
	?		

Date	ABC Inc. Women's Clothing Sales: Thousands of dollars	Regression Predictions	Winters' Predictions
Feb-10	7872	8,108.64	7,872.00
May-10	9360	9,737.64	9,172.26
Aug-10	8559	8,817.21	8,375.77
Nov-10	11209	11,039.54	10,895.52
Feb-11	8244	8,131.92	8,283.54
May-11	9911	9,762.44	9,672.25
Aug-11	9022	8,840.12	8,862.20
Nov-11	11539	11,062.05	11,517.82
Feb-12	8640	8,151.10	8,511.76
May-12	10398	9,779.13	10,163.01
Aug-12	9413	8,855.86	9,285.01
Nov-12	11886	11,079.36	11,977.78
Feb-13	8658	8,173.08	8,822.79
May-13	10172	9,811.03	10,302.45
Aug-13	9071	8,877.87	9,149.97
Nov-13	10450	11,087.59	11,545.74
Feb-14	7552	8,164.40	7,940.43
May-14	9353	9,790.58	9,082.21
Aug-14	8452	8,860.38	8,325.65
Nov-14	10423	11,079.38	10,410.71
Feb-15	7933	8,169.86	7,806.03
May-15	9487	9,800.19	9,574.93
Aug-15	8611	8,876.56	8,501.86
Nov-15	10938	11,097.07	10,572.75
Feb-16		8,183.66	8,152.07
May-16		9,810.96	9,814.13
Aug-16		8,887.33	8,826.02
Nov-16		11,108.00	10,378.00

This portion of the table shows the historical period and the fitted values from two different forecast methods.

These are forecast values.

Integrative Case

The Gap

FORECASTING THE GAP SALES DATA WITH A COMBINATION MODEL

The sales of The Gap stores for the 44 quarters covering 2006Q1 through 2017Q4 are again shown in the graph below. Recall that The Gap sales data are quite seasonal and are increasing over time. Use the full 2006Q1 through 2017Q4 data to construct your forecast models. (c5Gap)

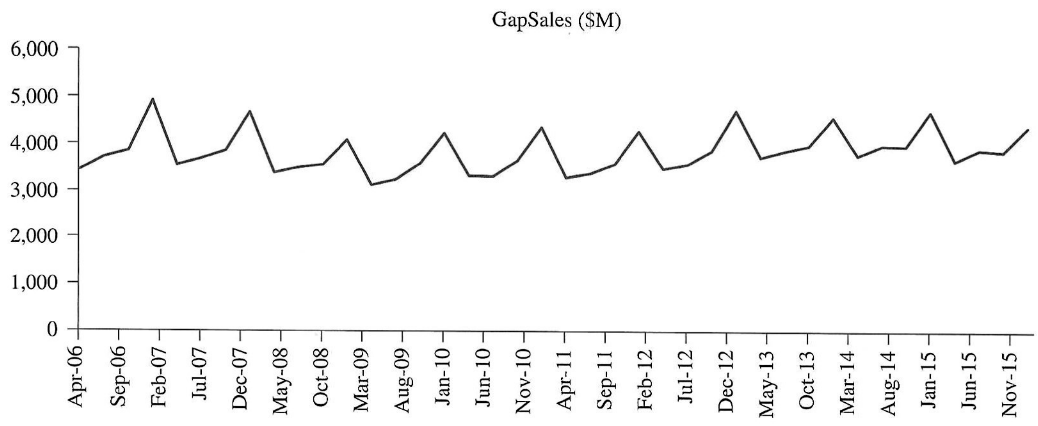

Case Questions

1. Assume you would like to combine a Winters' model and a multiple-regression model. For Gap sales, a Winters' forecast result and a regression forecast result are provided in the c5Gap Excel file. You do not have to do those forecasts. Use regression analyses to check for bias and to determine the best weights for each individual model.
2. Combine the two methods (i.e., the Winters' and the multiple-regression models) using the weights found in question 1.
3. The MAPE for the Winters' model in the historical period is 2.60 percent. The comparable MAPE for the regression model is 2.87 percent. Calculate the mean absolute percentage error for the combined model in the historical period, and comment on any improvement.

Solutions to Case Questions

1. To see if both models may reasonably be used in a combined forecast, run the regression that uses The Gap sales as the dependent variable and the two forecasts (one from the Winters' model and the other from the multiple-regression model) as the explanatory variables. The regression with a constant term indicates that the constant term is not significantly different from zero (at a 95 percent confidence level) because its

t-statistic of 0.682 is smaller than 1.96. Also, you see that the P-value of 0.499 is greater than the desired significance level of 0.05. Thus, we may combine the models without creating a bias.

Regression with a Constant Term (c5Gap)

	Coefficients	Standard Error	*t* Stat	P-value
Intercept	103.086	151.177	0.682	**0.499**
Multiple Regression	0.418	0.110	3.808	0.000
Winters'	0.554	0.098	5.657	0.000

Note that the constant term is not statistically different from zero.

Regression without a Constant Term (c5Gap)

	Coefficients	Standard Error	*t* Stat	P-value
Intercept	0	#N/A	#N/A	#N/A
Multiple Regression	**0.453**	0.097	4.677	0.000
Winters'	**0.546**	0.097	5.652	0.000
Sum =	0.999			

2. The two models are combined by running the same regression through the origin, as shown in the lower part of the regression results in question 1. Here the dependent variable is again Gap sales. Note that the weight on the Winters' forecast (i.e., 0.546) is larger than the weight on the multiple-regression forecast (i.e., 0.453). This is to be expected since the Winters' forecast alone has a MAPE of 2.57 percent compared with the multiple-regression forecast MAPE of 2.87 percent. In this situation, the MAPEs are not dramatically different, and we see that the weights assigned to each method in the combination are also fairly close.

   Note the very close association of the forecast with the original data, as shown in the graph below:

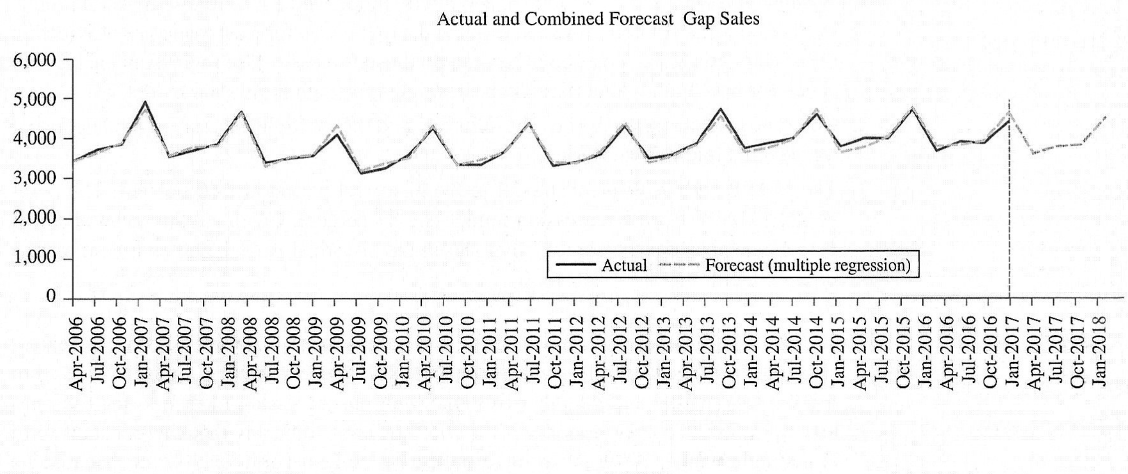

3. The combined forecast and the two candidate models can be compared by using the MAPE each model calculated on known historic data. The MAPEs are:

	RMSE
Winters' model	2.57%
Regression model	2.87%
Combination model	2.14%

Since the lowest MAPE calculated belongs to the combination model, it appears that there may be some support for forecast improvement from the combined model.

USING FORECASTX™ TO COMBINE FORECASTS

The first thing you need to do is to set up a data file with dates in column A, the historical sales (or other variable you want to forecast) in column B, one of the forecasts to be combined in column C, and the other forecast to be combined in column D. This was illustrated in Figure 5A.4.

As another example, an abbreviated portion of the Gap sales data and two forecasts are shown below in the format you need in order to combine the Winters' and regression forecasts:

Dates	Actual Gap Sales Data (M$)	Winters'	Regression
Apr-2006	3,441.00	3,441.00	3,432.73
Jul-2006	3,714.00	3,646.70	3,627.45
Oct-2006	3,851.00	3,824.94	3,951.77
Jan-2007	4,919.00	4,793.62	4,659.00
.	.	.	.
.	.	.	.
.	.	.	.
Apr-2014	3,774.00	3,650.75	3,600.17
Jul-2014	3,981.00	3,775.77	3,744.44
Oct-2014	3,972.00	4,054.75	3,994.29
Jan-2015	4,708.00	4,751.40	4,769.62
Apr-2015	3,657.00	3,825.59	3,744.35
Jul-2015	3,898.00	3,816.05	3,813.18
Oct-2015	3,857.00	3,899.29	3,940.54
Jan-2016	4,385.00	4,617.29	4,608.63
Apr-2016		3,573.47	3,606.92
Jul-2016		3,762.95	3,773.15
Oct-2016		3,746.51	3,914.18
Jan-2017		4,385.00	4,597.56

As usual, begin by opening your data file in Excel. Place your cursor in one of the sales data cells such as B5. Then start ForecastX™. In the **Data Capture** dialog box, identify the data you want to use, as shown here. Note that in this example, we have a sheet that has the date, the actual values for the Gap sales, and then the two forecasts we want to combine. It is important that in the **Data Capture** box for the process of combining forecast you have the **Forecast Periods** set to the same number of periods as the number of empty cells for the dependent variable.

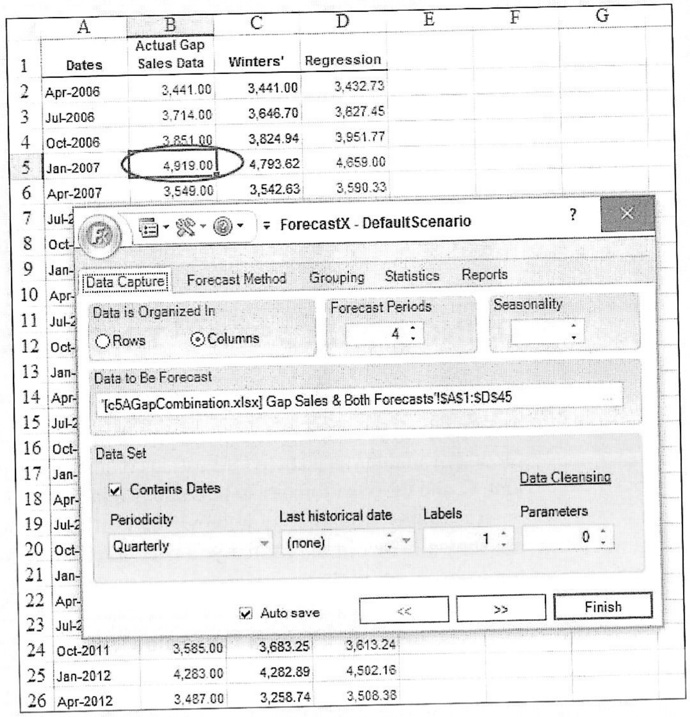

Source: John Galt Solutions

Then click the **Forecast Method** tab. In the **Method Selection** dialog box, click the down arrow in the **Forecasting Technique** box and select **Multiple Regression.** Make sure the desired variable is selected as the **Dependent Series,** which is the actual value of **Gap sales** in this example.

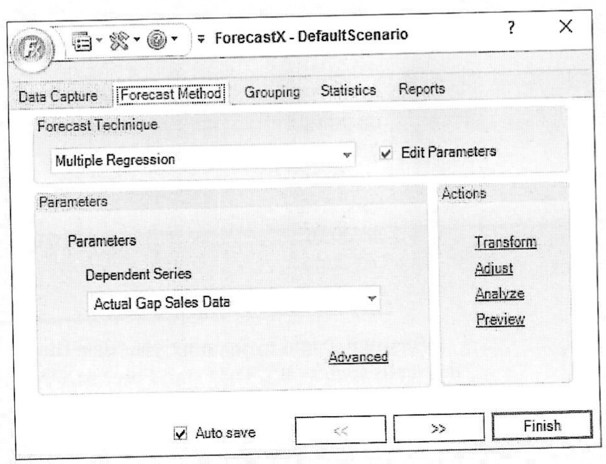

Source: John Galt Solutions

Click on **Advanced**. The following dialog box will appear.

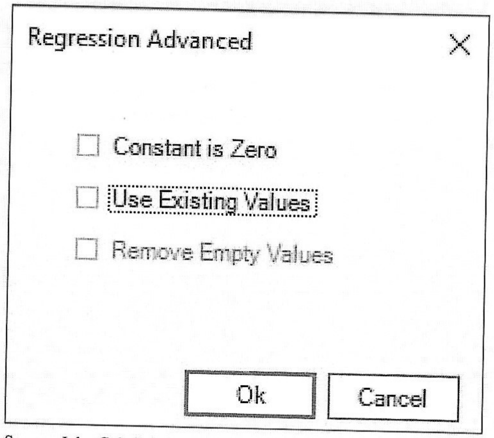

Source: John Galt Solutions

Make sure none of the boxes is checked. Click on **Ok** to return to the **Forecast Method** tab.

Then click the **Statistics** tab. In this dialog box, select the statistics that you desire. Here you are not particularly interested in the summary statistics. What you are looking for in this regression is whether the intercept is significantly different than zero, so you will want to get the coefficient table with P-values. That will be the next step.

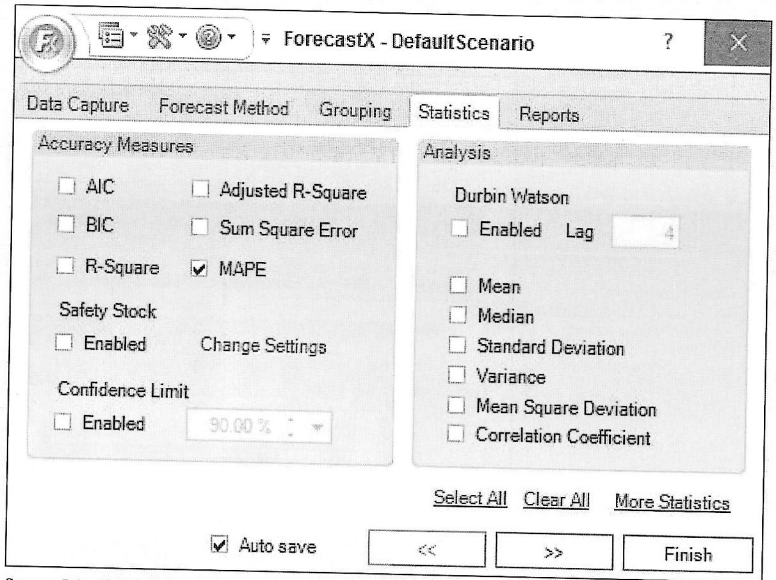

Source: John Galt Solutions

Remember that there are additional choices if you click the **More Statistics** button near the bottom right corner. In this dialog box, select **Regression** and check **Coefficient table** and **P-value**. Click **Ok,** and you will return to the **Statistics** tab.

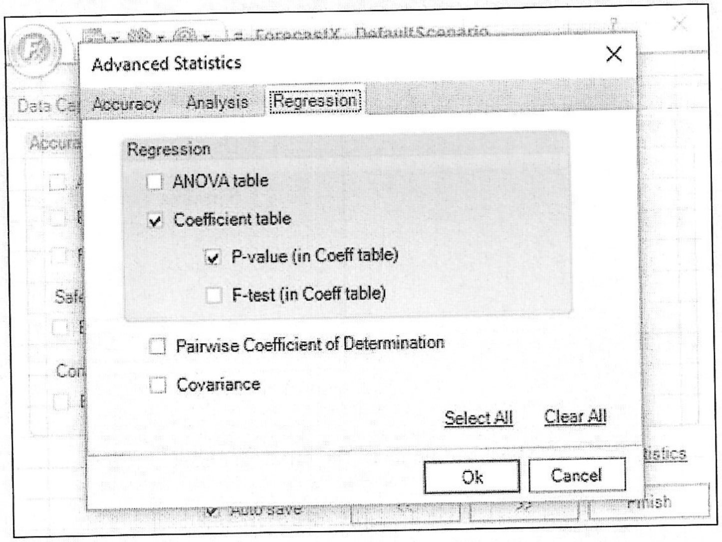

Source: John Galt Solutions

Now you are ready to click the **Reports** tab. In the **Reports** dialog box, for this application select the **Audit** tab. Then click **Finish**. The results will allow you to evaluate whether or not the constant term is significantly different than zero.

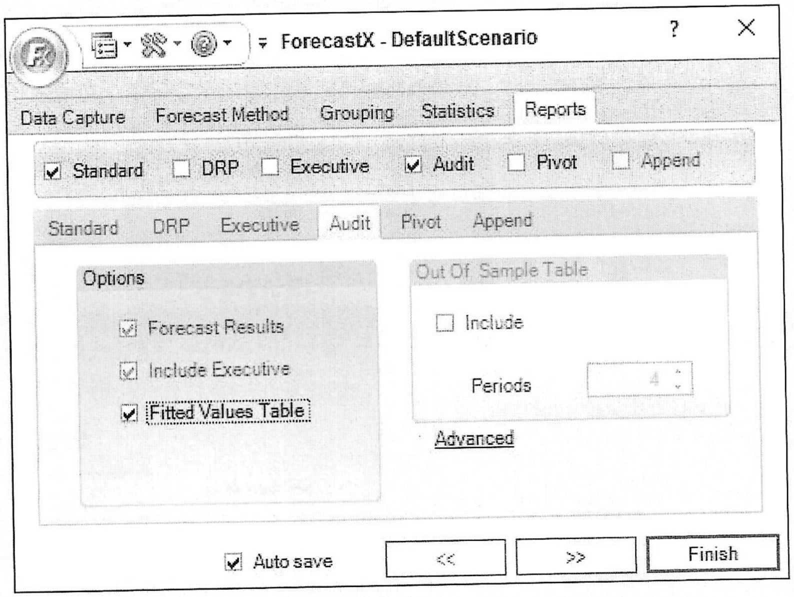

Source: John Galt Solutions

If the P-value for the intercept (or constant term or, in ForecastX™, the name of the dependent variable) is greater than 0.05, you know that the intercept is essentially zero. In this example, you see the P-value is 0.47.

Audit Trail -- Coefficient Table (Multiple Regression Selected)					
Series Description	Included in model	Coefficient	Standard error	T-test	P-value
Actual Gap Sales Data	Dependent	1,156.65	1,584.70	0.73	0.47
Winters'	Yes	2.01	1.04	1.94	0.06
Regression	Yes	-1.40	1.15	-1.22	0.23

Source: John Galt Solutions

If the P-value is less than 0.05, you would not want to do the combination because it would create a biased forecast. If the P-value is greater than 0.05, you may continue with the combining process.

To continue with a combination, redo the regression forcing the equation through the origin. To do this, in the **Method Selection** screen, click the **Advanced** button at the bottom. This time in the Regression Advanced box, check all three boxes. The regression coefficients in the resulting model are the optimum weights for the combined forecast, and the results provided by ForecastX™ are the combined forecast values.

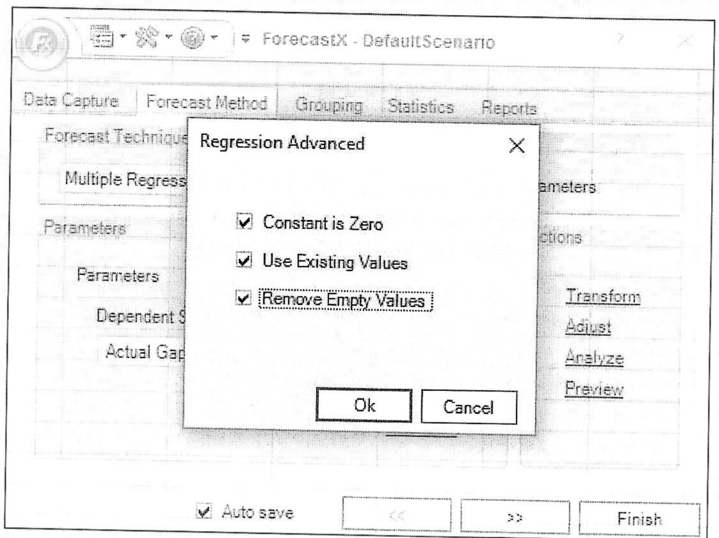

Source: John Galt Solutions

What happens when each of the above three boxes is checked? **First**, when the "Constant is Zero" box is checked, it tells ForecastX™ to do the regression model but to force the intercept (constant) to be zero. **Second**, when "Use Existing Values" is checked, it tells ForecastX™ to use the values of the independent variables that are already in the Excel sheet, rather than estimating those values with some extrapolation forecast method. Near the start of this discussion about using ForecastX™ to combine forecasts, you saw an abbreviated data sheet. Look at the last four rows, and you will see that values are already there for both the Winters' and regression forecasts. **Third**, "Remove Empty Values" tells ForecastX™ not to consider any empty cells, such as the last four cells in column A of this example.

It is a good idea when doing multiple regression forecasts to request both the Audit and the Standard Report, as shown below:.

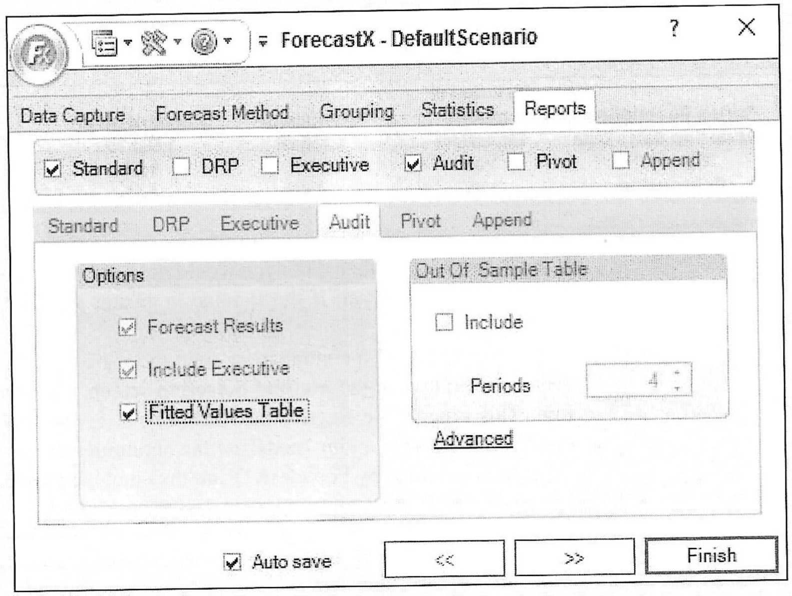

Source: John Galt Solutions

Let us now look at some results. Below is the graph and equation from the **Audit Report**. You see that the actual values are shown to fall to zero during the four-quarter forecast horizon. This is because when developing the graph, ForecastX™ **does** treat the empty cells as having zeros. That can be fixed in this graph, but it is cumbersome to do.

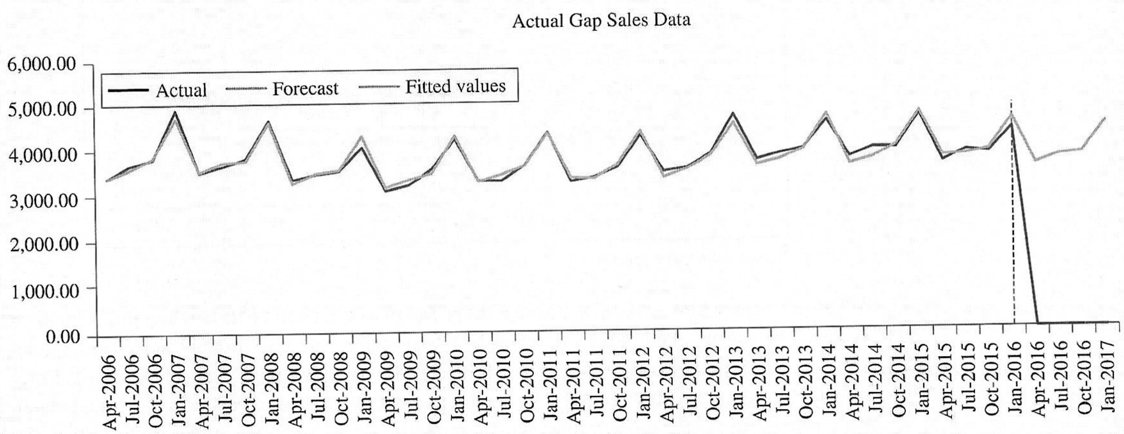

Now let us look at the result from the **Standard Report.** First is the graph and output as you will get them from ForecastX™.

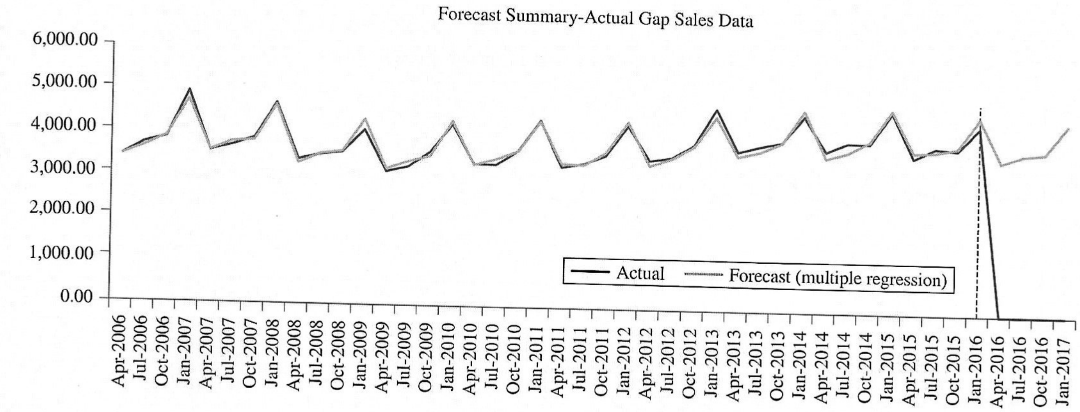

No Filter | Reconcile | Hide | << | < | Actual Gape Sales Data | Current Serie | > | >>

Actual Gap Sales Data = 0 + ((Winters') * 0.544645) + ((Regression) * 0.454242)

Title0 Type	Actual Gap Sales Data Adjust	Actual Gap Sales Data Forecast (Multiple Regression)	Winters' Adjust	Winters' Forecast	Regression Adjust	Regression Forecast
Apr-2016	0.00	*3,584.69*	3,573.47	*3,573.47*	3,606.92	*3,606.92*
Jul-2016	0.00	*3,763.40*	3,762.95	*3,762.95*	3,773.15	*3,773.15*
Oct-2016	0.00	*3,818.51*	3,746.51	*3,746.51*	3,914.18	*3,914.18*
Jan-2017	0.00	*4,476.68*	4,385.00	*4,385.00*	4,597.56	*4,597.56*
MAPE	2.17% 1.59	2.17% 1.59	100.00% 0.00	100.00% 0.00	100.00% 0.00	100.00% 0.00

Source: John Galt Solutions

You see again that the actual values are shown to fall to zero, when in fact, we know they are missing (not zero). We will fix that problem soon. Before we "fix" the graph, look at the values for the Winters' and regression forecasts during the four-quarter forecast horizon. They are exactly as shown in the abbreviated data table at the start of this section. We see that ForecastX™ did use the values provided because the **Use Existing Values** box was checked.

Now, how can we "fix" the graph? It is easy to get rid of the actual line that falls to zero. In the above output, simply delete the zeros for the forecast variable (Gap sales) during the forecast horizon. The result will look as follows:

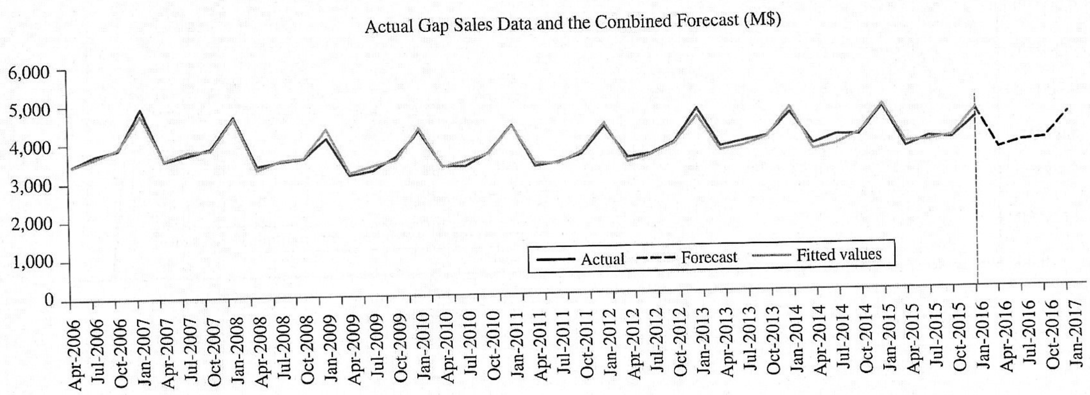

No Filter | Reconcile | Hide | << | < | Actual Gape Sales Data | Current Serie | > | >>

Actual Gap Sales Data = 0 + ((Winters') * 0.544645) + ((Regression) * 0.454242)

Title0 Type	Actual Gap Sales Data Adjust	Actual Gap Sales Data Forecast (Multiple Regression)	Winters' Adjust	Winters' Forecast	Regression Adjust	Regression Forecast
Apr-2016	0.00	*3,584.69*	3,573.47	*3,573.47*	3,606.92	*3,606.92*
Jul-2016	0.00	*3,763.40*	3,762.95	*3,762.95*	3,773.15	*3,773.15*
Oct-2016	0.00	*3,818.51*	3,746.51	*3,746.51*	3,914.18	*3,914.18*
Jan-2017	0.00	*4,476.68*	4,385.00	*4,385.00*	4,597.56	*4,597.56*
MAPE	2.17% 1.59	2.17% 1.59	100.00% 0.00	100.00% 0.00	100.00% 0.00	100.00% 0.00

Source: John Galt Solutions

Here you see that the lines have been reformatted also. This can often help make the lines more clearly different. The graph title was also edited to better reflect what the graph shows. Since all the ForecastX™ output is in Excel sheets, doing this kind of editing is not difficult.

Suggested Readings

Armstrong, J. Scott. "Combining Forecasts: The End of the Beginning or the Beginning of the End?" *International Journal of Forecasting* 5, no. 4 (1989), pp. 585–88.

———. *Long-Range Forecasting from Crystal Ball to Computer.* 2nd ed. New York: John Wiley & Sons, 1985.

Bates, J. M.; and C. W. J. Granger. "The Combination of Forecasts." *Operational Research Quarterly* 20, no. 4 (1969), pp. 451–68.

Bessler, David A.; and Jon A. Brandt. "Composite Forecasting: An Application with U.S. Hog Prices." *American Journal of Agricultural Economics* 63 (1981), pp. 135–40.

Chase, Charles W., Jr. "Composite Forecasting: Combining Forecasts for Improved Accuracy." *Journal of Business Forecasting* 19, no. 2 (Summer 2000), pp. 2, 20–22.

Clemen, Robert T. "Combining Forecasts: A Review and Annotated Bibliography." *International Journal of Forecasting* 5, no. 4 (1989), pp. 559–83.

Collopy, Fred; and J. Scott Armstrong. "Expert Opinions about Extrapolation and the Mystery of the Overlooked Discontinuities." *International Journal of Forecasting* 8, no. 4 (December 1992), pp. 575–82.

Diebold, Francis X. "Forecast Combination and Encompassing: Reconciling Two Divergent Literatures." *International Journal of Forecasting* 5, no. 4 (1989), pp. 589–92.

Fair, Roy C. *Predicting Presidential Elections and Other Things*. Stanford, CA: Stanford University Press, 2002.

Flores, Benito E.; David L. Olson; and Christopher Wolfe. "Judgmental Adjustment of Forecasts: A Comparison of Methods." *International Journal of Forecasting* 7, no. 4 (1992), pp. 421–33.

Fullerton, Thomas M., Jr. "A Composite Approach to Forecasting State Government Revenues: Case Study of the Idaho Sales Tax." *International Journal of Forecasting* 5, no. 3 (1989), pp. 373–80.

Goodwin, Paul. "Connect or Combine? Mechanically Integrating Judgmental Forecasts with Statistical Methods." *International Journal of Forecasting* 16, no. 2 (April–June 2000), pp. 261–75.

Hayes, Samuel P., Jr. "The Predictive Ability of Voters." *Journal of Social Psychology* 7 (1936), pp. 183–91.

Hendry, David F.; and Michael P. Clements. "Pooling of Forecasts." *Econometrics Journal* 5 (2002), pp. 1–26.

Hogarth, Robin M. "On Combining Diagnostic 'Forecasts': Thoughts and Some Evidence." *International Journal of Forecasting* 5, no. 4 (1989), pp. 593–97.

Li, Fuchun; and Greg Tkacz. "Combining Forecasts with Nonparametric Kernel Regressions." *Studies in Nonlinear Dynamics & Econometrics* 8, no. 4 (2004), article 2, http://www.bepress.com/snde/vol8/iss4/art2.

Lobo, Gerald I. "Analysis and Comparison of Financial Analysts' Time Series, and Combined Forecasts of Annual Earnings." *Journal of Business Research* 24 (1992), pp. 269–80.

Mahmoud, Essam. "Combining Forecasts: Some Managerial Issues." *International Journal of Forecasting* 5, no. 4 (1989), pp. 599–600.

Makridakis, Spyros. "Why Combining Works." *International Journal of Forecasting* 5 (1989), pp. 601–603.

Nash, John C. and Mary M. Nash. *Practical Forecasting for Managers*. London: Oxford University Press, 2001.

Nelson, Charles R. "A Benchmark for the Accuracy of Econometric Forecasts of GNP." *Business Economics* 19, no. 3 (April 1984), pp. 52–58.

Swanson, N. R.; and T. Zeng. "Choosing among Competing Econometric Forecasts: Regression-Based Forecast Combining Model Selection." *Journal of Forecasting* 20 (2001), pp. 425–40.

Wilson, J. Holton; and Deborah Allison-Koerber. "Combining Subjective and Objective Forecasts Improve Results." *Journal of Business Forecasting* 11, no. 3 (1992), pp. 3–8.

Winkler, Robert L.; and Robert T. Clemen. "Multiple Experts vs. Multiple Methods: Combining Correlation Assessments." *Decision Analysis* 1, no. 3 (September 2004), pp. 167–76.

Exercises

1. Explain why a combined model might be better than any of the original contributing models. Could there be cases in which a combined model would show no gain in forecast accuracy over the original models? Give an example where this situation might be likely to occur.
2. Outline the process for combining forecast models explained in this appendix.
3. Develop an example to show how to set up a data file to apply regression analysis to combine forecasts.
4. Your company produces a favorite summertime food product, and you have been placed in charge of forecasting shipments of this product. The historical shipments data below represent your company's past experience with the product.

Date	Shipments	Date	Shipments
Apr-2014	13,838	Jun-2015	21,056
May-2014	15,137	Jul-2015	13,509
Jun-2014	23,713	Aug-2015	9,729
Jul-2014	17,141	Sep-2015	13,454
Aug-2014	7,107	Oct-2015	13,426
Sep-2014	9,225	Nov-2015	17,792
Oct-2014	10,950	Dec-2015	19,026
Nov-2014	14,752	Jan-2016	9,432
Dec-2014	18,871	Feb-2016	6,356
Jan-2015	11,329	Mar-2016	12,893
Feb-2015	6,555	Apr-2016	19,379
Mar-2015	9,335	May-2016	14,542
Apr-2015	10,845	Jun-2016	18,043
May-2015	15,185	Jul-2016	10,803

a. Since the data appear to have strong seasonality, estimate a Winters' model using ForecastX. Request the mean absolute percentage error.

b. You also have access to a survey of purchasers' intentions for your product. This information has been collected for some time, and it has proved to be quite accurate for predicting shipments in the past.

Date	Purchasers' Intention Survey	Date	Purchasers' Intention Survey
Apr-2014	139	Jun-2015	246
May-2014	150	Jul-2015	142
Jun-2014	262	Aug-2015	91
Jul-2014	172	Sep-2015	121
Aug-2014	76	Oct-2015	119
Sep-2014	97	Nov-2015	176
Oct-2014	78	Dec-2015	155
Nov-2014	146	Jan-2016	99
Dec-2014	176	Feb-2016	73
Jan-2015	104	Mar-2016	112
Feb-2015	63	Apr-2016	189
Mar-2015	93	May-2016	140
Apr-2015	117	Jun-2016	206
May-2015	149	Jul-2016	128

Use a regression model with shipments as a function of purchasers' intentions to make a second forecast using ForecastX. Again request the mean absolute percentage error for this model.

c. After checking for bias, combine the forecasts in parts (*a*) and (*b*), and determine if a combined model may forecast better than either single model based on its MAPE.

Chapter Six

©VLADGRIN/Getty Images

Explanatory Models 2. Time-Series Decomposition

The information provided by time-series decomposition is consistent with the way managers tend to look at data and often helps them to get a better handle on data movements by providing concrete measurements for factors that are otherwise not quantified.

Many business and economic time series contain underlying components that, when examined individually, can help the forecaster better understand data movements and therefore make better forecasts. As discussed in Chapter 2, these components include the long-term trend, seasonal fluctuations, cyclical movements, and irregular (or random) fluctuations. Time-series decomposition models can be used to identify such underlying components by breaking the series into its component parts and then reassembling the parts to construct a forecast.

These models are among the oldest forecasting techniques available and yet remain popular today. Their popularity is due primarily to three factors. First, in many situations, time-series decomposition models provide excellent forecasts. Second, these models are relatively easy to understand and to explain to forecast users. This enhances the likelihood that the forecasts will be correctly interpreted and properly used. Third, the information provided by time-series decomposition is consistent with the way managers tend to look at data and often helps them to get a better handle on data movements by providing concrete measurements for factors that are otherwise not quantified.

There are a number of different methods for decomposing a time series. The one we will use is usually referred to as *classical time-series decomposition* and involves the ratio-to-moving-average technique. The classical time-series decomposition model uses the concepts of moving averages presented in Chapter 3 and trend projections discussed in Chapter 4. It also accounts for seasonality in a multiplicative way that is similar to what you have seen in Winters' exponential smoothing and the way we used seasonal indices in earlier chapters.[1]

[1] Remember that you have also accounted for seasonality using dummy variables in regression models. That method uses additive factors rather than multiplicative ones to account for seasonal patterns.