

# Chapter 1

## A Forecasting Overview

---

### *In This Chapter*

- ▶ Knowing the different methods of forecasting
  - ▶ Arranging your data in an order Excel can use
  - ▶ Getting acquainted with the Analysis ToolPak
  - ▶ Going it alone
- 

**A** sales forecast is like a weather forecast: It's an educated guess at what the future will bring. You can forecast all sorts of things — poppy-seed sales, stock market futures, the weather — in all sorts of ways: You can make your own best guess; you can compile and composite other people's guesses; or you can forecast on the basis of wishful thinking.

Unfortunately, all three of these options are less than desirable. If you want to do better more often than you do worse, you need to take advantage of some better options. Lucky for you, there are different ways to forecast, ways that have proven their accuracy over and over. They're a little more time consuming than guessing, but in the long run I've spent more time explaining bad guesses than doing the forecasts right in the first place.

Microsoft Excel was originally developed as a spreadsheet application, suited to figuring payment amounts, interest rates, account balances, and so on. But as Microsoft added more and more functions — for example, `AVERAGE` and `TREND` and inventory-management stuff — Excel became more of a multipurpose analyst than a single-purpose calculator.

Excel has the tools you need to make forecasts, whether you want to prepare something quick and dirty (and who doesn't from time to time?) or something sophisticated enough for a boardroom presentation.

The tools are there. You just need to know which tool to choose and know how to use it. You need to know how to arrange data for the tool. And you need to know how to interpret what the tool tells you — whether that tool's a basic one or something more advanced.

## Understanding Excel Forecasts

If you want to forecast the future — next quarter's sales, for example — you need to get a handle on what's happened in the past. So you always start with what's called a *baseline* (that is, past history — how many poppy seeds a company sold last year, where the market futures wound up last month, what the temperature was today).

Unless you're going to just roll the dice and make a guess, you need a baseline for a forecast. Today follows yesterday. What happens tomorrow generally follows the pattern of what happened today, last week, last month, last quarter, last year. If you look at what's already happened, you're taking a solid step toward forecasting what's going to happen next. (Part I of this book talks about forecast baselines and why they work.)

An Excel forecast isn't any different from forecasts you make with a specialized program. But Excel *is* a very useful application for making sales forecasts, for a variety of reasons:

- ✔ **You often have sales history recorded in an Excel worksheet.** When you already keep your sales history in Excel, basing your forecast on the existing sales history is easy — you've already got your hands on it.
- ✔ **Excel's charting features make it much easier to visualize what's going on in your sales history and how that history defines your forecasts.**
- ✔ **Excel has tools (found in what's called the Analysis ToolPak) that make generating forecasts easier.** You still have to know what you're doing and what the tools are doing — you don't want to just jam the numbers through some analysis tool and take the result at face value, without understanding what the tool's up to. But that's what this book is here for.
- ✔ **You can take more control over how the forecast is created by skipping the Analysis ToolPak's forecasting tools and entering the formulas yourself.** As you get more experience with forecasting, you'll probably find yourself doing that more and more.

You can choose from several different forecasting methods, and it's here that judgment begins. The three most frequently used methods, in no special order, are moving averages, exponential smoothing, and regression.

### *Method #1: Moving averages*

Moving averages may be your best choice if you have no source of information other than sales history — but you *do* need to know your sales history. Later in this chapter, I show you more of the logic behind using moving averages. The underlying idea is that market forces push your sales up or down. By averaging your sales results from month to month, quarter to quarter, or

year to year, you can get a better idea of the longer-term trend that's influencing your sales results.

For example, you find the average sales results of the last three months of last year — October, November, and December. Then you find the average of the next three-month period — November, December, and January (and then December, January, and February; and so on). Now you're getting an idea of the general direction that your sales are taking.

## ***Method #2: Exponential smoothing***

Exponential smoothing is closely related to moving averages. Just as with moving averages, exponential smoothing uses past history to forecast the future. You use what happened last week, last month, and last year to forecast what will happen next week, next month, or next year.

The difference is that when you use smoothing, you take into account how bad your previous forecast was — that is, you admit that the forecast was a little screwed up. (Get used to that — it happens.) The nice thing about exponential smoothing is that you take the error in your last forecast and use that error, so you hope, to improve your next forecast.

If your last forecast was too low, exponential smoothing kicks your next forecast up. If your last forecast was too high, exponential smoothing kicks the next one down.

The basic idea is that exponential smoothing corrects your next forecast in a way that would have made your *prior* forecast a better one. That's a good idea, and it usually works well.

## ***Method #3: Regression***

When you use regression to make a forecast, you're relying on one variable to predict another. For example, when the Federal Reserve raises short-term interest rates, you might rely on that variable to forecast what's going to happen to bond prices or the cost of mortgages. In contrast to moving averages or exponential smoothing, regression relies on a *different* variable to tell you what's likely to happen next — something other than your own sales history.

## ***Getting the Data Ready***

Which method of forecasting you use does make a difference, but regardless of your choice, in Excel you have to set up your baseline data in a particular

way. Excel prefers it if your data is in the form of a *list*. In Part II, I fill you in on how to arrange your data so that it best feeds your forecasts, but following is a quick overview.



## Using lists

There's nothing mysterious about an Excel list. A list is something very much like a database. Your Excel worksheet has columns and rows, and if you put a list there, you just need to manage three requirements:

- ✓ **Keep different variables in different columns.** For example, you can put sales dates in one column, sales amounts in another column, sales reps' names in another, product lines in yet another.
- ✓ **Keep different records in different rows.** When it comes to recording sales information, keep different sales records in different rows. Put information about a sale that was made on January 15 in one row, and information about a sale made on January 16 in a different row.
- ✓ **Put the names of the variables in the list's first row.** For example, you might put "Sales Date" in column A, "Revenue" in column B, "Sales Rep" in column C, and "Product" in column D.

Figure 1-1 shows a typical Excel list.

Why bother with lists? Because many Excel tools, including the ones you use to make forecasts, rely on lists. Charts — which help you visualize what's going on with your sales — rely on lists. Pivot tables — which are the most powerful way you have for summarizing your sales results in Excel — rely heavily on lists. The Analysis ToolPak — a very useful way of making forecasts — relies on lists, too.

You find a lot more about creating and using lists in Chapter 6. In the meantime, just keep in mind that a list has different variables in different columns, and different records in different rows.

## Ordering your data

"Ordering your data" may sound a little like "coloring inside the lines." The deal is that you have to tell Excel how much you sold in 1999, and then how much in 2000, and in 2001, and so on. If you're going to do that, you have to put the data in chronological order.

The screenshot shows a Microsoft Excel spreadsheet titled "Sales.xls" with a menu bar (File, Edit, View, Insert, Format, Tools, Data, Window, Help) and a toolbar. The spreadsheet contains a table with the following data:

	A	B	C	D	E	F	G	H	I	J	K	L	M
1	Sales Date	Revenue	Sales Rep	Product									
2	10/6/2005	\$ 5,150.38	Wilson	Software									
3	10/17/2005	\$ 2,739.29	Anderson	Hardware									
4	10/16/2005	\$ 7,524.88	James	Software									
5	10/7/2005	\$ 6,127.69	Anderson	Services									
6	10/13/2005	\$ 6,625.87	Wilson	Maintenance									
7	10/18/2005	\$ 6,762.55	Turgidson	Maintenance									
8	10/21/2005	\$ 2,227.09	Anderson	Software									
9	10/8/2005	\$ 9,136.42	Jones	Software									
10	10/18/2005	\$ 4,122.04	Coulter	Software									
11	10/10/2005	\$ 6,141.19	Anderson	Maintenance									
12	10/22/2005	\$ 3,057.58	Jones	Hardware									
13	10/13/2005	\$ 6,695.79	Tafoya	Software									
14	10/10/2005	\$ 1,066.58	Tafoya	Hardware									
15	10/21/2005	\$ 7,781.21	Turgidson	Warranty									
16	10/15/2005	\$ 5,571.73	Coulter	Hardware									
17	10/7/2005	\$ 6,147.57	Turgidson	Hardware									
18	10/20/2005	\$ 4,477.92	Wilson	Hardware									
19	10/5/2005	\$ 7,603.13	James	Maintenance									
20	10/14/2005	\$ 7,900.95	Turgidson	Services									
21	10/19/2005	\$ 3,627.69	James	Services									
22	10/11/2005	\$ 7,307.49	Coulter	Warranty									
23	10/11/2005	\$ 4,690.90	Turgidson	Software									
24	10/22/2005	\$ 1,551.35	Coulter	Services									
25	10/16/2005	\$ 1,113.69	Wilson	Warranty									
26	10/14/2005	\$ 7,447.74	Anderson	Warranty									
27	10/17/2005	\$ 3,353.07	Tafoya	Services									
28	10/12/2005	\$ 5,369.40	James	Hardware									
29	10/12/2005	\$ 6,401.29	Jones	Services									
30	10/8/2005	\$ 1,707.77	Coulter	Maintenance									
31	10/20/2005	\$ 5,252.47	Tafoya	Maintenance									
32	10/15/2005	\$ 6,204.11	Jones	Maintenance									
33	10/9/2005	\$ 4,202.45	Wilson	Services									
34	10/9/2005	\$ 1,711.89	James	Warranty									
35	10/19/2005	\$ 6,192.25	Jones	Warranty									

**Figure 1-1:**  
You don't have to keep the records in date order — you can handle that later.

The best — and I mean the *best* — way to put your data in chronological order in Excel is by way of pivot tables. A pivot table takes individual records that are in a list (or in an external database) and combines the records in ways that you control. You may have a list showing a year's worth of sales, including the name of the sales rep, the product sold, the date of sale, and the sales revenue. If so, you can very quickly create a pivot table that totals sales revenue by sales rep and by product across quarters. Using pivot tables, you can summarize tens of thousands of records, quite literally within seconds. If you haven't used pivot tables before, this book not only introduces the subject but also makes you dream about them in the middle of the night.

There are a couple of wonderful things about pivot tables:

- ✔ **They can accumulate for you all your sales data — or, for that matter, your data on the solar wind, but this book is about sales forecasting.** If you gather information on a sale-by-sale basis, and you then want to know how much your reps sold on a given day, in a given week, and so on, a pivot table is the best way to do so.
- ✔ **You can use a pivot table as the basis for your next forecast, which saves you a bunch of time.**
- ✔ **They have a unique way of helping you group your historical data — by day, by week, by month, by quarter, by year, you name it.** Chapter 8 gives you much more information on pivot tables, including troubleshooting some common problems.

## Making Basic Forecasts

Part III gets into the business of making actual forecasts, ones that are based on historical data (that is, what's gone on before). You see how to use the Analysis ToolPak to make forecasts that you can back up with actuals — given that you've looked at Part II and set up your actuals correctly. (Your *actuals* are the actual sales results that show up in the company's accounting records — say, when the company recognizes the revenue.)

The Analysis ToolPak (often abbreviated ATP) is a gizmo that has shipped with Excel ever since 1995. The ATP is a convenient way to make forecasts, as well as to do general data analysis. The three principal tools Excel's ATP gives you to make forecasts are:

- ✓ Moving Averages
- ✓ Exponential Smoothing
- ✓ Regression



### All the news that's fit to print — and then some

You may not be familiar with newsgroups, and if you're not, I think you should be. **Newsgroups** are what used to be called bulletin boards before the Internet and the Web came along. They're places where you can post questions (for example, about Excel) and a little later on read answers that other users have posted. (Unfortunately, because newsgroups are largely public, you also see a lot of garbage.) One good way to get familiar with newsgroups is to use Google:

1. **Go to `www.google.com`.**
2. **Click on Groups, and then click Advanced Groups Search.**

The Advanced Groups Search link is in really small print to the right of the Search Groups button.

3. **Enter `Microsoft.Public.Excel.*` in the Return Only Messages from the Group box.**

The asterisk returns all public newsgroups beginning with `Microsoft.Public.Excel`.

4. **Enter a word or phrase you want to look for in one of the Find Messages boxes.**

For example, you can type **add-in** to search specifically for newsgroups about Excel add-ins.

5. **Click the Google Search button.**

You'll see all the questions — and all the answers — concerning the word or phrase you entered that have shown up in those newsgroups.

Newsgroups are a great way to find out about Excel — or, for that matter, any other application.

Those are the three principal forecasting methods, and they form the basis for the more-advanced techniques and models. So it's no coincidence that these tools have the same names as the forecasting methods mentioned earlier in this chapter.



The Analysis ToolPak is an *add-in*. An add-in does tasks, like forecasting, on your behalf. An add-in is much like the other tools that are a part of Excel — the difference is that you can choose whether to install an add-in. For example, you can't choose whether the Goal Seek tool (Tools ⇨ Goal Seek) is available to you. If you decide to install Excel on your computer, Goal Seek is just part of the package. Add-ins are different. You can decide whether to install them. When you're installing Excel — and in most cases this means when you're installing Microsoft Office — you get to decide which add-ins you want to use.

The following sections offer a brief introduction to the three ATP tools.



Given a good baseline, the ATP can turn a forecast back to you. And then you're responsible for evaluating the forecast, for deciding whether it's a credible one, for thinking the forecast over in terms of what you know about your business model. After all, Excel just calculates — you're expected to do the thinking.

## Putting moving averages to work for you

You may already be familiar with moving averages. They have two main characteristics, as the name makes clear:

- ✓ **They move.** More specifically, they move over time. The first moving average may involve Monday, Tuesday, and Wednesday; in that case, the second moving average would involve Tuesday, Wednesday, and Thursday; the third Wednesday, Thursday, and Friday, and so on.
- ✓ **They're averages.** The first moving average may be the average of Monday's, Tuesday's, and Wednesday's sales. Then the second moving average would be the average of Tuesday's, Wednesday's, and Thursday's sales, and so on.

The basic idea, as with all forecasting methods, is that something regular and predictable is going on — often called the *signal*. Sales of ski boots regularly rise during the fall and winter, and predictably fall during the spring and summer. Beer sales regularly rise on NFL Sundays and predictably fall on other days of the week.

But something else is going on, something irregular and unpredictable — often called *noise*. If a local sporting goods store has a sale on, discounting ski boots from May through July, you and your friends may buy new boots during the spring and summer, even though the regular sales pattern (the

signal) says that people buy boots during the fall and winter. As a forecaster, you can't predict this special sale. It's random and tends to depend on things like overstock. It's noise.

Let's say you run a liquor store, and a Thursday night college football game that looked like it would be the Boring Game of the Week when you were scheduling your purchases in September has suddenly in November turned into one with championship implications. You may be caught short if you scheduled your purchases to arrive at your store the following Saturday, when the signal in the baseline leads you to expect your sales to peak. That's *noise* — the difference between what you predict and what actually happens. By definition, noise is unpredictable, and for a forecaster it's a pain.

If the noise is random, it averages out. Some months, your stores will be discounting ski boots for less than the cost of an arthroscopy. Some months, a new and really cool model will come on line, and they'll be taking every possible advantage. The peaks and valleys even out. Some weeks there will be an extra game or two and you'll sell (and therefore need) more bottles of beer. Some weeks there'll be a dry spell from Monday through Friday, you won't need so much beer, and you won't want to bear the carrying costs of beer you're not going to sell for a while.



The idea is that the noise averages out, and that what moving averages show you is the signal. To misquote Johnny Mercer, if you accentuate the signal and eliminate the noise, you latch on to a pretty good forecast.

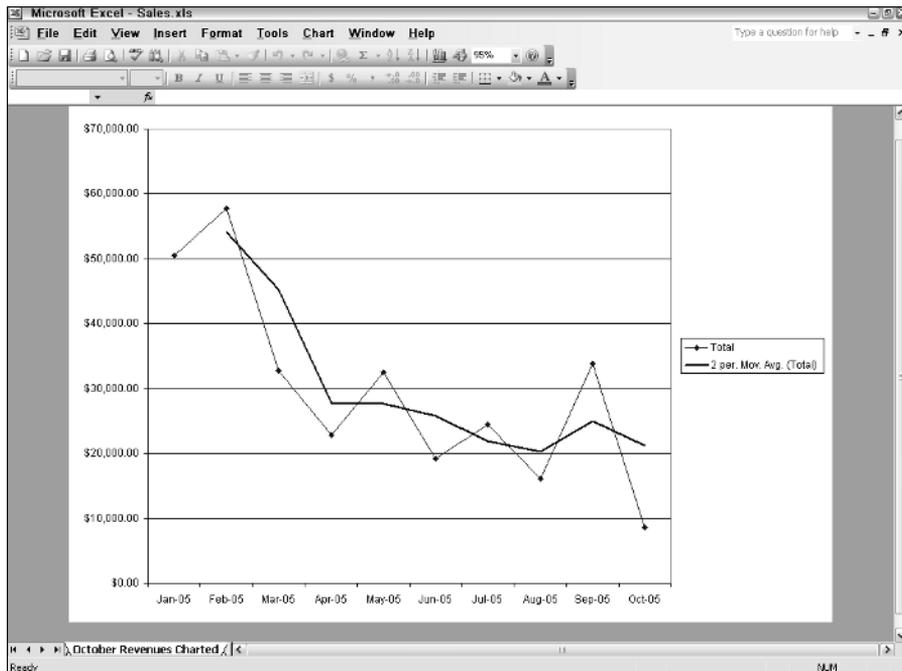
So with moving averages, you take account of the signal — the fact that you sell more ski boots during certain months and fewer during other months, or that you sell more beer on weekends than on weekdays. At the same time you want to let the random noises — also termed *errors* — cancel one another out. You do that by averaging what's already happened in two, three, four, or more previous consecutive months. The signal in those months is emphasized by the averaging, and that averaging also tends to minimize the noise.

Suppose you decide to base your moving averages on two-month records. That is, you'll average January and February, and then February and March, and then March and April, and so on. So you're getting a handle on the signal by averaging two consecutive months and reducing the noise at the same time. Then, if you want to forecast what will happen in May, you hope to be able to use the signal — that is, the average of what's happened in March and April.

Figure 1-2 shows an example of the monthly sales results and of the two-month moving average.

Chapter 14 goes into more detail about using moving averages for forecasting.

**Figure 1-2:** The moving average shows the general direction of the sales (the signal), and deemphasizes the random variations (the noise).



## Making sense of exponential smoothing

I know, the term *exponential smoothing* sounds intimidating and pretentious. I guess it's both — although I'm not responsible for it, I swear. (If you really want, you can find out why it's called that in Chapter 15.) In any event, don't worry about what it's called — it's just a kind of self-correcting moving average.

Suppose that in June, you forecast \$100,000 in sales for July. When the July sales results are in, you find that your July forecast of \$100,000 was \$25,000 too low — you actually made \$125,000 in sales. Now you need to forecast your sales for August. The idea behind this approach to forecasting is to adjust your August forecast in a way that would have made the *July* forecast more accurate. That is, because your July forecast was too low, you increase your August forecast above what it would have been otherwise.

More generally:

- ✓ If your most recent forecast turned out to be an underestimate, you adjust your next forecast upward.
- ✓ If your most recent forecast turned out to be an overestimate, you adjust your next forecast downward.

You don't make these adjustments just by guessing. There are formulas that help out, and the ATP's Exponential Smoothing tool can enter the formulas for you. Or you can roll your own formulas if you want. Turn to Chapter 15 to see how to do that.

Figure 1-3 shows what you would forecast if your prior forecast (for July) was too low — then you boost your forecast for August.

And if your prior, July forecast was too high (you optimist!), you cool your jets a little bit in your August forecast, as shown in Figure 1-4.

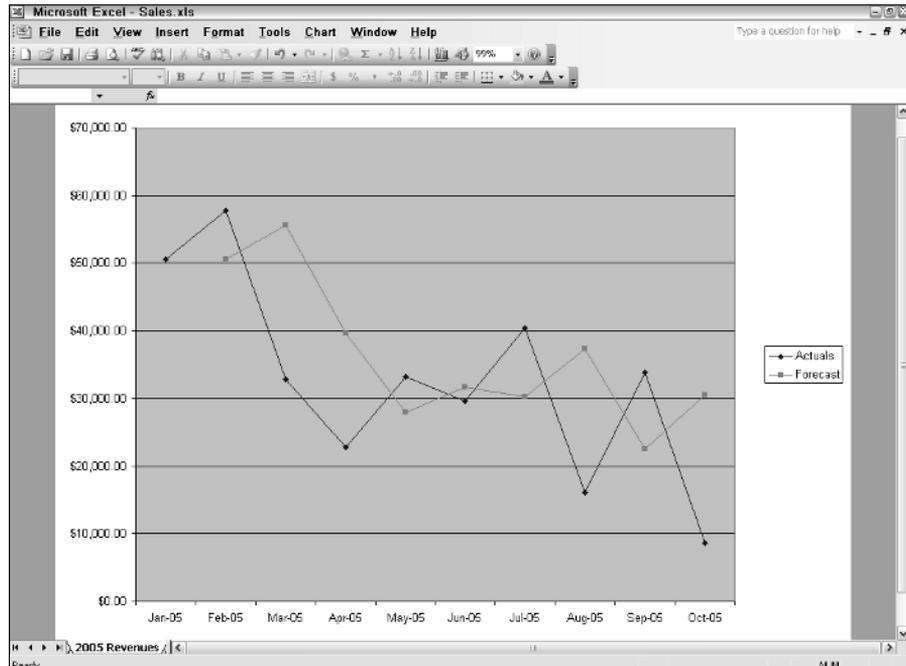
## Using regression to get what you want

The term *regression* doesn't sound as bad as *exponential smoothing*, but it is — I admit — more complicated, at least in terms of the math.

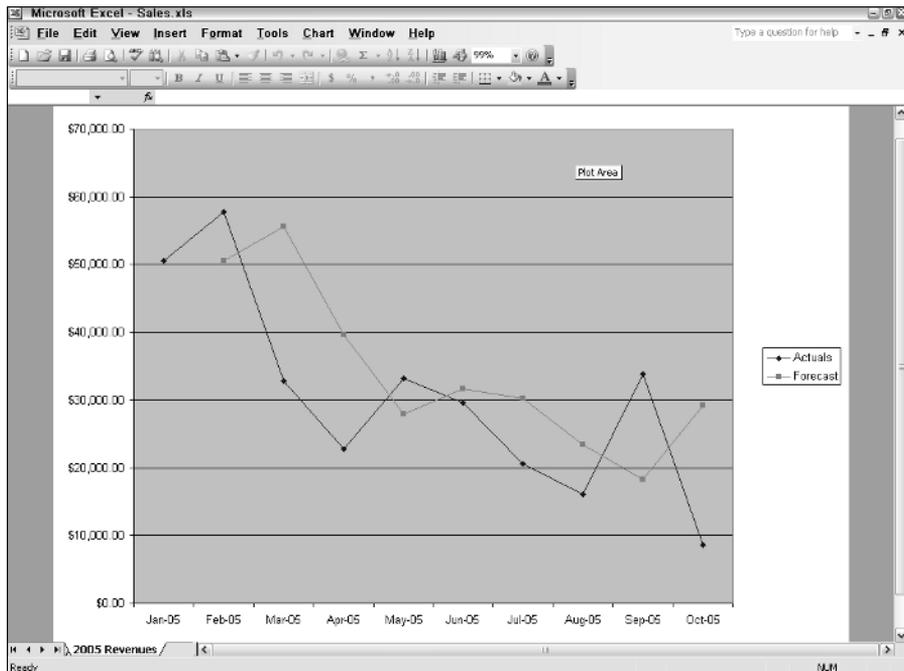
And that's why the ATP is convenient. It takes responsibility for the math, just as it does with moving averages and exponential smoothing. **Remember:** You still have to give a good baseline to the ATP to get accurate results.

Here's a quick look at forecasting with regression. (You can find a more detailed look in Chapter 11.)

**Figure 1-3:** Here's what happens if your forecast for July was an underestimate. Notice that the August forecast is kicked up.



**Figure 1-4:**  
Your prior  
forecast  
was too  
high, so  
exponential  
smoothing  
makes you  
back off  
your next  
forecast.



The idea behind regression is that one variable has a relationship with another variable. When you're a kid, for example, your height tends to have a relationship to your age. So if you want to forecast how tall you'll be next year — at least, until you quit growing — you can check how old you'll be next year.

Of course, people differ. When they're 15 years old, some people are 5 feet tall, some are 6 feet tall. On average, though, you can forecast with some confidence how tall someone will be at age 15. (And you can almost certainly forecast that a newborn kidlet is going to be under 2 feet tall, even wearing lifts.)

The same holds true with sales forecasting. Suppose your company sells consumer products. It's a good bet that the more advertising you do, the more you'll sell. At least it's worth checking out whether there's a relationship between the size of your advertising budget and the size of your sales revenue. If you find that there's a dependable relationship — and if you know how much your company is willing to spend on advertising — you're in a good position to forecast your sales.

Or suppose your company markets a specialty product, such as fire doors. (A *fire door* is one that's supposed to be resistant to fire for some period of time,

and there are a lot of them in office buildings.) Unlike consumer products, something such as a fire door doesn't have to be a particular off-the-shelf color or have a fresher-than-fresh aroma. If you're buying fire doors, you want to get the ones that meet the specs and are the cheapest.

So if you're selling fire doors, as long as your product meets the specs, you'd want to have a look at the relationship between the price of fire doors and how many are sold. Then you check with your marketing department to find out how much they want you to charge per door, and you can make your forecast accordingly.



The point is that more often than not you can find a dependable relationship between one variable (advertising dollars or unit price) and another (usually, sales revenue or units sold).

You use Excel's tools to quantify that relationship. In the case of regression forecasts, you give Excel a couple of baselines. To continue the examples we've used so far in this section:

- ✓ Historical advertising expenses and historical sales revenues
- ✓ How much you charged per fire door and how many doors you sold

If you give Excel good baselines, it will come back to you with a formula.

- ✓ Excel will give you a number to multiply times how much you expect to spend on advertising, and the result will be your expected sales revenue.
- ✓ Or, Excel will give you a number to multiply times the unit cost per door, and the result will be the number of doors you can expect to sell.



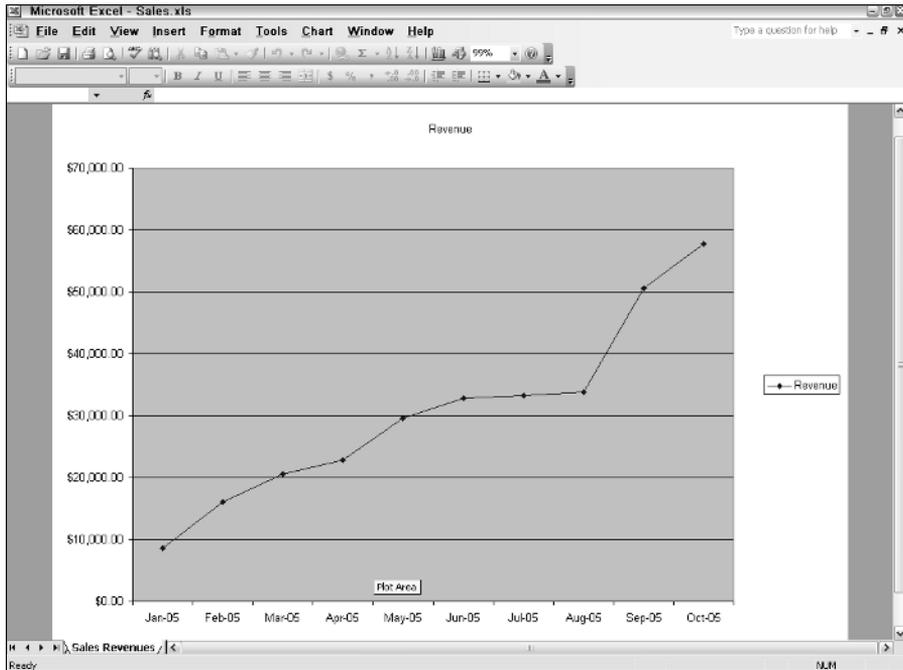
It's just a touch more complicated than that. Excel also gives you a number, called a *constant*, that you need to add to the result of the multiplication. But as Chapter 11 shows, you can get Excel to do that for you.

## Charting Your Data

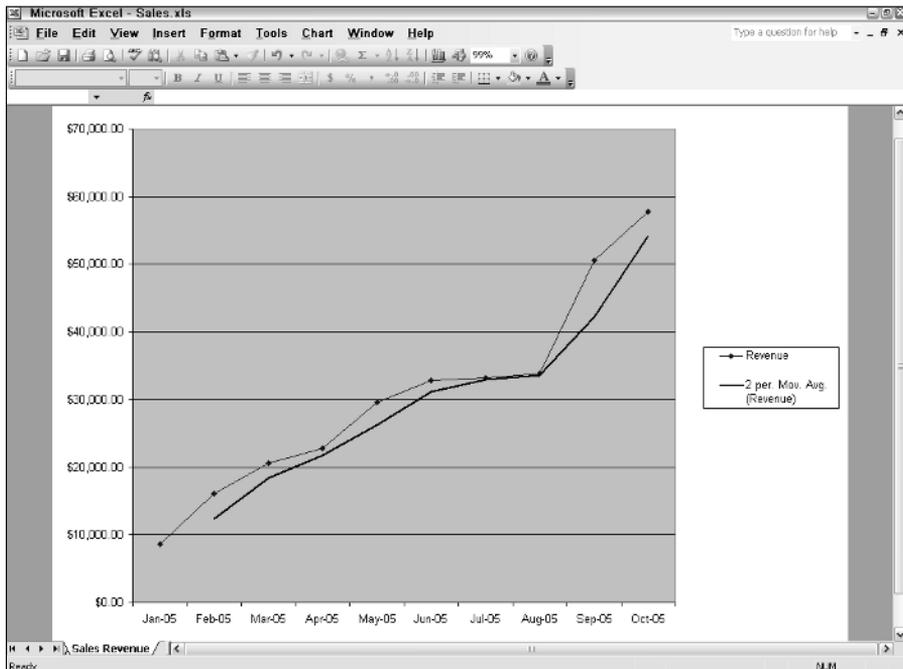
I've been doing this stuff for a long time, and I can't tell you how critical it is to chart your baseline and your forecast. Being able to visualize what's going on is important for several reasons.

Using Excel's charts, you can see how your actuals are doing (see Figure 1-5). And by charting your actuals, you can see how well your sales forecasts do against the actual sales results. Figure 1-6 shows a forecast that's based on moving averages, against the monthly actuals.

**Figure 1-5:**  
An Excel chart makes it much easier to see how your sales are doing.



**Figure 1-6:**  
Notice how the moving average lags behind the actual results.



## A word or two about Excel versions

Pivot tables were first introduced in Excel 95. This book has a good bit to say about basing charts on pivot tables, much of which has been made redundant by Excel 2000 and its pivot charts — beginning with Excel 2000, you can create not only pivot tables but also pivot charts.

Many companies continue to use Excel 97, though, and for good reason. The principal argument for upgrading to Excel 2000, 2002, or 2003 is that *other* software might rely on Excel to handle calculations, and that other software requires a

version of Excel that's more recent than 1997. A company that doesn't have to do the upgrade, though, usually doesn't — the cost-benefit equation doesn't say "Upgrade now!"

So, if like many others your company still uses Excel 97, you'll want to know how to create charts based on pivot tables. If your company uses Excel 2000 or later, Chapter 9 shows you how to create pivot charts, as well as how to create standard charts that are based on pivot tables.

By charting your baseline and your forecasts, you can:

- ✔ **See how your actual results are doing.** A chart is almost always more revealing than a table of numbers.
- ✔ **See how well your forecasts predict actual results.** Your eye is a good gauge of the quality of your forecasts.
- ✔ **See how well a different variable — advertising dollars or the Consumer Price Index — predicts the sales of your product.**

Yes, an R squared or some other summary statistic can give you a concise estimate of how well your forecasts are working. But there's nothing, *nothing*, like a chart to tell you if you're forecasting results or if you're forecasting junk. Chapter 9 shows you how to set up charts with Excel.

## Forecasting with Advanced Tools

There's a lot to be said for using the Analysis ToolPak to create your forecasts. The ATP is quick, it does the heavy lifting for you, and it's pretty comprehensive, taking care of the math and some of the charting.

But there's nothing like doing it yourself. When you wave goodbye to the ATP, you establish and maintain control over what's going on with the forecast. If you have formulas in your worksheet cells — formulas that support your forecasts — you can change those formulas as your forecasting needs change. And you can change — or add to — the baseline and immediately see what the effect doing so has on your forecast.

When the ATP does not give you formulas but static values instead, you can't experiment with the forecasts or see the effect of modifying the baseline. And the ATP's Regression tool gives you just the static values. The Exponential Smoothing tool is a little better, but it mixes formulas with static values. And the Moving Averages tool forces you to start over if you want to change the number of records in the baseline that make up a moving average.

Suppose that you have the number 3 in cell A1 and the number 5 in cell A2. In cell A3 you can enter the sum of those two numbers, 8. But if you now change the number 3 in cell A1 to, say, 103, you still have 8 in A3. It's a constant — a number, not a formula. It doesn't react to what's in cell A1 or A2: You're still going to see the number 8 in cell A3.

On the other hand, suppose you have this in cell A3:

```
=A1 + A2
```

That's a formula, not a constant, and it tells Excel to add whatever's in A1 to whatever's in A2. So if you change what's in A1, or what's in A2, Excel recalculates the result and shows it — in this example — in A3.

The point to keep in mind is that the ATP's regression tool gives you numbers, not formulas. It calculates your forecast, and the underlying figures, and writes numbers onto your worksheet. That means, regardless of how you change the numbers in your baseline, you're still going to be looking at the same forecast as offered by the ATP.

But — and it's a big one — if you make the forecast yourself instead of relying on the ATP, you can enter the formulas that the ATP denies you. Why is this important? By entering the formulas yourself, you have more control over what's going on with the forecast.

Relying on the ATP, which is a good toolbox, and one that you can generally trust, is perfectly okay. However, if you enter formulas, ones that react to changes in your baseline, you can make a change in the baseline and see what happens to the forecast. You can change this month's result from \$100,000 to \$75,000 and see whether your forecast for next month changes substantially. You can't do that with the ATP's Regression tool, because it doesn't give you formulas. To a smaller degree, the same is true of the Exponential Smoothing tool.

But the more important reason, the reason for you to consider entering the formulas yourself, is that you're relying on your own knowledge of how and why forecasting works. In Part IV, I show you how to use functions like LINEST and TREND to do your regression-based forecasts. You also see how to use array formulas to get the most out of those Excel functions.

You don't need to enter all the formulas yourself to make good forecasts. The ATP is a good tool. But if you do enter the formulas yourself, not only can you be more confident that you know what's going on with your forecast, but you can also exercise more control over what your forecast says is going to happen. In a business as tricky and trappy as forecasting, the more control you have, the better.