Wayne L. Winston

Marketing Analytics

Data-Driven Techniques with Microsoft Excel

WILEY

Marketing Analytics

Data-Driven Techniques with Microsoft® Excel®

Wayne L. Winston

WILEY

To my wonderful family: Gregory, Jennifer, and Vivian

Credits

Executive Editor
Robert Elliott

Project Editor
Victoria Swider

Technical Editor
Lopo Rego

Production Editor
Daniel Scribner

Copy Editor
San Dee Phillips

Editorial Manager
Mary Beth Wakefield

Freelancer Editorial Manager
Rosemarie Graham

Associate Director of Marketing
David Mayhew

Marketing Manager
Ashley Zurcher

Business Manager
Amy Knies

Vice President and Executive Group Publisher
Richard Swadley

Associate Publisher
Jim Minatel

Project Coordinator, Cover
Katie Crocker

Proofreaders
Josh Chase, Word One
Louise Watson, Word One

Indexer
Ron Strauss

Cover Image
Wiley

Cover Designer
Ryan Sneed

About the Author

Wayne Winston is Professor Emeritus at the Indiana University Kelley School of Business and is currently a Visiting Professor at the University of Houston Bauer College of Business. Wayne has won more than 45 teaching awards at Indiana University. These awards include six school-wide MBA teaching awards. He has authored 25 reference journal articles and a dozen books including, *Operations Research*: *Applications and Algorithms* (Cengage, 1987), *Practical Management Science* (Cengage, 2011), *Data Analysis and Decision-Making* (Cengage, 2013), *Simulation Modeling with @RISK* (Cengage, 2004), *Mathletics* (Princeton, 2009), and *Excel 2013 Data Analysis and Business Modeling* (O'Reilly, 2014). Wayne has also developed two online courses for Harvard Business School: Spreadsheet Modeling, and Mathematics for Management. He has taught Excel modeling and consulted for many organizations including the U.S. Army, the U.S. Navy, Broadcom, Cisco, Intel, Pfizer, Eli Lilly, Ford, GM, PWC, Microsoft, IAC, Deloitte Consulting, Booz Allen Hamilton, QAS, eBay, the Dallas Mavericks, and the New York Knicks. Lastly, Wayne is a two-time *Jeopardy*! champion.

About the Technical Editor

Lopo Rego joined the Kelley School of Business at Indiana University in 2011 as an Associate Professor of Marketing. Trained in Economics, he "converted to the dark side" during his MBA and has since been interested in understanding the association between marketing strategy and firm performance. This proved to be a life-long quest, leading him to Ann Arbor where he eventually earned his Ph.D. in Marketing at the University of Michigan's Ross School of Business. Not surprisingly, his research interests focus primarily in understanding how marketing decisions, strategies, and investments translate into firm performance, be it at the product-marketplace level, financial-accounting level or shareholder wealth level. Additionally, Lopo is interested in marketing analytics, namely in developing and analyzing marketing metrics that drive firm performance. His research has been published in such outlets as the *Journal of Marketing*, *Marketing Science*, *European Journal of Marketing*, *Journal of Empirical Generalisations in Marketing*, *Harvard Business Review*, *Journal of Research in Marketing,* and *Marketing Science Institute Working Paper Series*.

Acknowledgments

Of all my books, this one was probably the hardest to write. Thanks to my wonderful wife Vivian who was so nice to me when I got frustrated during the authoring process. Wiley acquisitions editor Robert Elliott was always encouraging and his input was a great help in shaping the final product. Wiley project editor Victoria Swider did a great job in pushing me to become a better writer.

Lastly, I must give a special note of thanks to my technical editor, Associate Professor of Marketing at the Kelly School of Business, Lopo Rego. Lopo did an amazing job of suggesting alternative wording and catching errors. He went above and beyond his role as technical editor, and I am truly indebted to him for his Herculean efforts.

Contents

Contents

Introduction

In the last 20 years, the use of analytic techniques in marketing has greatly increased. In April 2013, *Forbes* magazine reported a 67-percent growth in marketing-related analytics hires during the previous year and an amazing 136-percent growth during the previous 3 years.

Given this growth of interest in marketing analytics and my love of Excel modeling, I decided in 2004 to create a 7-week MBA elective in marketing analytics (K509) at the Indiana University Kelley School of Business. Although there are several excellent advanced marketing analytics books. (I am partial to *Database Marketing* by Robert Blattberg, Byung-Do Kim, and Scott Neslin (Springer, 2008).) I could not find an Excel-based book that provided a how-to-do-it approach suitable for an MBA elective or an advanced undergraduate course. With no suitable book in hand, I wrote up course notes that I used in classes for 10 years. The course has been wildly successful with nearly 65 percent of all MBA's at the Kelley School taking the class. In May 2013, I was honored to receive the Eli Lilly MBA teaching award as the best teacher in the MBA program, primarily for teaching K509. In November 2011, Robert Elliott of Wiley Publishing approached me about turning my notes into a book, and this book is the result. In addition to being utilized in K509, portions of the book have been used to teach marketing analytics to senior managers at Deloitte consulting, Booz Allen Hamilton consulting, and 3M marketing analysts.

How This Book Is Organized

Since I started using Excel in classes in 1992, I have become a total convert to teaching by example. This book is no exception. Virtually every chapter's primary focus is to teach you the concepts through how-to examples. Each example has the following components:

- Step-by-step instructions
- A downloadable Excel file containing data and solutions
- Screenshots of various steps and sections of the Excel file for clarity

The downloadable Excel files provide complete solutions to the examples, but the instructions encourage you to follow along and work through them on your own. If you follow along using the provided Excel files, you can work in empty

cells alongside the completed solution and compare your result with the provided solution to ensure your success.

The book has been organized around 11 topical areas.

Part I: Using Excel to Summarize Marketing Data

This part of the book introduces the marketing analyst to many Excel tools that can be used to analyze marketing problems: PivotTables (Chapter 1), charting (Chapter 2), and Excel statistical functions (Chapter 3), including the incredibly useful COUNTIF, COUNTIFS, SUMIF, SUMIFS, AVERAGEIF, and AVERAGEIFS functions.

Part II: Pricing

The determination of a profit maximizing pricing strategy is always difficult. In this section you learn how to quickly estimate demand curves and use the Excel Solver (Chapter 4) to determine profit maximizing prices. The Excel Solver is then used to optimize price bundling (Chapter 5), nonlinear pricing strategies (Chapter 6), and price-skimming strategies (Chapter 7). A brief introduction to revenue management, also known as yield management (Chapter 8), is also included.

Part III: Forecasting

Businesses need accurate forecasts of future sales. Sales forecasts drive decisions involving production schedules, inventory management, manpower planning, and many other parts of the business. In this section you first learn about two of the most used forecasting tools: simple linear (Chapter 9) and simple multiple regression (Chapters 10 and 11). Then you learn how to estimate the trend and seasonal aspects of sales (Chapter 12) and generate forecasts using two common extrapolation forecasting methods: the Ratio to Moving Average method (Chapter 13), and Winter's Method for exponential smoothing (Chapter 14) with trend and seasonality. Then you learn about neural networks (Chapter 15), a form of artificial intelligence whose role in marketing forecasting is rapidly growing.

Part IV: What Do Customers Want?

Every brand manager wants to know how various product attributes drive the sales of a product. For example, what is most important in a consumer's choice of car: price, brand, engine horsepower, styling, or fuel economy? In this section you learn how conjoint analysis (Chapter 16) and discrete choice (Chapter 18) can be used to rank the importance of product attributes and also rank levels of product attributes. For example, what type of styling on an SUV is most preferred? You also

learn about the widely used tool of logistic regression (Chapter 17), which is used to estimate probabilities involving situations in which two, or binary, outcomes must be forecasted. For example, how a person's demographic information can be used to predict the chance that he will subscribe to a magazine.

Part V: Customer Value

Companies cannot make intelligent decisions on how to spend money acquiring customers unless they understand the value of their customers. After all, spending $400 to acquire a customer who will generate $300 in long-term profits is a sure recipe for going out of business. In this section you learn how to measure customer value (Chapter 19), value companies based on the customer value concept (Chapter 20), incorporate uncertainty in customer value models (Chapter 21), and use your understanding of customer value to optimally allocate resources (Chapter 22) between acquisition and retention of customers.

Part VI: Market Segmentation

No matter what product you sell, your market consists of different market segments. For example, in Chapter 23 you will use cluster analysis to show that that every U.S. city can be classified into one of four demographic segments. In Chapter 25 you will learn how classification trees can be used to segment a market. You are also introduced to the exciting concepts behind collaborative filtering (Chapter 24), which is the basis for Amazon.com and Netflix recommendations.

Part VII: Forecasting New Product Sales

With little or no history about sales of a product, it is difficult to predict future product sales. Given a few data points, S curves (Chapter 26) can be used to predict future product sales. The famous Bass diffusion model (Chapter 27) explains how sales of products evolve over time and can be used to predict product sales even before a product comes to the market. The little-known Copernican Principle (Chapter 28) enables you to predict the remaining time for which a product will be sold.

Part VIII: Retailing

Analytic techniques can help retailers deal with many important issues. The concepts of market basket analysis and lift (Chapter 29) help retailers derive a store layout that maximizes sales from complementary products. Recency, frequency, and monetary value analysis (Chapter 30) helps direct mailers maximize profit from their mailings. The widely known SCAN*PRO (Chapter 31) model helps retailers

xxvi

determine how factors such as seasonality, price, and promotions influence product sales. In Chapter 32 you learn how to use analytic techniques to determine optimal allocation of store space between products and also optimize the use of a corporate sales force. Finally in Chapter 33 you learn how to forecast total sales of a product from a few data points.

Part IX: Advertising

Department store owner John Wanamaker said, "Half the money I spend on advertising is wasted; the trouble is I don't know which half." In Chapter 34 you learn how John Wanamaker could have used the ADSTOCK model to measure the effectiveness of his advertising expenditures. In Chapter 35 you learn how to allocate ads between the available media outlets to maximize the effectiveness of ads. Chapter 36 deals with the math behind online ad auctions.

Part X: Marketing Research Tools

Often the marketing analyst must deal with data sets involving many variables. Principal components (Chapter 37) and Multidimensional Scaling (Chapter 38) enable the marketing analysts to reduce data sets involving many variables to a few easily understood variables. Often the marketing analyst must classify objects into one of several groups. Naive Bayes and discriminant analysis (Chapter 39) are great tools for developing classification rules. When the marketing analyst wants to determine if a single factor or a pair of factors has a significant effect on product sales, ANOVA (Chapter 40 and Chapter 41) is a useful tool.

Part XI: The Internet and Social Marketing

In the last 20 years, the Internet has turned our world upside down, and marketing is no exception. Social media such as Facebook and Twitter create many interesting opportunities for the marketer, which require careful analysis. In Chapter 42 you learn how the theory of networks sheds light on how you can identify people who are the key to spreading the word about your product. Chapter 43 discusses the math behind Malcom Gladwell's bestselling book *The Tipping Point (Back Bay Books, 2002)*. Chapter 44 discusses the math behind videos (such as the notorious "Gangnam Style") going viral. Finally, in Chapter 45 you learn how text mining can be used to glean useful insight from Twitter, blogs, and Facebook posts.

Who Should Read This Book

There is plenty of material in this book for a one-semester course on marketing analytics at the advanced undergraduate or MBA level. I also believe the book can be useful to any corporate marketing analyst. With regard to prerequisites for the book, I assume you understand the Copy command in Excel. That is, you know when and where to put dollars signs in a formula. If you work hard, that's about all the prior knowledge needed to get a lot out of the book.

I always try to write my books in a modular fashion, so you can skip around and read about what interests you. If you don't want to read the book from start to finish, the following table should help you navigate the book.

Chapters	Chapter Prerequisites
Chapter 1: Slicing and Dicing Marketing Data with PivotTables	None
Chapter 2: Using Excel Charts to Summarize Marketing Data	1
Chapter 3: Using Excel Functions to Summarize Marketing Data	2
Chapter 4: Estimating Demand Curves and Using Solver to Optimize Price	None
Chapter 5: Price Bundling	4
Chapter 6: Nonlinear Pricing	5
Chapter 7: Price Skimming and Sales	5
Chapter 8: Revenue Management	4
Chapter 9: Simple Linear Regression and Correlation	3
Chapter 10: Using Multiple Regression to Forecast Sales	9
Chapter 11: Forecasting in the Presence of Special Events	10
Chapter 12: Modeling Trend and Seasonality	5 and 11
Chapter 13: Ratio to Moving Average Forecasting Method	3 and 12
Chapter 14: Winter's Method	12
Chapter 15: Using Neural Networks to Forecast Sales	10
Chapter 16: Conjoint Analysis	10

Continues

(continued)

Chapters	Chapter Prerequisites
Chapter 17: Logistic Regression	16
Chapter 18: Discrete Choice Analysis	17
Chapter 19: Calculating Lifetime Customer Value	3
Chapter 20: Using Customer Value to Value a Business	19
Chapter 21: Customer Value, Monte Carlo Simulation, and Marketing Decision Making	19
Chapter 22: Allocating Marketing Resources between Customer Acquisition and Retention	4 and 19
Chapter 23: Cluster Analysis	5
Chapter 24: Collaborative Filtering	23
Chapter 25: Using Classification Trees for Segmentation	24
Chapter 26: Using S Curves to Forecast Sales of a New Product	5 and 12
Chapter 27: The Bass Diffusion Model	26
Chapter 28: Using the Copernican Principle to Predict Duration of Future Sales	None
Chapter 29: Market Basket Analysis and Lift	19
Chapter 30: RFM Analysis and Optimizing Direct Mail Campaigns	29
Chapter 31: Using the SCAN*PRO Model and Its Variants	12
Chapter 32: Allocating Retail Space and Sales Resources	5
Chapter 33: Forecasting Sales from Few Data Points	31
Chapter 34: Measuring the Effectiveness of Advertising	31
Chapter 35: Media Selection Models	4, 21, and 34
Chapter 36: Pay Per Click (PPC) Online Advertising	None
Chapter 37: Principal Component Analysis (PCA)	10 and 23
Chapter 38: Multidimensional Scaling (MDS)	37
Chapter 39: Classification Algorithms: Naive Bayes Classifier and Discriminant Analysis	37 and 38

Chapters	Chapter Prerequisites
Chapter 40: Analysis of Variance: One-way ANOVA	10
Chapter 41: Analysis of Variance: Two-way ANOVA	40
Chapter 42: Networks	None
Chapter 43: The Mathematics Behind *The Tipping Point*	42
Chapter 44: Viral Marketing	10, 15, and 39
Chapter 45: Text Mining	3

For example, before reading Chapter 5 you need to have read Chapter 4; before reading Chapter 34 you need to have read Chapter 31, and so on.

Tools You Need

To work through the vast majority of the book, all you need is Excel 2007, 2010, or 2013. Chapters 15, 21, and 35 require use of the Palisade.com Decision Tools Suite. You can download a 15-day trial version of the suite from `www.Palisade.com`.

What's on the Website

From the book website (`www.wiley.com/go/marketinganalytics`) you can download all Excel files used in the book as well as answers to all of the Exercises at the end of each chapter.

Errata

We make every effort to ensure that there are no errors in the text or in the code. However, no one is perfect, and mistakes do occur. If you find an error in this book, like a spelling mistake or a calculation error, we would be very grateful for your feedback. By sending in errata you may save another reader hours of frustration and at the same time you will be helping us provide even higher quality information.

To submit errata for this book go to `http://support.wiley.com` and complete the form on the Ask a Question tab there to send us the error you have found. We'll check

the information and, if appropriate, post a message to the book's errata page and fix the problem in subsequent editions of the book.

Summary

A famous Chinese proverb (popularized by the late management guru Stephen Covey) states, "If you give a man a fish you feed him for a day. If you teach a man to fish you feed him for a lifetime." Hopefully this book will teach you enough about marketing analytics so you will be well equipped to develop your own quantitative marketing models for most problems that come your way. Happy modeling!

Using Excel to Summarize Marketing Data

1

Slicing and Dicing Marketing Data with PivotTables

In many marketing situations you need to analyze, or "slice and dice," your data to gain important marketing insights. Excel PivotTables enable you to quickly summarize and describe your data in many different ways. In this chapter you learn how to use PivotTables to perform the following:

- Examine sales volume and percentage by store, month and product type.
- Analyze the influence of weekday, seasonality, and the overall trend on sales at your favorite bakery.
- Investigate the effect of marketing promotions on sales at your favorite bakery.
- Determine the influence that demographics such as age, income, gender and geographic location have on the likelihood that a person will subscribe to *ESPN*: *The Magazine*.

Analyzing Sales at True Colors Hardware

To start analyzing sales you first need some data to work with. The `data` worksheet from the `PARETO.xlsx` file (available for download on the companion website) contains sales data from two local hardware stores (uptown store owned by Billy Joel and downtown store owned by Petula Clark). Each store sells 10 types of tape, 10 types of adhesive, and 10 types of safety equipment. Figure 1-1 shows a sample of this data.

Throughout this section you will learn to analyze this data using Excel PivotTables to answer the following questions:

- What percentage of sales occurs at each store?
- What percentage of sales occurs during each month?
- How much revenue does each product generate?
- Which products generate 80 percent of the revenue?

	Y	Z	AA	AB
7	Product	Month	Store	Price
8	Tape 10	April	downtown	$2.50
9	Safety 8	August	uptown	$10.00
10	Safety 2	February	uptown	$10.00
11	Safety 8	Novembe	uptown	$10.00
12	Tape 10	October	uptown	$2.50
13	Safety 8	January	uptown	$10.00
14	Safety 8	December	downtown	$10.00
15	Safety 1	Septembe	downtown	$12.00
16	Safety 2	May	uptown	$10.00
17	Adhesive	July	uptown	$7.00
18	Adhesive	March	uptown	$7.00
19	Safety 8	August	downtown	$10.00
20	Safety 8	October	downtown	$10.00
21	Tape 10	July	downtown	$2.50
22	Safety 2	February	downtown	$10.00

Figure 1-1: Hardware store data

Calculating the Percentage of Sales at Each Store

The first step in creating a PivotTable is ensuring you have headings in the first row of your data. Notice that Row 7 of the example data in the data worksheet has the headings Product, Month, Store, and Price. Because these are in place, you can begin creating your PivotTable. To do so, perform the following steps:

1. Place your cursor anywhere in the data cells on the data worksheet, and then click PivotTable in the Tables group on the Insert tab. Excel opens the Create PivotTable dialog box, as shown in Figure 1-2, and correctly guesses that the data is included in the range Y7:AB1333.
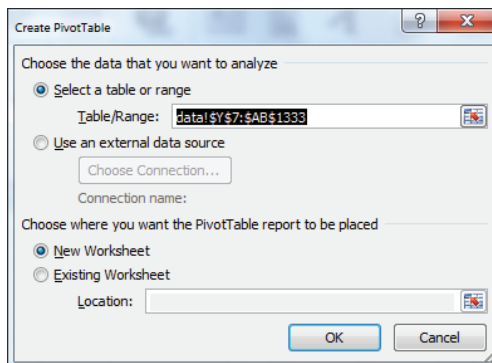
Figure 1-2: PivotTable Dialog Box

NOTE   If you select Use an External Data Source here, you could also refer to a database as a source for a PivotTable. In Exercise 14 at the end of the chapter you can practice creating PivotTables from data in different worksheets or even different workbooks.

2. Click OK and you see the PivotTable Field List, as shown in Figure 1-3.
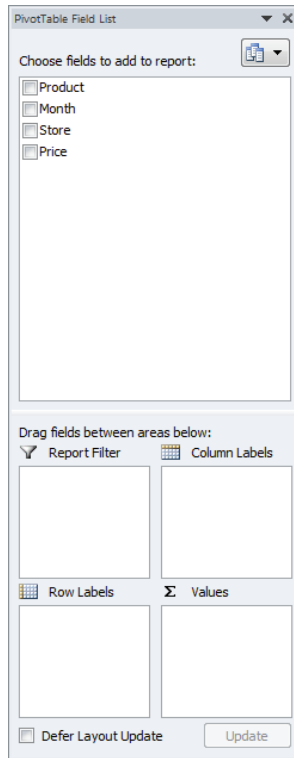
Figure 1-3: PivotTable Field List

3. Fill in the PivotTable Field List by dragging the PivotTable headings or *fields* into the boxes or *zones*. You can choose from the following four zones:

■ **Row Labels:** Fields dragged here are listed on the left side of the table in the order in which they are added to the box. In the current example, the Store field should be dragged to the Row Labels box so that data can be summarized by store.

■ **Column Labels:** Fields dragged here have their values listed across the top row of the PivotTable. In the current example no fields exist in the Column Labels zone.

■ **Values:** Fields dragged here are summarized mathematically in the PivotTable. The Price field should be dragged to this zone. Excel tries to guess the type of calculation you want to perform on a field. In this example Excel guesses that you want all Prices to be summed. Because you want to compute total revenue, this is correct. If you want to change the method of calculation for a data field to an average, a count, or something else, simply double-click the data field or choose Value Field Settings. You learn how to use the Value Fields Setting command later in this section.

■ **Report Filter:** Beginning in Excel 2007, *Report Filter* is the new name for the Page Field area. For fields dragged to the Report Filter zone, you can easily pick any subset of the field values so that the PivotTable shows calculations based only on that subset. In Excel 2010 or Excel 2013 you can use the exciting *Slicers* to select the sub-set of fields used in PivotTable calculations. The use of the Report Filter and Slicers is shown in the "Report Filter and Slicers" section of this chapter.

> **NOTE**   To see the field list, you need to be in a field in the PivotTable. If you do not see the field list, right-click any cell in the PivotTable, and select Show Field List.

Figure 1-4 shows the completed PivotTable Field List and the resulting PivotTable is shown in Figure 1-5 as well as on the `FirstorePT` worksheet.

Figure 1-5 shows the downtown store sold $4,985.50 worth of goods, and the uptown store sold $4,606.50 of goods. The total sales are $9592.

If you want a percentage breakdown of the sales by store, you need to change the way Excel displays data in the Values zone. To do this, perform these steps:

1. Right-click in the summarized data in the `FirstStorePT` worksheet and select Value Field Settings.
2. Select Show Values As and click the drop-down arrow on the right side of the dialog box.
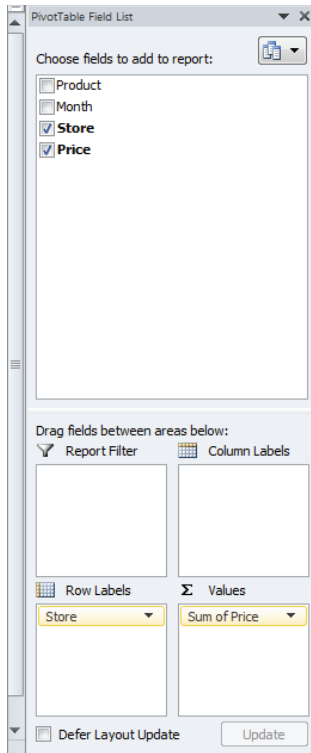3. Select the % of Column Total option, as shown in Figure 1-6.

Figure 1-4: Completed PivotTable Field List

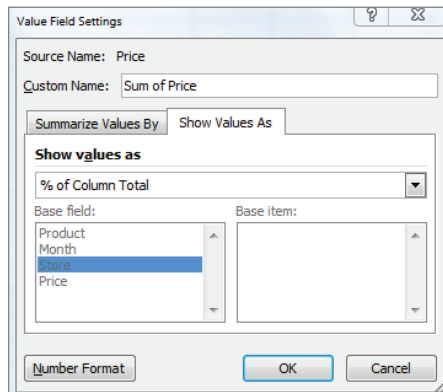

Figure 1-5: Completed PivotTable

Figure 1-6: Obtaining percentage breakdown by Store

Figure 1-7 shows the resulting PivotTable with the new percentage break-down by Store with 52 percent of the sales in the downtown store and 48 percent in the uptown store. You can also see this in the `revenue by store` worksheet of the `PARETO.xlsx` file.

NOTE   If you want a PivotTable to incorporate a different set of data, then under Options, you can select Change Data Source and select the new source data. To have a PivotTable incorporate changes in the original source data, simply right-click and select Refresh. If you are going to add new data below the original data and you want the PivotTable to include the new data when you select Refresh, you should use the Excel Table feature discussed in Chapter 2, "Using Excel Charts to Summarize Marketing Data."

	A	B
1		
2		
3	Row Labels ▾	Sum of Price
4	downtown	51.98%
5	uptown	48.02%
6	**Grand Total**	**100.00%**

Figure 1-7: Percentage breakdown by Store

Summarizing Revenue by Month

You can also use a PivotTable to break down the total revenue by month and cal-culate the percentage of sales that occur during each month. To accomplish this, perform the following steps:

1. Return to the `data` worksheet and bring up the PivotTable Field List by choos-ing Insert PivotTable.
2. Drag the Month field to the Row Labels zone and the Price field to the Values zone. This gives the total sales by month. Because you also want a percentage breakdown of sales by month, drag the Price field again to the Values zone.
3. As shown in Figure 1-8, right-click on the first column in the Values zone and choose Value Field Settings; then choose the % of Column Total option. You now see the percentage monthly breakdown of revenue.
4. Double-click the Column headings and change them to Percentage of Sales by Month and Total Revenue.

5. Finally, double-click again the Total Revenue Column; select Number Format, and choose the Currency option so the revenue is formatted in dollars.

	A	B	C
1			
2			
3	Row Labels ▾	Percentage of Sales by Month	Total Revenue
4	January	8.81%	$845.00
5	February	8.55%	$820.00
6	March	6.58%	$631.00
7	April	8.10%	$776.50
8	May	8.48%	$813.00
9	June	8.87%	$850.50
10	July	8.77%	$841.00
11	August	7.13%	$684.00
12	September	8.71%	$835.50
13	October	7.85%	$753.00
14	November	9.11%	$873.50
15	December	9.06%	$869.00
16	Grand Total	100.00%	$9,592.00

Figure 1-8: Monthly percentage breakdown of Revenue

You can see that $845 worth of goods was sold in January and 8.81 percent of the sales were in January. Because the percentage of sales in each month is approximately 1/12 (8.33 percent), the stores exhibit little seasonality. Part III, "Forecasting Sales of Existing Products," includes an extensive discussion of how to estimate seasonality and the importance of seasonality in marketing analytics.

Calculating Revenue for Each Product

Another important part of analyzing data includes determining the revenue generated by each product. To determine this for the example data, perform the following steps:

1. Return to the data worksheet and drag the Product field to the Row Labels zone and the Price field to the Values zone.
2. Double-click on the Price column, change the name of the Price column to Revenue, and then reformat the Revenue Column as Currency.
3. Click the drop-down arrow in cell A3 and select Sort A to Z so you can alphabetize the product list and obtain the PivotTable in the products worksheet, as shown in Figure 1-9.