

Camm * Cochran * Fry * Ohlmann * Anderson * Sweeney * Williams



Essentials of Business Analytics^{2e}

DESCRIPTIVE * PREDICTIVE * PRESCRIPTIVE

Copyright 2017 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. WCN 02-200-208

This is an electronic version of the print textbook. Due to electronic rights restrictions, some third party content may be suppressed. Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. The publisher reserves the right to remove content from this title at any time if subsequent rights restrictions require it. For valuable information on pricing, previous editions, changes to current editions, and alternate formats, please visit www.cengage.com/highered to search by ISBN, author, title, or keyword for materials in your areas of interest.

Important notice: Media content referenced within the product description or the product text may not be available in the eBook version.



Essentials of Business Analytics

2e

Jeffrey D. Camm
Wake Forest University

James J. Cochran
University of Alabama

Michael J. Fry
University of Cincinnati

Jeffrey W. Ohlmann
University of Iowa

David R. Anderson
University of Cincinnati

Dennis J. Sweeney
University of Cincinnati

Thomas A. Williams
Rochester Institute
of Technology



Australia • Brazil • Japan • Korea • Mexico • Singapore • Spain • United Kingdom • United States

**Essentials of Business Analytics,
Second Edition**

**Jeffrey D. Camm, James J. Cochran,
Michael J. Fry, Jeffrey W. Ohlmann,
David R. Anderson, Dennis J. Sweeney,
Thomas A. Williams**

Vice President, General Manager: Science,
Math & Quantitative Business: Balraj Kalsi

Product Director: Mike Schenk

Product Team Manager: Joe Sabatino

Product Manager: Aaron Arnsperger

Content Developer: Anne Merrill

Senior Marketing Manager: Nate Anderson

Content Project Manager: Jana Lewis

Media Developer: Chris Valentine

Manufacturing Planner: Ron Montgomery

Production Service: MPS Limited

Sr. Art Director: Michelle Kunkler

Internal Designer: Beckmeyer Design

Cover Designer: Beckmeyer Design

Cover Image: iStockphoto.com/alienforce

Intellectual Property Analyst: Brittani Morgan

Intellectual Property Project Manager:
Nick Barrows

© 2017, 2015 Cengage Learning®

ALL RIGHTS RESERVED. No part of this work covered by the copyright herein may be reproduced or distributed in any form or by any means, except as permitted by U.S. copyright law, without the prior written permission of the copyright owner.

For product information and technology assistance, contact us at
Cengage Learning Customer & Sales Support, 1-800-354-9706

For permission to use material from this text or product,
submit all requests online at **www.cengage.com/permissions**

Further permissions questions can be emailed to
permissionrequest@cengage.com

Unless otherwise noted, all items © Cengage Learning.

Microsoft Excel® is a registered trademark of Microsoft Corporation.
©2016 Microsoft.

Library of Congress Control Number: 2015958527

Package ISBN: 978-1-305-62773-4

Cengage Learning

20 Channel Center Street
Boston, MA 02210
USA

Cengage Learning is a leading provider of customized learning solutions with employees residing in nearly 40 different countries and sales in more than 125 countries around the world. Find your local representative at **www.cengage.com**.

Cengage Learning products are represented in Canada by
Nelson Education, Ltd.

To learn more about Cengage Learning Solutions, visit
www.cengage.com

Purchase any of our products at your local college store or at our
preferred online store **www.cengagebrain.com**

Printed in Canada

Print Number: 01 Print Year: 2016

Brief Contents

ABOUT THE AUTHORS XVII

PREFACE XXI

CHAPTER 1	Introduction 1
CHAPTER 2	Descriptive Statistics 17
CHAPTER 3	Data Visualization 76
CHAPTER 4	Descriptive Data Mining 137
CHAPTER 5	Probability: An Introduction to Modeling Uncertainty 169
CHAPTER 6	Statistical Inference 225
CHAPTER 7	Linear Regression 295
CHAPTER 8	Time Series Analysis and Forecasting 380
CHAPTER 9	Predictive Data Mining 438
CHAPTER 10	Spreadsheet Models 508
CHAPTER 11	Linear Optimization Models 540
CHAPTER 12	Integer Linear Optimization Models 595
CHAPTER 13	Nonlinear Optimization Models 641
CHAPTER 14	Monte Carlo Simulation 681
CHAPTER 15	Decision Analysis 752
APPENDIX A	Basics of Excel 810
APPENDIX B	Database Basics with Microsoft Access 823
APPENDIX C	Solutions to Even-Numbered Questions (online)

REFERENCES 863

INDEX 865

MindTap

Feel confident as you use the most **engaging** digital content available to transform today's students into **critical thinkers**.

Personalize your course to match the way you teach and your students learn...

Improve outcomes with real-time insight into student progress...

and

SAVE TIME.



All while your students **engage** with your course content, enjoy the **flexibility** of studying anytime and anywhere, **stay connected** with the MindTap Mobile app...

and

EARN BETTER GRADES.

Try a demo, access the research, learn about our LMS integration and Cengage Learning's unparalleled digital course support now at

www.cengage.com/mindtap

Copyright 2017 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. WCN 02-200-208

Contents

ABOUT THE AUTHORS XVII

PREFACE XXI

CHAPTER 1 Introduction 1

- 1.1 Decision Making 4
- 1.2 Business Analytics Defined 5
- 1.3 A Categorization of Analytical Methods and Models 5
 - Descriptive Analytics 5
 - Predictive Analytics 6
 - Prescriptive Analytics 6
- 1.4 Big Data 7
 - Volume 8
 - Velocity 8
 - Variety 9
 - Veracity 9
- 1.5 Business Analytics in Practice 11
 - Financial Analytics 11
 - Human Resource (HR) Analytics 12
 - Marketing Analytics 12
 - Health Care Analytics 12
 - Supply-Chain Analytics 13
 - Analytics for Government and Nonprofits 14
 - Sports Analytics 14
 - Web Analytics 15
- Summary 15
- Glossary 16

CHAPTER 2 Descriptive Statistics 17

- Analytics in Action: U.S. Census Bureau 18
- 2.1 Overview of Using Data: Definitions and Goals 18
- 2.2 Types of Data 20
 - Population and Sample Data 20
 - Quantitative and Categorical Data 20
 - Cross-Sectional and Time Series Data 20
 - Sources of Data 20
- 2.3 Modifying Data in Excel 23
 - Sorting and Filtering Data in Excel 23
 - Conditional Formatting of Data in Excel 25
- 2.4 Creating Distributions from Data 29
 - Frequency Distributions for Categorical Data 29
 - Relative Frequency and Percent Frequency Distributions 31

Frequency Distributions for Quantitative Data	31
Histograms	34
Cumulative Distributions	38
2.5 Measures of Location	39
Mean (Arithmetic Mean)	39
Median	40
Mode	41
Geometric Mean	41
2.6 Measures of Variability	44
Range	45
Variance	45
Standard Deviation	47
Coefficient of Variation	48
2.7 Analyzing Distributions	48
Percentiles	48
Quartiles	49
z-Scores	50
Empirical Rule	51
Identifying Outliers	53
Box Plots	53
2.8 Measures of Association Between Two Variables	55
Scatter Charts	56
Covariance	56
Correlation Coefficient	57
Summary	61
Glossary	61
Problems	63
Case Problem: Heavenly Chocolates Web Site Transactions	71
Appendix 2.1 Creating Box Plots with XLMiner	73
CHAPTER 3 Data Visualization	76
Analytics in Action: Cincinnati Zoo & Botanical Garden	77
3.1 Overview of Data Visualization	79
Effective Design Techniques	79
3.2 Tables	82
Table Design Principles	84
Crosstabulation	86
PivotTables in Excel	87
Recommended PivotTables in Excel	92
3.3 Charts	95
Scatter Charts	95
Recommended Charts in Excel	97

Line Charts	97
Bar Charts and Column Charts	102
A Note on Pie Charts and Three-Dimensional Charts	105
Bubble Charts	105
Heat Maps	106
Additional Charts for Multiple Variables	109
PivotCharts in Excel	111
3.4 Advanced Data Visualization	114
Advanced Charts	114
Geographic Information Systems Charts	116
3.5 Data Dashboards	117
Principles of Effective Data Dashboards	117
Applications of Data Dashboards	118
Summary	120
Glossary	120
Problems	121
Case Problem: All-Time Movie Box-Office Data	131
Appendix 3.1 Creating a Scatter-Chart Matrix and a Parallel-Coordinates Plot with XLMiner	133
 CHAPTER 4 Descriptive Data Mining	 137
Analytics in Action: Advice from a Machine	138
4.1 Data Preparation	139
Treatment of Missing Data	139
Identification of Outliers and Erroneous Data	140
Variable Representation	140
4.2 Cluster Analysis	141
Measuring Similarity Between Observations	142
Hierarchical Clustering	144
k-Means Clustering	148
Hierarchical Clustering Versus k-Means Clustering	149
4.3 Association Rules	150
Evaluating Association Rules	151
Summary	152
Glossary	153
Problems	154
Case Problem: Know Thy Customer	158
Appendix 4.1 Hierarchical Clustering with XLMiner	160
Appendix 4.2 k-Means Clustering with XLMiner	163
Appendix 4.3 Association Rules with XLMiner	166

CHAPTER 5 Probability: An Introduction to Modeling Uncertainty 169

Analytics in Action: National Aeronautics and Space Administration 170

5.1 Events and Probabilities 171

5.2 Some Basic Relationships of Probability 173

Complement of an Event 173

Addition Law 173

5.3 Conditional Probability 176

Independent Events 180

Multiplication Law 181

Bayes' Theorem 181

5.4 Random Variables 184

Discrete Random Variables 184

Continuous Random Variables 185

5.5 Discrete Probability Distributions 186

Custom Discrete Probability Distribution 186

Expected Value and Variance 188

Discrete Uniform Probability Distribution 192

Binomial Probability Distribution 192

Poisson Probability Distribution 194

5.6 Continuous Probability Distributions 198

Uniform Probability Distribution 198

Triangular Probability Distribution 201

Normal Probability Distribution 202

Exponential Probability Distribution 208

Summary 211

Glossary 212

Problems 213

Case Problem: Hamilton County Judges 223

CHAPTER 6 Statistical Inference 225

Analytics in Action: John Morrell & Company 226

6.1 Selecting a Sample 228

Sampling from a Finite Population 229

Sampling from an Infinite Population 231

6.2 Point Estimation 233

Practical Advice 235

6.3 Sampling Distributions 235

Sampling Distribution of $\bar{x}$ 237Sampling Distribution of $\bar{p}$ 243

6.4 Interval Estimation 246

Interval Estimation of the Population Mean 247

Interval Estimation of the Population Proportion 254

6.5 Hypothesis Tests 257

Developing Null and Alternative Hypotheses 258

Type I and Type II Errors 261

Hypothesis Test of the Population Mean 262

Hypothesis Test of the Population Proportion 274

Big Data, Statistical Inference, and Practical Significance 278

Summary 280

Glossary 281

Problems 282

Case Problem 1: Young Professional Magazine 291

Case Problem 2: Quality Associates, Inc. 292

CHAPTER 7 Linear Regression 295

Analytics in Action: Alliance Data Systems 296

7.1 Simple Linear Regression Model 297

Regression Model 297

Estimated Regression Equation 297

7.2 Least Squares Method 299

Least Squares Estimates of the Regression

Parameters 301

Using Excel's Chart Tools to Compute the Estimated
Regression Equation 304**7.3 Assessing the Fit of the Simple Linear Regression Model 305**

The Sums of Squares 305

The Coefficient of Determination 308

Using Excel's Chart Tools to Compute the Coefficient
of Determination 309**7.4 The Multiple Regression Model 310**

Regression Model 310

Estimated Multiple Regression Equation 310

Least Squares Method and Multiple Regression 311

Butler Trucking Company and Multiple Regression 312

Using Excel's Regression Tool to Develop the Estimated Multiple
Regression Equation 312**7.5 Inference and Regression 315**Conditions Necessary for Valid Inference in the Least Squares
Regression Model 316

Testing Individual Regression Parameters 322

Addressing Nonsignificant Independent Variables 324

Multicollinearity 325

Inference and Very Large Samples 328

7.6 Categorical Independent Variables 332

Butler Trucking Company and Rush Hour 332

Interpreting the Parameters	333
More Complex Categorical Variables	335
7.7 Modeling Nonlinear Relationships	337
Quadratic Regression Models	338
Piecewise Linear Regression Models	342
Interaction Between Independent Variables	345
7.8 Model Fitting	350
Variable Selection Procedures	350
Overfitting	351
Summary	353
Glossary	353
Problems	355
Case Problem: Alumni Giving	371
Appendix 7.1 Regression with XLMiner	374
CHAPTER 8 Time Series Analysis and Forecasting	380
Analytics in Action: ACCO Brands	381
8.1 Time Series Patterns	383
Horizontal Pattern	383
Trend Pattern	385
Seasonal Pattern	387
Trend and Seasonal Pattern	388
Cyclical Pattern	388
Identifying Time Series Patterns	390
8.2 Forecast Accuracy	390
8.3 Moving Averages and Exponential Smoothing	395
Moving Averages	395
Forecast Accuracy	399
Exponential Smoothing	400
Forecast Accuracy	403
8.4 Using Regression Analysis for Forecasting	405
Linear Trend Projection	405
Seasonality	407
Seasonality Without Trend	407
Seasonality with Trend	409
Using Regression Analysis as a Causal Forecasting Method	411
Combining Causal Variables with Trend and Seasonality Effects	414
Considerations in Using Regression in Forecasting	415
8.5 Determining the Best Forecasting Model to Use	415
Summary	416
Glossary	416
Problems	417

Case Problem: Forecasting Food and Beverage Sales 426

Appendix 8.1 Using Excel Forecast Sheet 427

Appendix 8.2 Forecasting with XLMiner 433

CHAPTER 9 Predictive Data Mining 438

Analytics in Action: Orbitz 439

9.1 Data Sampling 440

9.2 Data Partitioning 441

9.3 Accuracy Measures 442

Evaluating the Classification of Categorical Outcomes 442

Evaluating the Estimation of Continuous Outcomes 448

9.4 Logistic Regression 449

9.5 *k*-Nearest Neighbors 453

Classifying Categorical Outcomes with *k*-Nearest
Neighbors 453

Estimating Continuous Outcomes with *k*-Nearest Neighbors 455

9.6 Classification and Regression Trees 456

Classifying Categorical Outcomes with a Classification Tree 456

Estimating Continuous Outcomes with a Regression Tree 462

Ensemble Methods 463

Summary 466

Glossary 468

Problems 469

Case Problem: Grey Code Corporation 477

Appendix 9.1 Data Partitioning with XLMiner 479

Appendix 9.2 Logistic Regression Classification with XLMiner 481

Appendix 9.3 *k*-Nearest Neighbor Classification and Estimation with
XLMiner 488

Appendix 9.4 Single Classification and Regression Trees with
XLMiner 494

Appendix 9.5 Random Forests of Classification or Regression Trees with
XLMiner 503

CHAPTER 10 Spreadsheet Models 508

Analytics in Action: Procter & Gamble 509

10.1 Building Good Spreadsheet Models 510

Influence Diagrams 510

Building a Mathematical Model 510

Spreadsheet Design and Implementing the Model in a
Spreadsheet 512

10.2 What-If Analysis 515

Data Tables 515

Goal Seek 519

10.3	Some Useful Excel Functions for Modeling	521
	SUM and SUMPRODUCT	521
	IF and COUNTIF	523
	VLOOKUP	525
10.4	Auditing Spreadsheet Models	527
	Trace Precedents and Dependents	527
	Show Formulas	528
	Evaluate Formulas	528
	Error Checking	529
	Watch Window	529
	Summary	532
	Glossary	532
	Problems	532
	Case Problem: Retirement Plan	539
CHAPTER 11 Linear Optimization Models		540
	Analytics in Action: MeadWestvaco Corporation	541
11.1	A Simple Maximization Problem	542
	Problem Formulation	543
	Mathematical Model for the Par, Inc. Problem	545
11.2	Solving the Par, Inc. Problem	546
	The Geometry of the Par, Inc. Problem	546
	Solving Linear Programs with Excel Solver	548
11.3	A Simple Minimization Problem	552
	Problem Formulation	552
	Solution for the M&D Chemicals Problem	553
11.4	Special Cases of Linear Program Outcomes	555
	Alternative Optimal Solutions	555
	Infeasibility	556
	Unbounded	558
11.5	Sensitivity Analysis	560
	Interpreting Excel Solver Sensitivity Report	560
11.6	General Linear Programming Notation and More Examples	562
	Investment Portfolio Selection	563
	Transportation Planning	565
	Advertising Campaign Planning	569
11.7	Generating an Alternative Optimal Solution for a Linear Program	575
	Summary	576
	Glossary	577
	Problems	578
	Case Problem: Investment Strategy	587

Appendix 11.1 Solving Linear Optimization Models Using Analytic Solver Platform 589

CHAPTER 12 Integer Linear Optimization Models 595

Analytics in Action: Petrobras 596

12.1 Types of Integer Linear Optimization Models 596

12.2 Eastborne Realty, An Example of Integer Optimization 597
The Geometry of Linear All-Integer Optimization 598

12.3 Solving Integer Optimization Problems with Excel Solver 600

A Cautionary Note About Sensitivity Analysis 604

12.4 Applications Involving Binary Variables 605

Capital Budgeting 605

Fixed Cost 606

Bank Location 611

Product Design and Market Share Optimization 613

12.5 Modeling Flexibility Provided by Binary Variables 617

Multiple-Choice and Mutually Exclusive Constraints 617

k Out of n Alternatives Constraint 617

Conditional and Corequisite Constraints 618

12.6 Generating Alternatives in Binary Optimization 618

Summary 620

Glossary 621

Problems 622

Case Problem: Applecore Children's Clothing 632

Appendix 12.1 Solving Integer Linear Optimization Problems Using Analytic Solver Platform 635

CHAPTER 13 Nonlinear Optimization Models 641

Analytics in Action: Intercontinental Hotels 642

13.1 A Production Application: Par, Inc. Revisited 642

An Unconstrained Problem 643

A Constrained Problem 643

Solving Nonlinear Optimization Models Using Excel Solver 646

Sensitivity Analysis and Shadow Prices in Nonlinear Models 647

13.2 Local and Global Optima 648

Overcoming Local Optima with Excel Solver 651

13.3 A Location Problem 653

13.4 Markowitz Portfolio Model 654

13.5 Forecasting Adoption of a New Product 659

Summary	663
Glossary	664
Problems	664
Case Problem: Portfolio Optimization with Transaction Costs	672
Appendix 13.1 Solving Nonlinear Optimization Problems with Analytic Solver Platform	676

CHAPTER 14 Monte Carlo Simulation 681

Analytics in Action: Cook County Hospital ICU	682
14.1 Risk Analysis for Sanotronics LLC	683
Base-Case Scenario	683
Worst-Case Scenario	684
Best-Case Scenario	684
Sanotronics Spreadsheet Model	684
Use of Probability Distributions to Represent Random Variables	685
Generating Values for Random Variables with Excel	687
Executing Simulation Trials with Excel	690
Measuring and Analyzing Simulation Output	692
14.2 Simulation Modeling for Land Shark Inc.	694
Spreadsheet Model for Land Shark	695
Generating Values for Land Shark's Random Variables	697
Executing Simulation Trials and Analyzing Output	698
14.3 Simulation Considerations	700
Verification and Validation	700
Advantages and Disadvantages of Using Simulation	701
Summary	701
Glossary	702
Problems	703
Case Problem: Four Corners	714
Appendix 14.1 Land Shark Inc. Simulation with Analytic Solver Platform	716
Appendix 14.2 Distribution Fitting with Analytic Solver Platform	724
Appendix 14.3 Simulation Optimization with Analytic Solver Platform	734
Appendix 14.4 Correlating Random Variables with Analytic Solver Platform	740
Appendix 14.5 Probability Distributions for Random Variables	747

CHAPTER 15 Decision Analysis 752

Analytics in Action: Phytopharm	753
15.1 Problem Formulation	754

	Payoff Tables	755
	Decision Trees	755
15.2	Decision Analysis Without Probabilities	756
	Optimistic Approach	756
	Conservative Approach	757
	Minimax Regret Approach	757
15.3	Decision Analysis with Probabilities	759
	Expected Value Approach	759
	Risk Analysis	761
	Sensitivity Analysis	762
15.4	Decision Analysis with Sample Information	763
	Expected Value of Sample Information	769
	Expected Value of Perfect Information	769
15.5	Computing Branch Probabilities with Bayes' Theorem	770
15.6	Utility Theory	773
	Utility and Decision Analysis	775
	Utility Functions	778
	Exponential Utility Function	782
	Summary	783
	Glossary	784
	Problems	786
	Case Problem: Property Purchase Strategy	797
	Appendix 15.1 Using Analytic Solver Platform to Create Decision Trees	799
APPENDIX A	Basics of Excel	810
APPENDIX B	Database Basics with Microsoft Access	823
APPENDIX C	Solutions to Even-Numbered Questions (online)	
REFERENCES		863
INDEX		865

About the Authors

Jeffrey D. Camm. Jeffrey D. Camm is the Inmar Presidential Chair and Associate Dean of Analytics in the School of Business at Wake Forest University. Born in Cincinnati, Ohio, he holds a B.S. from Xavier University (Ohio) and a Ph.D. from Clemson University. Prior to joining the faculty at Wake Forest, he was on the faculty of the University of Cincinnati. He has also been a visiting scholar at Stanford University and a visiting professor of business administration at the Tuck School of Business at Dartmouth College.

Dr. Camm has published over 30 papers in the general area of optimization applied to problems in operations management and marketing. He has published his research in *Science*, *Management Science*, *Operations Research*, *Interfaces*, and other professional journals. Dr. Camm was named the Dornoff Fellow of Teaching Excellence at the University of Cincinnati and he was the 2006 recipient of the INFORMS Prize for the Teaching of Operations Research Practice. A firm believer in practicing what he preaches, he has served as an operations research consultant to numerous companies and government agencies. From 2005 to 2010 he served as editor-in-chief of *Interfaces*.

James J. Cochran. James J. Cochran is Professor of Applied Statistics and the Rogers-Spivey Faculty Fellow at The University of Alabama. Born in Dayton, Ohio, he earned his B.S., M.S., and M.B.A. degrees from Wright State University and a Ph.D. from the University of Cincinnati. He has been at The University of Alabama since 2014 and has been a visiting scholar at Stanford University, Universidad de Talca, the University of South Africa, and Pole Universitaire Leonard de Vinci.

Professor Cochran has published over three dozen papers in the development and application of operations research and statistical methods. He has published his research in *Management Science*, *The American Statistician*, *Communications in Statistics—Theory and Methods*, *Annals of Operations Research*, *European Journal of Operational Research*, *Journal of Combinatorial Optimization*, *Interfaces*, *Statistics and Probability Letters*, and other professional journals. He was the 2008 recipient of the INFORMS Prize for the Teaching of Operations Research Practice and the 2010 recipient of the Mu Sigma Rho Statistical Education Award. Professor Cochran was elected to the International Statistics Institute in 2005 and named a Fellow of the American Statistical Association in 2011. He also received the Founders Award in 2014 and the Karl E. Peace Award in 2015 from the American Statistical Association. A strong advocate for effective operations research and statistics education as a means of improving the quality of applications to real problems, Professor Cochran has organized and chaired teaching effectiveness workshops in Montevideo, Uruguay; Cape Town, South Africa; Cartagena, Colombia; Jaipur, India; Buenos Aires, Argentina; and Nairobi, Kenya. He has served as an operations research consultant to numerous companies and not-for-profit organizations. He served as editor-in-chief of *INFORMS Transactions on Education* from 2006–2012 and is on the editorial board of *Interfaces*, *International Transactions in Operational Research*, and *Significance*.

Michael J. Fry. Michael J. Fry is Professor and Head of the Department of Operations, Business Analytics, and Information Systems in the Carl H. Lindner College of Business at the University of Cincinnati. Born in Killeen, Texas, he earned a B.S. from Texas A&M University, and M.S.E. and Ph.D. degrees from the University of Michigan. He has been at the University of Cincinnati since 2002, where he has been named a Lindner Research Fellow and has served as Assistant Director and Interim Director of the Center for Business Analytics. He has also been a visiting professor at the Samuel Curtis Johnson Graduate School of Management at Cornell University and the Sauder School of Business at the University of British Columbia.

Professor Fry has published over twenty research papers in journals such as *Operations Research*, *M&SOM*, *Transportation Science*, *Naval Research Logistics*, *IIE Transactions*, and *Interfaces*. His research interests are in applying quantitative management methods to the areas of supply chain analytics, sports analytics, and public-policy operations. He has worked with many different organizations for his research, including Dell, Inc., Copeland Corporation, Starbucks Coffee Company, Great American Insurance Group, the Cincinnati Fire Department, the State of Ohio Election Commission, the Cincinnati Bengals, and the Cincinnati Zoo. In 2008, he was named a finalist for the Daniel H. Wagner Prize for Excellence in Operations Research Practice, and he has been recognized for both his research and teaching excellence at the University of Cincinnati.

Jeffrey W. Ohlmann. Jeffrey W. Ohlmann is Associate Professor of Management Sciences and Huneke Research Fellow in the Tippie College of Business at the University of Iowa. Born in Valentine, Nebraska, he earned a B.S. from the University of Nebraska, and M.S. and Ph.D. degrees from the University of Michigan. He has been at the University of Iowa since 2003.

Professor Ohlmann's research on the modeling and solution of decision-making problems has produced over a dozen research papers in journals such as *Operations Research*, *Mathematics of Operations Research*, *INFORMS Journal on Computing*, *Transportation Science*, the *European Journal of Operational Research*, and *Interfaces*. He has collaborated with companies such as Transfreight, LeanCor, Cargill, the Hamilton County Board of Elections, and three National Football League franchises. Due to the relevance of his work to industry, he was bestowed the George B. Dantzig Dissertation Award and was recognized as a finalist for the Daniel H. Wagner Prize for Excellence in Operations Research Practice.

David R. Anderson. David R. Anderson is Professor Emeritus of Quantitative Analysis in the Carl H. Lindner College of Business at the University of Cincinnati. Born in Grand Forks, North Dakota, he earned his B.S., M.S., and Ph.D. degrees from Purdue University. Professor Anderson has served as Head of the Department of Quantitative Analysis and Operations Management and as Associate Dean of the College of Business Administration. In addition, he was the coordinator of the College's first Executive Program.

At the University of Cincinnati, Professor Anderson has taught introductory statistics for business students as well as graduate-level courses in regression analysis, multivariate analysis, and management science. He has also taught statistical courses at the Department of Labor in Washington, D.C. He has been honored with nominations and awards for excellence in teaching and excellence in service to student organizations.

Professor Anderson has coauthored 10 textbooks in the areas of statistics, management science, linear programming, and production and operations management. He is an active consultant in the field of sampling and statistical methods.

Dennis J. Sweeney. Dennis J. Sweeney is Professor Emeritus of Quantitative Analysis and Founder of the Center for Productivity Improvement at the University of Cincinnati. Born in Des Moines, Iowa, he earned a B.S.B.A. degree from Drake University and his M.B.A. and D.B.A. degrees from Indiana University, where he was an NDEA Fellow. During 1978–1979, Professor Sweeney worked in the management science group at Procter & Gamble; during 1981–1982, he was a visiting professor at Duke University. Professor Sweeney served as Head of the Department of Quantitative Analysis and as Associate Dean of the College of Business Administration at the University of Cincinnati.

Professor Sweeney has published more than 30 articles and monographs in the areas of management science and statistics. The National Science Foundation, IBM, Procter & Gamble, Federated Department Stores, Kroger, and Cincinnati Gas & Electric have funded his research, which has been published in *Management Science*, *Operations Research*, *Mathematical Programming*, *Decision Sciences*, and other journals.

Professor Sweeney has coauthored 10 textbooks in the areas of statistics, management science, linear programming, and production and operations management.

Thomas A. Williams. Thomas A. Williams is Professor Emeritus of Management Science in the College of Business at Rochester Institute of Technology. Born in Elmira, New York, he earned his B.S. degree at Clarkson University. He did his graduate work at Rensselaer Polytechnic Institute, where he received his M.S. and Ph.D. degrees.

Before joining the College of Business at RIT, Professor Williams served for seven years as a faculty member in the College of Business Administration at the University of Cincinnati, where he developed the undergraduate program in Information Systems and then served as its coordinator. At RIT he was the first chairman of the Decision Sciences Department. He teaches courses in management science and statistics, as well as graduate courses in regression and decision analysis.

Professor Williams is the coauthor of 11 textbooks in the areas of management science, statistics, production and operations management, and mathematics. He has been a consultant for numerous *Fortune* 500 companies and has worked on projects ranging from the use of data analysis to the development of large-scale regression models.

Preface

E*ssentials of Business Analytics* 2E is designed to introduce the concept of business analytics to undergraduate and graduate students. This textbook contains one of the first collections of materials that are essential to the growing field of business analytics. In Chapter 1 we present an overview of business analytics and our approach to the material in this textbook. In simple terms, business analytics helps business professionals make better decisions based on data. We discuss models for summarizing, visualizing, and understanding useful information from historical data in Chapters 2 through 6. Chapters 7 through 9 introduce methods for both gaining insights from historical data as well as predicting possible future outcomes. Chapter 10 covers the use of spreadsheets for examining data and building decision models. In Chapters 11 through 12 we discuss optimization models to help decision makers choose the best decision based on the available data. Chapter 13 presents material that some may consider more advanced forms of optimization (nonlinear optimization models), although these models are extremely useful and widely applicable to many business situations. In any case, some instructors may choose to omit covering Chapter 13. In Chapter 14 we introduce the concept of simulation models for understanding the effect of uncertainty on decisions. Chapter 15 is an overview of decision analysis approaches for incorporating a decision maker's views about risk into decision making. In Appendix A we present optional material for students who need to learn the basics of using Microsoft Excel. The use of databases and manipulating data in Microsoft Access is discussed in Appendix B.

This textbook can be used by students who have previously taken a course on basic statistical methods as well as students who have not had a prior course in statistics. The expanded material in the second edition of *Essentials of Business Analytics* also makes it amenable to a two-course sequence in business statistics and analytics. All statistical concepts contained in this textbook are presented from a business analytics perspective using practical business examples. Chapters 2, 5, 6 and 7 provide an introduction to basic statistical concepts that form the foundation for more advanced analytics methods. Chapters 3, 4 and 9 cover additional topics of data visualization and data mining that are not traditionally part of most introductory business statistics courses, but they are exceedingly important and commonly used in current business environments. Chapter 10 and Appendix A provide the foundational knowledge students need to use Microsoft Excel for analytics applications. Chapters 11 through 15 build upon this spreadsheet knowledge to present additional topics that are used by many organizations that are leaders in the use of prescriptive analytics to improve decision making.

Updates in the Second Edition

The second edition of *Essentials of Business Analytics* is a major revision of the first edition. We have added several new chapters, expanded the coverage of existing chapters, and updated all chapters based on changes in the software used with this textbook. Stylistically, the 2nd edition of *Essentials of Business Analytics* also has an entirely new look. We have added full-color figures throughout the textbook that make many chapters much more meaningful and easier to read.

- **New Chapters on Probability and Statistical Inference.** Chapters 5 and 6 are new to this edition. Chapter 5 covers an introduction to probability for those students who are not familiar with basic probability concepts such as random variables, conditional probability, Bayes' theorem, and probability distributions. Chapter 6 presents statistical inference topics such as sampling, sampling distributions, interval estimation, and hypothesis testing. These two chapters extend the basic statistical coverage

of *Essentials of Business Analytics* (in conjunction with Chapter 2 on Descriptive Statistics and Chapter 7 on Linear Regression) so that the textbook includes a full coverage of introductory business statistics for students who are unfamiliar with these concepts.

- **Expanded Data Mining Coverage.** The Data Mining chapter from the first edition has been broken into two chapters: Chapter 4 on Descriptive Data Mining and Chapter 9 on Predictive Data Mining. This allows us to cover additional material related to these concepts and to also clearly delineate the different forms of data mining based on their intended result. Chapter 4 on Descriptive Data Mining covers unsupervised learning methods such as clustering and association rules where the user is interested in identifying relationships among observations rather than predicting specific outcome variables. Chapter 4 also covers very important topics related to data preparation including missing data, outliers, and variable representation. Chapter 9 on Predictive Data Mining introduces supervised learning methods that are used to predict an outcome based on a set of input variables. The methods covered in Chapter 9 include logistic regression, k -nearest neighbors clustering, and classification and regression trees. Additional data preparation methods such as data sampling and data portioning are also covered in this chapter.
- **Revision of Linear Regression Chapter.** Based on user feedback from the first edition, Chapter 7's coverage of linear regression has been substantially revised to streamline the exposition with a focus on intuitive understanding without sacrificing technical accuracy. The appendix of this chapter has been expanded to demonstrate the construction of prediction intervals using the Analytic Solver Platform software.
- **New Appendix to Chapter 8.** Chapter 8 on Time Series Analysis and Forecasting now includes an appendix on Excel 2016's new Forecast Sheet tool for implementing Holt-Winters additive seasonal smoothing model.
- **Revision of Simulation Chapter.** As with all other chapters, the coverage of Analytics Solver Platform has been moved to the appendix. All material in the body of the chapter uses only native Excel to implement Monte Carlo simulations.
- **Coverage of Analytic Solver Platform (ASP) Moved to Chapter Appendices.** All coverage of the Excel add-in, Analytics Solver Platform, has been moved to the chapter appendices. This means that instructors can now cover all the material in the bodies of the chapters using only native Excel functionality. ASP is used most heavily in the data mining and simulation chapters, so the result of this change is that the chapter appendices are quite long for Chapters 4, 9, and 14. However, this change makes it easier for an instructor to tailor a course's coverage of data mining concepts and the execution of these concepts.
- **Updates to ASP.** All examples, problems, and solutions have been updated in response to changes in the ASP software. Frontline Systems, the developer of ASP, implemented a major rewrite of the code base that powers ASP shortly after the release of the first edition of *Essentials of Business Analytics*. This new code base is much faster and more stable than the previous releases of ASP, but it also completely changed the output given by ASP in many cases. All the material related to ASP is updated to correspond to Analytic Solver Platform V2016 (16.0.0).
- **Incorporation of Excel 2016.** Most updates in Excel 2016 are relatively minor as they relate to its use for statistics and analytics. However, Excel 2016 does have new options for creating Charts in Excel and for implementing forecasting methods. Excel 2016 allows for the creation of box plots, tree maps, and several other data visualization tools that could not be created in previous versions of Excel. Excel's new Forecast Sheet tool implements a time series forecasting model known as the Holt-Winters

additive seasonal smoothing model; this is covered in the appendix to Chapter 8. Several other minor updates to the Ribbon and tabs have also been made in Excel 2016. All material in the second edition of this textbook is easily accessible for students using earlier versions of Excel. For Excel tools that are only implementable in Excel 2016, we include these either in a chapter appendix (such as Forecast Sheet in Chapter 8 appendix) or with margin notes explaining how the same action can be executed in Excel 2013.

- **Additional Excel Features Incorporated.** Several other features that were introduced in Excel 2013 have been more fully incorporated in this edition. Chapter 2 introduces the Quick Analysis button in Excel, and Chapter 3 now makes full use of the Chart Buttons in Excel. The Quick Analysis button is a shortcut method for accomplishing many common Excel formatting and other tasks. The Chart Buttons make it much easier to format, edit, and analyze charts in Excel. Chapter 3 also now also includes coverage of the Recommended PivotTables and Recommended Charts tools in Excel.
- **New Style and More Color.** The second edition of *Essentials of Business Analytics* includes full color figures and a new color template throughout the text. This makes much of the material covered, such as Chapter 3 on Data Visualization, much easier for students to interpret and understand.

Continued Features and Pedagogy

The style of this textbook is based on the other classic textbooks written by the Anderson, Sweeney, and Williams (ASW) team. Some of the specific features that we use in this textbook are listed below.

- **Integration of Microsoft Excel:** Excel has been thoroughly integrated throughout this textbook. For many methodologies, we provide instructions for how to perform calculations both by hand and with Excel. In other cases where realistic models are practical only with the use of a spreadsheet, we focus on the use of Excel to describe the methods to be used.
- **Notes and Comments:** At the end of many sections, we provide Notes and Comments to give the student additional insights about the methods presented in that section. These insights include comments on the limitations of the presented methods, recommendations for applications, and other matters. Additionally, margin notes are used throughout the textbook to provide additional insights and tips related to the specific material being discussed.
- **Analytics in Action:** Each chapter contains an Analytics in Action article. These articles present interesting examples of the use of business analytics in practice. The examples are drawn from many different organizations in a variety of areas including healthcare, finance, manufacturing, marketing, and others.
- **DATAfiles and MODELfiles:** All data sets used as examples and in student exercises are also provided online as files available for download by the student. DATAfiles are Excel files that contain data needed for the examples and problems given in the textbook. MODELfiles contain additional modeling features such as extensive use of Excel formulas or the use of Excel Solver or Analytic Solver Platform.
- **Problems and Cases:** With the exception of Chapter 1, each chapter contains an extensive selection of problems to help the student master the material presented in that chapter. The problems vary in difficulty and most relate to specific examples of the use of business analytics in practice. Answers to even-numbered problems are provided in

an online supplement for student access. With the exception of Chapter 1, each chapter also includes an in-depth case study that connects many of the different methods introduced in the chapter. The case studies are designed to be more open-ended than the chapter problems, but enough detail is provided to give the student some direction in solving the cases.

MindTap

MindTap is a customizable digital course solution that includes an interactive eBook, auto-graded exercises from the textbook, and author-created video walkthroughs of key chapter concepts and select examples that use Analytic Solver platform. All of these materials offer students better access to understand the materials within the course. For more information on MindTap, please contact your Cengage representative.

For Students

Online resources are available to help the student work more efficiently. The resources can be accessed through **www.cengagebrain.com**.

- **Analytic Solver Platform:** Instructions to download an educational version of Frontline Systems' Analytic Solver Platform are included with the purchase of this textbook. These instructions can be found within the inside front cover of this text.

For Instructors

Instructor resources are available to adopters on the Instructor Companion Site, which can be found and accessed at **www.cengage.com**, including:

- **Solutions Manual:** The Solutions Manual, prepared by the authors, includes solutions for all problems in the text. It is available online as well as print.
- **Solutions to Case Problems:** These are also prepared by the authors and contain solutions to all case problems presented in the text.
- **PowerPoint Presentation Slides:** The presentation slides contain a teaching outline that incorporates figures to complement instructor lectures.
- **Test Bank:** Cengage Learning Testing Powered by Cognero is a flexible, online system that allows you to:
 - author, edit, and manage test bank content from multiple Cengage Learning solutions,
 - create multiple test versions in an instant, and
 - deliver tests from your LMS, your classroom, or wherever you want. The Test Bank is also available in Microsoft Word.

Acknowledgements

We would like to acknowledge the work of our reviewers, who provided comments and suggestions for improvement of this text. Thanks to:

Matthew D. Bailey
Bucknell University

Phillip Beaver
Daniels College of Business University of Denver

M. Khurram S. Bhutta
Ohio University

Q B. Chung
Villanova University

Elizabeth A. Denny
University of Kentucky

Mike Taein Eom
University of Portland

Yvette Njan Essounga
Fayetteville State University

Lawrence V. Fulton
Texas State University

Ed Wasil
American University

Paolo Catasti
Virginia Commonwealth University

James F. Hoelscher
Lincoln Memorial University

Eric Huggins
Fort Lewis College

Faizul Huq
Ohio University

Marco Lam
York College of Pennsylvania

Ram Pakath
University of Kentucky

Susan Palocsay
James Madison University

Dothan Truong
Embry-Riddle Aeronautical University

Kai Wang
Wake Technical Community College

We are indebted to our product director Joe Sabatino and our product manager, Aaron Arnsparger; our marketing manager, Nathan Anderson, and our content developer, Anne Merrill; our content project manager, Jana Lewis; our media developer, Chris Valentine; and others at Cengage Learning for their counsel and support during the preparation of this text.

Jeffrey D. Camm
James J. Cochran
Michael J. Fry
Jeffrey W. Ohlmann
David R. Anderson
Dennis J. Sweeney
Thomas A. Williams

Chapter 1

Introduction

CONTENTS

1.1 DECISION MAKING

1.2 BUSINESS ANALYTICS DEFINED

1.3 A CATEGORIZATION OF ANALYTICAL METHODS AND MODELS

Descriptive Analytics

Predictive Analytics

Prescriptive Analytics

1.4 BIG DATA

Volume

Velocity

Variety

Veracity

1.5 BUSINESS ANALYTICS IN PRACTICE

Financial Analytics

Human Resource (HR) Analytics

Marketing Analytics

Health Care Analytics

Supply-Chain Analytics

Analytics for Government and Nonprofits

Sports Analytics

Web Analytics

You apply for a loan for the first time. How does the bank assess the riskiness of the loan it might make to you? How does Amazon.com know which books and other products to recommend to you when you log in to their web site? How do airlines determine what price to quote to you when you are shopping for a plane ticket? How can doctors better diagnose and treat you when you are ill or injured?

You may be applying for a loan for the first time, but millions of people around the world have applied for loans before. Many of these loan recipients have paid back their loans in full and on time, but some have not. The bank wants to know whether you are more like those who have paid back their loans or more like those who defaulted. By comparing your credit history, financial situation, and other factors to the vast database of previous loan recipients, the bank can effectively assess how likely you are to default on a loan.

Similarly, Amazon.com has access to data on millions of purchases made by customers on its web site. Amazon.com examines your previous purchases, the products you have viewed, and any product recommendations you have provided. Amazon.com then searches through its huge database for customers who are similar to you in terms of product purchases, recommendations, and interests. Once similar customers have been identified, their purchases form the basis of the recommendations given to you.

Prices for airline tickets are frequently updated. The price quoted to you for a flight between New York and San Francisco today could be very different from the price that will be quoted tomorrow. These changes happen because airlines use a pricing strategy known as revenue management. Revenue management works by examining vast amounts of data on past airline customer purchases and using these data to forecast future purchases. These forecasts are then fed into sophisticated optimization algorithms that determine the optimal price to charge for a particular flight and when to change that price. Revenue management has resulted in substantial increases in airline revenues.

Finally, consider the case of being evaluated by a doctor for a potentially serious medical issue. Hundreds of medical papers may describe research studies done on patients facing similar diagnoses, and thousands of data points exist on their outcomes. However, it is extremely unlikely that your doctor has read every one of these research papers or is aware of all previous patient outcomes. Instead of relying only on her medical training and knowledge gained from her limited set of previous patients, wouldn't it be better for your doctor to have access to the expertise and patient histories of thousands of doctors around the world?

A group of IBM computer scientists initiated a project to develop a new decision technology to help in answering these types of questions. That technology is called Watson, named after the founder of IBM, Thomas J. Watson. The team at IBM focused on one aim: how the vast amounts of data now available on the Internet can be used to make more data-driven, smarter decisions.

Watson became a household name in 2011, when it famously won the television game show, *Jeopardy!* Since that proof of concept in 2011, IBM has reached agreements with the health insurance provider WellPoint (now part of Anthem), the financial services company Citibank, and Memorial Sloan-Kettering Cancer Center to apply Watson to the decision problems that they face.

Watson is a system of computing hardware, high-speed data processing, and analytical algorithms that are combined to make data-based recommendations. As more and more data are collected, Watson has the capability to learn over time. In simple terms, according to IBM, Watson gathers hundreds of thousands of possible solutions from a huge data bank, evaluates them using analytical techniques, and proposes only the best solutions for consideration. Watson provides not just a single solution, but rather a range of good solutions with a confidence level for each.

For example, at a data center in Virginia, to the delight of doctors and patients, Watson is already being used to speed up the approval of medical procedures. Citibank is beginning

to explore how to use Watson to better serve its customers, and cancer specialists at more than a dozen hospitals in North America are using Watson to assist with the diagnosis and treatment of patients.¹


This book is concerned with data-driven decision making and the use of analytical approaches in the decision-making process. Three developments spurred recent explosive growth in the use of analytical methods in business applications. First, technological advances—such as improved point-of-sale scanner technology and the collection of data through e-commerce, Internet social networks, and data generated from personal electronic devices—produce incredible amounts of data for businesses. Naturally, businesses want to use these data to improve the efficiency and profitability of their operations, better understand their customers, price their products more effectively, and gain a competitive advantage. Second, ongoing research has resulted in numerous methodological developments, including advances in computational approaches to effectively handle and explore massive amounts of data, faster algorithms for optimization and simulation, and more effective approaches for visualizing data. Third, these methodological developments were paired with an explosion in computing power and storage capability. Better computing hardware, parallel computing, and, more recently, cloud computing (the remote use of hardware and software over the Internet) have enabled businesses to solve big problems more quickly and more accurately than ever before.

In summary, the availability of massive amounts of data, improvements in analytic methodologies, and substantial increases in computing power have all come together to result in a dramatic upsurge in the use of analytical methods in business and a reliance on the discipline that is the focus of this text: business analytics. Figure 1.1 shows the job trend for analytics from 2006 to 2015. The chart from indeed.com shows the percentage of job ads that contain the word *analytics* and illustrates that demand has grown and continues to be strong for analytical skills.

Business analytics is a crucial area of study for students looking to enhance their employment prospects. It has been predicted that by 2018 there will be a shortage of more than 1.5 million business managers with adequate training in analytics in the United States alone.² As stated in the Preface, the purpose of this text is to provide

It is difficult to know for sure the cause of the large spike in analytics job ads in 2008. We do note, however, that the thought-provoking book Competing on Analytics by Davenport and Harris was published in 2007.

FIGURE 1.1 Analytics Job Trend According to Indeed.com



¹"IBM's Watson Is Learning Its Way to Saving Lives," Fastcompany web site, December 8, 2012; "IBM's Watson Targets Cancer and Enlists Prominent Providers in the Fight," ModernHealthcare web site, May 5, 2015.

²J. Manyika et al., "Big Data: The Next Frontier for Innovation, Competition and Productivity," McKinsey Global Institute Report, 2011.

students with a sound conceptual understanding of the role that business analytics plays in the decision-making process. To reinforce the applications orientation of the text and to provide a better understanding of the variety of applications in which analytical methods have been used successfully, Analytics in Action articles are presented throughout the book. Each Analytics in Action article summarizes an application of analytical methods in practice.

1.1 Decision Making

It is the responsibility of managers to plan, coordinate, organize, and lead their organizations to better performance. Ultimately, managers' responsibilities require that they make strategic, tactical, or operational decisions. **Strategic decisions** involve higher-level issues concerned with the overall direction of the organization; these decisions define the organization's overall goals and aspirations for the future. Strategic decisions are usually the domain of higher-level executives and have a time horizon of three to five years. **Tactical decisions** concern how the organization should achieve the goals and objectives set by its strategy, and they are usually the responsibility of midlevel management. Tactical decisions usually span a year and thus are revisited annually or even every six months. **Operational decisions** affect how the firm is run from day to day; they are the domain of operations managers, who are the closest to the customer.

Consider the case of the Thoroughbred Running Company (TRC). Historically, TRC had been a catalog-based retail seller of running shoes and apparel. TRC sales revenues grew quickly as it changed its emphasis from catalog-based sales to Internet-based sales. Recently, TRC decided that it should also establish retail stores in the malls and downtown areas of major cities. This strategic decision will take the firm in a new direction that it hopes will complement its Internet-based strategy. TRC middle managers will therefore have to make a variety of tactical decisions in support of this strategic decision, including how many new stores to open this year, where to open these new stores, how many distribution centers will be needed to support the new stores, and where to locate these distribution centers. Operations managers in the stores will need to make day-to-day decisions regarding, for instance, how many pairs of each model and size of shoes to order from the distribution centers and how to schedule their sales personnel's work time.

Regardless of the level within the firm, *decision making* can be defined as the following process:

1. Identify and define the problem.
2. Determine the criteria that will be used to evaluate alternative solutions.
3. Determine the set of alternative solutions.
4. Evaluate the alternatives.
5. Choose an alternative.

If I were given one hour to save the planet, I would spend 59 minutes defining the problem and one minute resolving it.

—Albert Einstein

Step 1 of decision making, identifying and defining the problem, is the most critical. Only if the problem is well-defined, with clear metrics of success or failure (step 2), can a proper approach for solving the problem (steps 3 and 4) be devised. Decision making concludes with the choice of one of the alternatives (step 5).

There are a number of approaches to making decisions: tradition ("We've always done it this way"), intuition ("gut feeling"), and rules of thumb ("As the restaurant owner, I schedule twice the number of waiters and cooks on holidays"). The power of each of these approaches should not be underestimated. Managerial experience and intuition are valuable inputs to making decisions, but what if relevant data were available to help us make more informed decisions? With the vast amounts of data now generated and stored

electronically, it is estimated that the amount of data stored by businesses more than doubles every two years. How can managers convert these data into knowledge that they can use to be more efficient and effective in managing their businesses?

1.2 Business Analytics Defined

What makes decision making difficult and challenging? Uncertainty is probably the number one challenge. If we knew how much the demand will be for our product, we could do a much better job of planning and scheduling production. If we knew exactly how long each step in a project will take to be completed, we could better predict the project's cost and completion date. If we knew how stocks will perform, investing would be a lot easier.

Another factor that makes decision making difficult is that we often face such an enormous number of alternatives that we cannot evaluate them all. What is the best combination of stocks to help me meet my financial objectives? What is the best product line for a company that wants to maximize its market share? How should an airline price its tickets so as to maximize revenue?

Business analytics is the scientific process of transforming data into insight for making better decisions.³ Business analytics is used for data-driven or fact-based decision making, which is often seen as more objective than other alternatives for decision making.

As we shall see, the tools of business analytics can aid decision making by creating insights from data, by improving our ability to more accurately forecast for planning, by helping us quantify risk, and by yielding better alternatives through analysis and optimization. A study based on a large sample of firms that was conducted by researchers at MIT's Sloan School of Management and the University of Pennsylvania, concluded that firms guided by data-driven decision making have higher productivity and market value and increased output and profitability.⁴

Some firms and industries use the simpler term, analytics. Analytics is often thought of as a broader category than business analytics, encompassing the use of analytical techniques in the sciences and engineering as well. In this text, we use business analytics and analytics synonymously.

1.3 A Categorization of Analytical Methods and Models

Business analytics can involve anything from simple reports to the most advanced optimization techniques (methods for finding the best course of action). Analytics is generally thought to comprise three broad categories of techniques: descriptive analytics, predictive analytics, and prescriptive analytics.

Descriptive Analytics

Descriptive analytics encompasses the set of techniques that describes what has happened in the past. Examples are data queries, reports, descriptive statistics, data visualization including data dashboards, some data-mining techniques, and basic what-if spreadsheet models.

A **data query** is a request for information with certain characteristics from a database. For example, a query to a manufacturing plant's database might be for all records of shipments to a particular distribution center during the month of March. This query provides descriptive information about these shipments: the number of shipments, how much was included in each shipment, the date each shipment was sent, and so on. A report summarizing relevant historical information for management might be conveyed by the use of descriptive statistics (means, measures of variation, etc.) and data-visualization tools (tables, charts, and maps). Simple descriptive statistics and data-visualization techniques can be used to find patterns or relationships in a large database.

Appendix B at the end of this book describes how to use Microsoft Access to conduct data queries.

³We adopt the definition of analytics developed by the Institute for Operations Research and the Management Sciences (INFORMS).

⁴E. Brynjolfsson, L. M. Hitt, and H. H. Kim, "Strength in Numbers: How Does Data-Driven Decisionmaking Affect Firm Performance?" (April 18, 2013). Available at SSRN, http://papers.ssrn.com/sol3/papers.cfm?abstract_id=1819486.

Data dashboards are collections of tables, charts, maps, and summary statistics that are updated as new data become available. Dashboards are used to help management monitor specific aspects of the company's performance related to their decision-making responsibilities. For corporate-level managers, daily data dashboards might summarize sales by region, current inventory levels, and other company-wide metrics; front-line managers may view dashboards that contain metrics related to staffing levels, local inventory levels, and short-term sales forecasts.

Data mining is the use of analytical techniques for better understanding patterns and relationships that exist in large data sets. For example, by analyzing text on social network platforms like Twitter, data-mining techniques (including cluster analysis and sentiment analysis) are used by companies to better understand their customers. By categorizing certain words as positive or negative and keeping track of how often those words appear in tweets, a company like Apple can better understand how its customers are feeling about a product like the Apple Watch.

Predictive Analytics

Predictive analytics consists of techniques that use models constructed from past data to predict the future or ascertain the impact of one variable on another. For example, past data on product sales may be used to construct a mathematical model to predict future sales. This mode can factor in the product's growth trajectory and seasonality based on past patterns. A packaged-food manufacturer may use point-of-sale scanner data from retail outlets to help in estimating the lift in unit sales due to coupons or sales events. Survey data and past purchase behavior may be used to help predict the market share of a new product. All of these are applications of predictive analytics.

Linear regression, time series analysis, some data-mining techniques, and simulation, often referred to as risk analysis, all fall under the banner of predictive analytics. We discuss all of these techniques in greater detail later in this text.

Data mining, previously discussed as a descriptive analytics tool, is also often used in predictive analytics. For example, a large grocery store chain might be interested in developing a targeted marketing campaign that offers a discount coupon on potato chips. By studying historical point-of-sale data, the store may be able to use data mining to predict which customers are the most likely to respond to an offer on discounted chips by purchasing higher-margin items such as beer or soft drinks in addition to the chips, thus increasing the store's overall revenue.

Simulation involves the use of probability and statistics to construct a computer model to study the impact of uncertainty on a decision. For example, banks often use simulation to model investment and default risk in order to stress-test financial models. Simulation is also often used in the pharmaceutical industry to assess the risk of introducing a new drug.

Prescriptive Analytics

Prescriptive analytics differs from descriptive or predictive analytics in that **prescriptive analytics** indicates a best course of action to take; that is, the output of a prescriptive model is a best decision. The airline industry's use of revenue management is an example of a prescriptive analytics. Airlines use past purchasing data as inputs into a model that recommends the best pricing strategy across all flights in order to maximize revenue.

Other examples of prescriptive analytics are portfolio models in finance, supply network design models in operations, and price-markdown models in retailing. Portfolio models use historical investment return data to determine which mix of investments will yield the highest expected return while controlling or limiting exposure to risk. Supply-network design models provide data about plant and distribution center locations that will

TABLE 1.1 Coverage of Business Analytics Topics in This Text

Chapter	Title	Descriptive	Predictive	Prescriptive
1	Introduction	●	●	●
2	Descriptive Statistics	●		
3	Data Visualization	●		
4	Descriptive Data Mining	●		
5	Probability: An Introduction to Modeling Uncertainty	●		
6	Statistical Inference	●		
7	Linear Regression		●	
8	Time Series and Forecasting		●	
9	Predictive Data Mining		●	
10	Spreadsheet Models	●		
11	Linear Optimization Models			●
12	Integer Optimization Models			●
13	Nonlinear Optimization Models			●
14	Simulation		●	●
15	Decision Analysis			●

minimize costs while still meeting customer service requirements. Given historical data, retail price markdown models yield revenue-maximizing discount levels and the timing of discount offers when goods have not sold as planned. All of these models are known as **optimization models**, that is, models that give the best decision subject to the constraints of the situation.


Another type of modeling in the prescriptive analytics category is **simulation optimization**, which combines the use of probability and statistics to model uncertainty with optimization techniques to find good decisions in highly complex and highly uncertain settings. Finally, the techniques of **decision analysis** can be used to develop an optimal strategy when a decision maker is faced with several decision alternatives and an uncertain set of future events. Decision analysis also employs **utility theory**, which assigns values to outcomes based on the decision maker's attitude toward risk, loss, and other factors.

In this text we cover all three areas of business analytics: descriptive, predictive, and prescriptive. Table 1.1 shows how the chapters cover the three categories.

1.4 Big Data

Walmart handles over 1 million purchase transactions per hour. Facebook processes more than 250 million picture uploads per day. Six billion cell-phone owners around the world generate vast amounts of data by calling, texting, tweeting, and browsing the web on a daily basis.⁵ As Google CEO Eric Schmidt has noted, the amount of data currently created every 48 hours is equivalent to the entire amount of data created from the dawn of civilization until the year 2003. It is through technology that we have truly been thrust into the data age. Because data can now be collected electronically, the amounts of it available are staggering. The Internet, cell phones, retail checkout scanners, surveillance video, and sensors on everything from aircraft to cars to bridges allow us to collect and store vast amounts of data in real time.

⁵SAS White Paper, "Big Data Meets Big Data Analytics," SAS Institute, 2012.

FIGURE 1.2 The 4 Vs of Big Data

Source: IBM

In the midst of all of this data collection, the new term *big data* has been created. There is no universally accepted definition of big data. However, probably the most accepted and most general definition is that **big data** is any set of data that is too large or too complex to be handled by standard data-processing techniques and typical desktop software. IBM describes the phenomenon of big data through the four Vs: volume, velocity, variety, and veracity, as shown in Figure 1.2.⁶

Volume

Because data are collected electronically, we are able to collect more of it. To be useful, these data must be stored, and this storage has led to vast quantities of data. Many companies now store in excess of 100 terabytes of data (a terabyte of data is 100,000 gigabytes).

Velocity

Real-time capture and analysis of data present unique challenges both in how data are stored and the speed with which those data can be analyzed for decision making. For example, the New York Stock Exchange collects 1 terabyte of data in a single trading session, and having current data and real-time rules for trades and predictive modeling are important for managing stock portfolios.

⁶IBM web site: http://www.ibmbigdatahub.com/sites/default/files/infographic_file/4-Vs-of-big-data.jpg.

Variety

In addition to the sheer volume and speed with which companies now collect data, more complicated types of data are now available and are proving to be of great value to businesses. Text data are collected by monitoring what is being said about a company's products or services on social media platforms such as Twitter. Audio data are collected from service calls (on a service call, you will often hear "this call may be monitored for quality control"). Video data collected by in-store video cameras are used to analyze shopping behavior. Analyzing information generated by these nontraditional sources is more complicated in part because of the processing required to transform the data into a numerical form that can be analyzed.

Veracity

Veracity has to do with how much uncertainty is in the data. For example, the data could have many missing values, which makes reliable analysis a challenge. Inconsistencies in units of measure and the lack of reliability of responses in terms of bias also increase the complexity of the data.

Businesses have realized that understanding big data can lead to a competitive advantage. Although big data represents opportunities, it also presents challenges in terms of data storage and processing, security and available analytical talent.

The four Vs indicate that big data creates challenges in terms of how these complex data can be captured, stored, and processed; secured; and then analyzed. Traditional databases more or less assume that data fit into nice rows and columns, but that is not always the case with big data. Also, the sheer volume (the first V) often means that it is not possible to store all of the data on a single computer. This has led to new technologies like **Hadoop**—an open-source programming environment that supports big data processing through distributed storage and distributed processing on clusters of computers. Essentially, Hadoop provides a divide-and-conquer approach to handling massive amounts of data, dividing the storage and processing over multiple computers. **MapReduce** is a programming model used within Hadoop that performs the two major steps for which it is named: the map step and the reduce step. The map step divides the data into manageable subsets and distributes it to the computers in the cluster (often termed nodes) for storing and processing. The reduce step collects answers from the nodes and combines them into an answer to the original problem. Without technologies like Hadoop and MapReduce, and relatively inexpensive computer power, processing big data would not be cost-effective; in some cases, processing might not even be possible.

While some sources of big data are publicly available (Twitter, weather data, etc.), much of it is private information. Medical records, bank account information, and credit card transactions, for example, are all highly confidential and must be protected from computer hackers. **Data security**, the protection of stored data from destructive forces or unauthorized users, is of critical importance to companies. For example, credit card transactions are potentially very useful for understanding consumer behavior, but compromise of these data could lead to unauthorized use of the credit card or identity theft. Data security company Datacastle estimated that the average cost of a data breach for a company in 2012 was \$7.2 million. Since 2014, companies such as Target, Anthem, JPMorgan Chase, and Home Depot have faced major data breaches costing millions of dollars.

The complexities of the 4 Vs have increased the demand for analysts, but a shortage of qualified analysts has made hiring more challenging. More companies are searching for **data scientists**, who know how to effectively process and analyze massive amounts of data

because they are well trained in both computer science and statistics. Next we discuss three examples of how companies are collecting big data for competitive advantage.

Kroger Understands Its Customers⁷ Kroger is the largest retail grocery chain in the United States. It sends over 11 million pieces of direct mail to its customers each quarter. The quarterly mailers each contain 12 coupons that are tailored to each household based on several years of shopping data obtained through its customer loyalty card program. By collecting and analyzing consumer behavior at the individual household level and better matching its coupon offers to shopper interests, Kroger has been able to realize a far higher redemption rate on its coupons. In the six-week period following distribution of the mailers, over 70% of households redeem at least one coupon, leading to an estimated coupon revenue of \$10 billion for Kroger.

MagicBand at Disney⁸ The Walt Disney Company has begun offering a wristband to visitors to its Orlando, Florida, Disney World theme park. Known as the MagicBand, the wristband contains technology that can transmit more than 40 feet and can be used to track each visitor's location in the park in real time. The band can link to information that allows Disney to better serve its visitors. For example, prior to the trip to Disney World, a visitor might be asked to fill out a survey on his or her birth date and favorite rides, characters, and restaurant table type and location. This information, linked to the MagicBand, can allow Disney employees using smartphones to greet you by name as you arrive, offer you products they know you prefer, wish you a happy birthday, have your favorite characters show up as you wait in line or have lunch at your favorite table. The MagicBand can be linked to your credit card, so there is no need to carry cash or a credit card. And during your visit, your movement throughout the park can be tracked and the data can be analyzed to better serve you during your next visit to the park.

General Electric and the Internet of Things⁹ The **Internet of Things (IoT)** is the technology that allows data, collected from sensors in all types of machines, to be sent over the Internet to repositories where it can be stored and analyzed. This ability to collect data from products has enabled the companies that produce and sell those products to better serve their customers and offer new services based on analytics. For example, each day General Electric (GE) gathers nearly 50 million pieces of data from 10 million sensors on medical equipment and aircraft engines it has sold to customers throughout the world. In the case of aircraft engines, through a service agreement with its customers, GE collects data each time an airplane powered by its engines takes off and lands. By analyzing these data, GE can better predict when maintenance is needed, which helps customers to avoid unplanned maintenance and downtime and helps ensure safe operation. GE can also use the data to better control how the plane is flown, leading to a decrease in fuel cost by flying more efficiently. In 2014, GE realized approximately \$1.1 billion in revenue from the IoT.

Although big data is clearly one of the drivers for the strong demand for analytics, it is important to understand that in some sense big data issues are a subset of analytics. Many very valuable applications of analytics do not involve big data, but rather traditional data sets that are very manageable by traditional database and analytics software. The key to analytics is that it provides useful insights and better decision making using the data that are available—whether those data are “big” or “small.”

⁷Based on “Kroger Knows Your Shopping Patterns Better Than You Do,” *Forbes.com*, October 23, 2013.

⁸Based on “Disney’s \$1 Billion Bet on a Magical Wristband,” *Wired.com*, March 10, 2015.

⁹Based on “G.E. Opens Its Big Data Platform,” *NYTimes.com*, October 9, 2014.

1.5 Business Analytics in Practice


Business analytics involves tools as simple as reports and graphs to those that are as sophisticated as optimization, data mining, and simulation. In practice, companies that apply analytics often follow a trajectory similar to that shown in Figure 1.3. Organizations start with basic analytics in the lower left. As they realize the advantages of these analytic techniques, they often progress to more sophisticated techniques in an effort to reap the derived competitive advantage. Therefore, predictive and prescriptive analytics are sometimes referred to as **advanced analytics**. Not all companies reach that level of usage, but those that embrace analytics as a competitive strategy often do.

Analytics has been applied in virtually all sectors of business and government. Organizations such as Procter & Gamble, IBM, UPS, Netflix, Amazon.com, Google, the Internal Revenue Service, and General Electric have embraced analytics to solve important problems or to achieve a competitive advantage. In this section, we briefly discuss some of the types of applications of analytics by application area.

Financial Analytics

Applications of analytics in finance are numerous and pervasive. Predictive models are used to forecast financial performance, to assess the risk of investment portfolios and projects, and to construct financial instruments such as derivatives. Prescriptive models are used to construct optimal portfolios of investments, to allocate assets, and to create optimal capital budgeting plans. For example, GE Asset Management uses optimization models to decide how to invest its own cash received from insurance policies and other financial products, as well as the cash of its clients, such as Genworth Financial. The estimated benefit from the optimization models was \$75 million over a five-year period.¹⁰ Simulation is also often used to assess risk in the financial sector; one example is the deployment by

FIGURE 1.3 The Spectrum of Business Analytics



Source: Adapted from SAS.

¹⁰L. C. Chalermkraivuth et al., "GE Asset Management, Genworth Financial, and GE Insurance Use a Sequential-Linear Programming Algorithm to Optimize Portfolios," *Interfaces* 35, no. 5 (September–October 2005): 370–80.

Hypo Real Estate International of simulation models to successfully manage commercial real estate risk.¹¹

Human Resource (HR) Analytics

A relatively new area of application for analytics is the management of an organization's human resources (HR). The HR function is charged with ensuring that the organization: (1) has the mix of skill sets necessary to meet its needs, (2) is hiring the highest-quality talent and providing an environment that retains it, and (3) achieves its organizational diversity goals. Google refers to its HR Analytics function as “people analytics.” Google has analyzed substantial data on their own employees to determine the characteristics of great leaders, to assess factors that contribute to productivity, and to evaluate potential new hires. Google also uses predictive analytics to continually update their forecast of future employee turnover and retention.¹²

Marketing Analytics

Marketing is one of the fastest-growing areas for the application of analytics. A better understanding of consumer behavior through the use of scanner data and data generated from social media has led to an increased interest in marketing analytics. As a result, descriptive, predictive, and prescriptive analytics are all heavily used in marketing. A better understanding of consumer behavior through analytics leads to the better use of advertising budgets, more effective pricing strategies, improved forecasting of demand, improved product-line management, and increased customer satisfaction and loyalty. For example, each year, NBCUniversal uses a predictive model to help support its annual upfront market—a period in late May when each television network sells the majority of its on-air advertising for the upcoming television season. Over 200 NBC sales and finance personnel use the results of the forecasting model to support pricing and sales decisions.¹³

In another example of high-impact marketing analytics, automobile manufacturer Chrysler teamed with J.D. Power and Associates to develop an innovative set of predictive models to support its pricing decisions for automobiles. These models help Chrysler to better understand the ramifications of proposed pricing structures (a combination of manufacturer's suggested retail price, interest rate offers, and rebates) and, as a result, to improve its pricing decisions. The models have generated an estimated annual savings of \$500 million.¹⁴

Figure 1.4 shows the Google Trends graph for Marketing, Financial, and HR Analytics. While interest in each of these three areas of business is increasing, the graph clearly shows the pronounced increase in the interest in marketing analytics.

Health Care Analytics

The use of analytics in health care is on the increase because of pressure to simultaneously control costs and provide more effective treatment. Descriptive, predictive, and prescriptive analytics are used to improve patient, staff, and facility scheduling; patient flow; purchasing; and inventory control. A study by McKinsey Global Institute (MGI) and McKinsey & Company¹⁵

¹¹Y. Jafry, C. Marrison, and U. Umkehrer-Neudeck, “Hypo International Strengthens Risk Management with a Large-Scale, Secure Spreadsheet-Management Framework,” *Interfaces* 38, no. 4 (July–August 2008): 281–88.

¹²J. Sullivan, “How Google Is Using People Analytics to Completely Reinvent HR,” Talent Management and HR web site, February 26, 2013.


¹³S. Bollapragada et al., “NBC-Universal Uses a Novel Qualitative Forecasting Technique to Predict Advertising Demand,” *Interfaces* 38, no. 2 (March–April 2008): 103–11.

¹⁴J. Silva-Risso et al., “Chrysler and J. D. Power: Pioneering Scientific Price Customization in the Automobile Industry,” *Interfaces* 38, no. 1 (January–February 2008): 26–39.

¹⁵J. Manyika et al., “Big Data: The Next Frontier for Innovation, Competition and Productivity,” McKinsey Global Institute Report, 2011.

FIGURE 1.4

Google Trends for Marketing, Financial, and Human Resource (HR) Analytics, 2006–2015



estimates that the health care system in the United States could save more than \$300 billion per year by better utilizing analytics; these savings are approximately the equivalent of the entire gross domestic product of countries such as Finland, Singapore, and Ireland.

The use of prescriptive analytics for diagnosis and treatment is relatively new, but it may prove to be the most important application of analytics in health care. For example, working with the Georgia Institute of Technology, Memorial Sloan-Kettering Cancer Center developed a real-time prescriptive model to determine the optimal placement of radioactive seeds for the treatment of prostate cancer.¹⁶ Using the new model, 20–30% fewer seeds are needed, resulting in a faster and less invasive procedure.

Supply-Chain Analytics

One of the earliest applications of analytics was in logistics and supply-chain management. The core service of companies such as UPS and FedEx is the efficient delivery of goods, and analytics has long been used to achieve efficiency. The optimal sorting of goods, vehicle and staff scheduling, and vehicle routing are all key to profitability for logistics companies such as UPS and FedEx.

Companies can benefit from better inventory and processing control and more efficient supply chains. Analytic tools used in this area span the entire spectrum of analytics. For example, the women's apparel manufacturer Bernard Claus, Inc., has successfully used descriptive analytics to provide its managers a visual representation of the status of its supply chain.¹⁷ ConAgra Foods uses predictive and prescriptive analytics to better plan capacity utilization by incorporating the inherent uncertainty in commodities pricing. ConAgra realized a 100% return on their investment in analytics in under three months—an unheard of result for a major technology investment.¹⁸

¹⁶E. Lee and M. Zaider, "Operations Research Advances Cancer Therapeutics," *Interfaces* 38, no. 1 (January–February 2008): 5–25.

¹⁷T. H. Davenport, ed., *Enterprise Analytics* (Upper Saddle River, NJ: Pearson Education Inc., 2013).

¹⁸"ConAgra Mills: Up-to-the-Minute Insights Drive Smarter Selling Decisions and Big Improvements in Capacity Utilization," IBM Smarter Planet Leadership Series. Available at: <http://www.ibm.com/smarterplanet/us/en/leadership/conagra/>, retrieved December 1, 2012.

Analytics for Government and Nonprofits

Government agencies and other nonprofits have used analytics to drive out inefficiencies and increase the effectiveness and accountability of programs. Indeed, much of advanced analytics has its roots in the U.S. and English military dating back to World War II. Today, the use of analytics in government is becoming pervasive in everything from elections to tax collection. For example, the New York State Department of Taxation and Finance has worked with IBM to use prescriptive analytics in the development of a more effective approach to tax collection. The result was an increase in collections from delinquent payers of \$83 million over two years.¹⁹ The U.S. Internal Revenue Service has used data mining to identify patterns that distinguish questionable annual personal income tax filings. In one application, the IRS combines its data on individual taxpayers with data received from banks, on mortgage payments made by those taxpayers. When taxpayers report a mortgage payment that is unrealistically high relative to their reported taxable income, they are flagged as possible underreporters of taxable income. The filing is then further scrutinized and may trigger an audit.

Likewise, nonprofit agencies have used analytics to ensure their effectiveness and accountability to their donors and clients. Catholic Relief Services (CRS) is the official international humanitarian agency of the U.S. Catholic community. The CRS mission is to provide relief for the victims of both natural and human-made disasters and to help people in need around the world through its health, educational, and agricultural programs. CRS uses an analytical spreadsheet model to assist in the allocation of its annual budget based on the impact that its various relief efforts and programs will have in different countries.²⁰

Sports Analytics

The use of analytics in sports has gained considerable notoriety since 2003 when renowned author Michael Lewis published *Moneyball*. Lewis' book tells the story of how the Oakland Athletics used an analytical approach to player evaluation in order to assemble a competitive team with a limited budget. The use of analytics for player evaluation and on-field strategy is now common, especially in professional sports. Professional sports teams use analytics to assess players for the amateur drafts and to decide how much to offer players in contract negotiations;²¹ professional motorcycle racing teams use sophisticated optimization for gearbox design to gain competitive advantage;²² and teams use analytics to assist with on-field decisions such as which pitchers to use in various games of a Major League Baseball playoff series.

The use of analytics for off-the-field business decisions is also increasing rapidly. Ensuring customer satisfaction is important for any company, and fans are the customers of sports teams. The Cleveland Indians professional baseball team used a type of predictive modeling known as conjoint analysis to design its premium seating offerings at Progressive Field based on fan survey data. Using prescriptive analytics, franchises across several major sports dynamically adjust ticket prices throughout the season to reflect the relative attractiveness and potential demand for each game.

¹⁹G. Miller et al., "Tax Collection Optimization for New York State," *Interfaces* 42, no. 1 (January–February 2013): 74–84.

²⁰J. Gamvros, R. Nidel, and S. Raghavan, "Investment Analysis and Budget Allocation at Catholic Relief Services," *Interfaces* 36, no. 5 (September–October 2006): 400–406.

²¹N. Streib, S. J. Young, and J. Sokol, "A Major League Baseball Team Uses Operations Research to Improve Draft Preparation," *Interfaces* 42, no. 2 (March–April 2012): 119–30.

²²J. Amoros, L. F. Escudero, J. F. Monge, J. V. Segura, and O. Reinoso, "TEAM ASPAR Uses Binary Optimization to Obtain Optimal Gearbox Ratios in Motorcycle Racing," *Interfaces* 42, no. 2 (March–April 2012): 191–98.

Web Analytics

Web analytics is the analysis of online activity, which includes, but is not limited to, visits to web sites and social media sites such as Facebook and LinkedIn. Web analytics obviously has huge implications for promoting and selling products and services via the Internet. Leading companies apply descriptive and advanced analytics to data collected in online experiments to determine the best way to configure web sites, position ads, and utilize social networks for the promotion of products and services. Online experimentation involves exposing various subgroups to different versions of a web site and tracking the results. Because of the massive pool of Internet users, experiments can be conducted without risking the disruption of the overall business of the company. Such experiments are proving to be invaluable because they enable the company to use trial-and-error in determining statistically what makes a difference in their web site traffic and sales.

S U M M A R Y

This introductory chapter began with a discussion of decision making. Decision making can be defined as the following process: (1) identify and define the problem; (2) determine the criteria that will be used to evaluate alternative solutions; (3) determine the set of alternative solutions; (4) evaluate the alternatives; and (5) choose an alternative. Decisions may be strategic (high-level, concerned with the overall direction of the firm), tactical (midlevel, concerned with how to achieve the strategic goals of the firm), or operational (day-to-day decisions that must be made to run the company).

Uncertainty and an overwhelming number of alternatives are two key factors that make decision making difficult. Business analytics approaches can assist by identifying and mitigating uncertainty and by prescribing the best course of action from a very large number of alternatives. In short, business analytics can help us make better-informed decisions.

There are three categories of analytics: descriptive, predictive, and prescriptive. Descriptive analytics describes what has happened and includes tools such as reports, data visualization, data dashboards, descriptive statistics, and some data-mining techniques. Predictive analytics consists of techniques that use past data to predict future events and include regression, data mining, forecasting, and simulation. Prescriptive analytics uses data to determine a best course of action. This class of analytical techniques includes simulation, decision analysis, and optimization. Descriptive and predictive analytics can help us better understand the uncertainty and risk associated with our decision alternatives. Predictive and prescriptive analytics, also often referred to as advanced analytics, can help us make the best decision when facing a myriad of alternatives.

Big data is a set of data that is too large or too complex to be handled by standard data-processing techniques or typical desktop software. The increasing prevalence of big data is leading to an increase in the use of analytics. The Internet, retail scanners, and cell phones are making huge amounts of data available to companies, and these companies want to better understand these data. Business analytics is helping them understand these data and use them to make better decisions.

We concluded this chapter with a discussion of various application areas of analytics. Our discussion focused on financial analytics, human resource analytics, marketing analytics, health care analytics, supply-chain analytics, analytics for government and nonprofit organizations, sports analytics, and web analytics. However, the use of analytics is rapidly spreading to other sectors, industries, and functional areas of organizations. Each remaining chapter in this text will provide a real-world vignette in which business analytics is applied to a problem faced by a real organization.

G L O S S A R Y

Advanced analytics Predictive and prescriptive analytics.

Big data Any set of data that is too large or too complex to be handled by standard data-processing techniques and typical desktop software.

Business analytics The scientific process of transforming data into insight for making better decisions.

Data dashboard A collection of tables, charts, and maps to help management monitor selected aspects of the company's performance.

Data mining The use of analytical techniques for better understanding patterns and relationships that exist in large data sets.

Data query A request for information with certain characteristics from a database.

Data scientists Analysts trained in both computer science and statistics who know how to effectively process and analyze massive amounts of data.

Data security Protecting stored data from destructive forces or unauthorized users.

Decision analysis A technique used to develop an optimal strategy when a decision maker is faced with several decision alternatives and an uncertain set of future events.

Descriptive analytics Analytical tools that describe what has happened.

Hadoop An open-source programming environment that supports big data processing through distributed storage and distributed processing on clusters of computers.

Internet of Things (IoT) The technology that allows data collected from sensors in all types of machines to be sent over the Internet to repositories where it can be stored and analyzed.

MapReduce Programming model used within Hadoop that performs the two major steps for which it is named: the map step and the reduce step. The map step divides the data into manageable subsets and distributes it to the computers in the cluster for storing and processing. The reduce step collects answers from the nodes and combines them into an answer to the original problem.

Operational decision A decision concerned with how the organization is run from day to day.

Optimization model A mathematical model that gives the best decision, subject to the situation's constraints.

Predictive analytics Techniques that use models constructed from past data to predict the future or to ascertain the impact of one variable on another.

Prescriptive analytics Techniques that analyze input data and yield a best course of action.

Simulation The use of probability and statistics to construct a computer model to study the impact of uncertainty on the decision at hand.

Simulation optimization The use of probability and statistics to model uncertainty, combined with optimization techniques, to find good decisions in highly complex and highly uncertain settings.

Strategic decision A decision that involves higher-level issues and that is concerned with the overall direction of the organization, defining the overall goals and aspirations for the organization's future.

Tactical decision A decision concerned with how the organization should achieve the goals and objectives set by its strategy.

Utility theory The study of the total worth or relative desirability of a particular outcome that reflects the decision maker's attitude toward a collection of factors such as profit, loss, and risk.

Chapter 2

Descriptive Statistics

CONTENTS

2.1 OVERVIEW OF USING DATA: DEFINITIONS AND GOALS

2.2 TYPES OF DATA

- Population and Sample Data
- Quantitative and Categorical Data
- Cross-Sectional and Time Series Data
- Sources of Data

2.3 MODIFYING DATA IN EXCEL

- Sorting and Filtering Data in Excel
- Conditional Formatting of Data in Excel

2.4 CREATING DISTRIBUTIONS FROM DATA

- Frequency Distributions for Categorical Data
- Relative Frequency and Percent Frequency Distributions
- Frequency Distributions for Quantitative Data
- Histograms
- Cumulative Distributions

2.5 MEASURES OF LOCATION

- Mean (Arithmetic Mean)
- Median
- Mode
- Geometric Mean

2.6 MEASURES OF VARIABILITY

- Range
- Variance
- Standard Deviation
- Coefficient of Variation

2.7 ANALYZING DISTRIBUTIONS

- Percentiles
- Quartiles
- z-Scores
- Empirical Rule
- Identifying Outliers
- Box Plots

2.8 MEASURES OF ASSOCIATION BETWEEN TWO VARIABLES

- Scatter Charts
- Covariance
- Correlation Coefficient

APPENDIX 2.1: CREATING BOX PLOTS WITH XLMINER

ANALYTICS IN ACTION

U.S. Census Bureau

The Bureau of the Census is part of the U.S. Department of Commerce and is more commonly known as the U.S. Census Bureau. The U.S. Census Bureau collects data related to the population and economy of the United States using a variety of methods and for many purposes. These data are essential to many government and business decisions.

Probably the best-known data collected by the U.S. Census Bureau is the decennial census, which is an effort to count the total U.S. population. Collecting these data is a huge undertaking involving mailings, door-to-door visits, and other methods. The decennial census collects categorical data such as the sex and race of the respondents, as well as quantitative data such as the number of people living in the household. The data collected in the decennial census are used to determine the number of representatives assigned to each state, the number of Electoral College votes apportioned to each state, and how federal government funding is divided among communities.

The U.S. Census Bureau also administers the Current Population Survey (CPS). The CPS is a cross-sectional monthly survey of a sample of 60,000 households used to estimate employment and unemployment rates in different geographic areas. The CPS has been administered since 1940, so an extensive time series of employment and unemployment data now exists. These

data drive government policies such as job assistance programs. The estimated unemployment rates are watched closely as an overall indicator of the health of the U.S. economy.

The data collected by the U.S. Census Bureau are also very useful to businesses. Retailers use data on population changes in different areas to plan new store openings. Mail-order catalog companies use the demographic data when designing targeted marketing campaigns. In many cases, businesses combine the data collected by the U.S. Census Bureau with their own data on customer behavior to plan strategies and to identify potential customers. The U.S. Census Bureau is one of the most important providers of data used in business analytics.

In this chapter, we first explain the need to collect and analyze data and identify some common sources of data. Then we discuss the types of data that you may encounter in practice and present several numerical measures for summarizing data. We cover some common ways of manipulating and summarizing data using spreadsheets. We then develop numerical summary measures for data sets consisting of a single variable. When a data set contains more than one variable, the same numerical measures can be computed separately for each variable. In the two-variable case, we also develop measures of the relationship between the variables.

2.1 Overview of Using Data: Definitions and Goals

Data are the facts and figures collected, analyzed, and summarized for presentation and interpretation. Table 2.1 shows a data set containing 2013 information for stocks in the Dow Jones Industrial Index (or simply “the Dow”). The Dow is tracked by many financial advisors and investors as an indication of the state of the overall financial markets and the economy in the United States. The share prices for the 30 companies listed in Table 2.1 are the basis for computing the Dow Jones Industrial Average (DJI), which is tracked continuously by virtually every financial publication.

A characteristic or a quantity of interest that can take on different values is known as a **variable**; for the data in Table 2.1, the variables are Symbol, Industry, Share Price, and Volume. An **observation** is a set of values corresponding to a set of variables; each row in Table 2.1 corresponds to an observation.

Practically every problem (and opportunity) that an organization (or individual) faces is concerned with the impact of the possible values of relevant variables on the business outcome. Thus, we are concerned with how the value of a variable can vary; **variation** is the difference in a variable measured over observations (time, customers, items, etc.).

TABLE 2.1 Data for Dow Jones Industrial Index Companies

Company	Symbol	Industry	Share Price (\$)	Volume
Apple	AAPL	Technology	124.50	42,162,332
American Express	AXP	Financial	75.90	8,639,908
Boeing	BA	Manufacturing	144.06	2,454,976
Caterpillar	CAT	Manufacturing	76.10	9,175,903
Cisco Systems	CSCO	Technology	28.40	39,471,714
Chevron Corporation	CVX	Chemical, Oil, and Gas	90.60	11,158,429
DuPont	DD	Chemical, Oil, and Gas	56.94	5,352,510
Walt Disney	DIS	Entertainment	118.91	4,320,854
General Electric	GE	Conglomerate	25.75	31,124,222
Goldman Sachs	GS	Financial	207.35	2,453,880
The Home Depot	HD	Retail	113.59	4,427,770
IBM	IBM	Technology	159.75	3,778,186
Intel	INTC	Technology	28.06	31,621,031
Johnson & Johnson	JNJ	Pharmaceuticals	99.15	6,524,173
JPMorgan Chase	JPM	Banking	68.91	12,413,896
Coca-Cola	KO	Food and Drink	40.44	10,912,528
McDonald's	MCD	Food and Drink	96.10	5,554,624
3M	MMM	Conglomerate	149.33	3,433,648
Merck	MRK	Pharmaceuticals	57.41	7,847,073
Microsoft	MSFT	Technology	45.94	21,428,146
Nike	NKE	Consumer Goods	112.99	2,983,621
Pfizer	PFE	Pharmaceuticals	34.26	21,428,146
Procter & Gamble	PG	Consumer Goods	80.29	5,660,786
Travelers	TRV	Insurance	105.27	1,604,998
UnitedHealth Group	UNH	Healthcare	117.94	3,840,567
United Technologies	UTX	Conglomerate	99.31	6,588,011
Visa	V	Financial	74.80	21,196,114
Verizon	VZ	Telecommunications	46.04	19,528,682
Wal-Mart	WMT	Retail	71.58	5,951,117
ExxonMobil	XOM	Chemical, Oil, and Gas	79.94	14,888,464

The role of descriptive analytics is to collect and analyze data to gain a better understanding of variation and its impact on the business setting. The values of some variables are under direct control of the decision maker (these are often called decision variables, as discussed in Chapters 11, 12, and 13). The values of other variables may fluctuate with uncertainty because of factors outside the direct control of the decision maker. In general, a quantity whose values are not known with certainty is called a **random variable**, or **uncertain variable**. Random variables are discussed in greater detail in Chapters 5 and 14. When we collect data, we are gathering past observed values, or realizations of a variable. By collecting these past realizations of one or more variables, our goal is to learn more about the variation of a particular business situation.

To ensure that the companies in the Dow form a representative sample, companies are periodically added and removed from the Dow. It is possible that the companies in the Dow today have changed from what is shown in Table 2.1.

2.2 Types of Data

Population and Sample Data

Data can be categorized in several ways based on how they are collected and the type collected. In many cases, it is not feasible to collect data from the **population** of all elements of interest. In such instances, we collect data from a subset of the population known as a **sample**. For example, with the thousands of publicly traded companies in the United States, tracking and analyzing all of these stocks every day would be too time consuming and expensive. The Dow represents a sample of 30 stocks of large public companies based in the United States, and it is often interpreted to represent the larger population of all publicly traded companies. It is very important to collect sample data that are representative of the population data so that generalizations can be made from them. In most cases (although not true of the Dow), a representative sample can be gathered by **random sampling** from the population data. Dealing with populations and samples can introduce subtle differences in how we calculate and interpret summary statistics. In almost all practical applications of business analytics, we will be dealing with sample data.

Quantitative and Categorical Data

Data are considered **quantitative data** if numeric and arithmetic operations, such as addition, subtraction, multiplication, and division, can be performed on them. For instance, we can sum the values for Volume in the Dow data in Table 2.1 to calculate a total volume of all shares traded by companies included in the Dow. If arithmetic operations cannot be performed on the data, they are considered **categorical data**. We can summarize categorical data by counting the number of observations or computing the proportions of observations in each category. For instance, the data in the Industry column in Table 2.1 are categorical. We can count the number of companies in the Dow that are in the telecommunications industry. Table 2.1 shows three companies in the financial industry: American Express, Goldman Sachs, and Visa. We cannot perform arithmetic operations on the data in the Industry column.

Cross-Sectional and Time Series Data

For statistical analysis, it is important to distinguish between cross-sectional data and time series data. **Cross-sectional data** are collected from several entities at the same, or approximately the same, point in time. The data in Table 2.1 are cross-sectional because they describe the 30 companies that comprise the Dow at the same point in time (July 2015). **Time series data** are collected over several time periods. Graphs of time series data are frequently found in business and economic publications. Such graphs help analysts understand what happened in the past, identify trends over time, and project future levels for the time series. For example, the graph of the time series in Figure 2.1 shows the DJI value from January 2005 to June 2015. The figure illustrates that the DJI was between 10,000 and 11,000 in 2005 and climbed to above 14,000 in 2007. However, the financial crisis in 2008 led to a significant decline in the DJI to between 6,000 and 7,000 by 2009. Since 2009, the DJI has been generally increasing and topped 18,000 in 2015.

Sources of Data

Data necessary to analyze a business problem or opportunity can often be obtained with an appropriate study; such statistical studies can be classified as either experimental or observational. In an *experimental study*, a variable of interest is first identified. Then one or more other variables are identified and controlled or manipulated to obtain data about how these variables influence the variable of interest. For example, if a pharmaceutical firm conducts an experiment to learn about how a new drug affects blood pressure, then blood pressure is the variable of interest. The dosage level of the new drug is another variable

FIGURE 2.1 Dow Jones Index Values Since 2005

that is hoped to have a causal effect on blood pressure. To obtain data about the effect of the new drug, researchers select a sample of individuals. The dosage level of the new drug is controlled by giving different dosages to the different groups of individuals. Before and after the study, data on blood pressure are collected for each group. Statistical analysis of these experimental data can help determine how the new drug affects blood pressure.

Nonexperimental, or observational, studies make no attempt to control the variables of interest. A survey is perhaps the most common type of observational study. For instance, in a personal interview survey, research questions are first identified. Then a questionnaire is designed and administered to a sample of individuals. Some restaurants use observational studies to obtain data about customer opinions with regard to the quality of food, quality of service, atmosphere, and so on. A customer opinion questionnaire used by Chops City Grill in Naples, Florida, is shown in Figure 2.2. Note that the customers who fill out the questionnaire are asked to provide ratings for 12 variables, including overall experience, the greeting by hostess, the table visit by the manager, overall service, and so on. The response categories of excellent, good, average, fair, and poor provide categorical data that enable Chops City Grill management to maintain high standards for the restaurant's food and service.

In some cases, the data needed for a particular application exist from an experimental or observational study that has already been conducted. For example, companies maintain a variety of databases about their employees, customers, and business operations. Data on employee salaries, ages, and years of experience can usually be obtained from internal personnel records. Other internal records contain data on sales, advertising expenditures, distribution costs, inventory levels, and production quantities. Most companies also maintain detailed data about their customers.

Anyone who wants to use data and statistical analysis to aid in decision making must be aware of the time and cost required to obtain the data. The use of existing data sources is desirable when data must be obtained in a relatively short period of time. If important data are not readily available from a reliable existing source, the additional time and cost involved in obtaining the data must be taken into account. In all cases, the decision maker should consider the potential contribution of the statistical analysis to the decision-making process. In Chapter 15 we discuss methods for determining the value of additional information that can be provided by collecting data. The cost of data acquisition and the subsequent statistical analysis should not exceed the savings generated by using the information to make a better decision.

FIGURE 2.2 Customer Opinion Questionnaire Used by Chops City Grill Restaurant



Date: _____

Server Name: _____

*O*ur customers are our top priority. Please take a moment to fill out our survey card, so we can better serve your needs. You may return this card to the front desk or return by mail. Thank you!

SERVICE SURVEY	Excellent	Good	Average	Fair	Poor
Overall Experience	☐	☐	☐	☐	☐
Greeting by Hostess	☐	☐	☐	☐	☐
Manager (Table Visit)	☐	☐	☐	☐	☐
Overall Service	☐	☐	☐	☐	☐
Professionalism	☐	☐	☐	☐	☐
Menu Knowledge	☐	☐	☐	☐	☐
Friendliness	☐	☐	☐	☐	☐
Wine Selection	☐	☐	☐	☐	☐
Menu Selection	☐	☐	☐	☐	☐
Food Quality	☐	☐	☐	☐	☐
Food Presentation	☐	☐	☐	☐	☐
Value for \$ Spent	☐	☐	☐	☐	☐

What comments could you give us to improve our restaurant?

Thank you, we appreciate your comments. —The staff of Chops City Grill.

NOTES + COMMENTS

1. Organizations that specialize in collecting and maintaining data make available substantial amounts of business and economic data. Companies can access these external data sources through leasing arrangements or by purchase. Dun & Bradstreet, Bloomberg, and Dow Jones & Company are three firms that provide extensive business database services to clients. Nielsen and Ipsos are two companies that have built successful businesses collecting and processing data that they sell to advertisers and product manufacturers. Data are also available from a variety of industry associations and special-interest organizations.

2. Government agencies are another important source of existing data. For instance, the web site data.gov was launched by the U.S. government in 2009 to make it easier for the public to access data collected by the U.S. federal government. The data.gov web site includes hundreds of thousands of data sets from a variety of U.S. federal departments and agencies. In general, the Internet is an important source of data and statistical information. One can obtain access to stock quotes, meal prices at restaurants, salary data, and a wide array of other information simply by performing an Internet search.

2.3 Modifying Data in Excel

Projects often involve so much data that it is difficult to analyze all of the data at once. In this section, we examine methods for summarizing and manipulating data using Excel to make the data more manageable and to develop insights.

Sorting and Filtering Data in Excel

Excel contains many useful features for sorting and filtering data so that one can more easily identify patterns. Table 2.2 contains data on the 20 top-selling automobiles in the United States in March 2011. The table shows the model and manufacturer of each automobile as well as the sales for the model in March 2011 and March 2010.

Figure 2.3 shows the data from Table 2.2 entered into an Excel spreadsheet, and the percent change in sales for each model from March 2010 to March 2011 has been calculated. This is done by entering the formula $= (D2-E2)/E2$ in cell F2 and then copying the contents of this cell to cells F3 to F20. (We cannot calculate the percent change in sales for the Ford Fiesta because it was not being sold in March 2010.)

Suppose that we want to sort these automobiles by March 2010 sales instead of by March 2011 sales. To do this, we use Excel's Sort function, as shown in the following steps.

Step 1. Select cells A1:F21

Step 2. Click the **Data** tab in the Ribbon



TABLE 2.2 20 Top-Selling Automobiles in United States in March 2011				
Rank (by March 2011 Sales)	Manufacturer	Model	Sales (March 2011)	Sales (March 2010)
1	Honda	Accord	33,616	29,120
2	Nissan	Altima	32,289	24,649
3	Toyota	Camry	31,464	36,251
4	Honda	Civic	31,213	22,463
5	Toyota	Corolla/Matrix	30,234	29,623
6	Ford	Fusion	27,566	22,773
7	Hyundai	Sonata	22,894	18,935
8	Hyundai	Elantra	19,255	8,225
9	Toyota	Prius	18,605	11,786
10	Chevrolet	Cruze/Cobalt	18,101	10,316
11	Chevrolet	Impala	18,063	15,594
12	Nissan	Sentra	17,851	8,721
13	Ford	Focus	17,178	19,500
14	Volkswagon	Jetta	16,969	9,196
15	Chevrolet	Malibu	15,551	17,750
16	Mazda	3	12,467	11,353
17	Nissan	Versa	11,075	13,811
18	Subaru	Outback	10,498	7,619
19	Kia	Soul	10,028	5,106
20	Ford	Fiesta	9,787	0

Source: *Manufacturers and Automotive News Data Center*

FIGURE 2.3

Data for 20 Top-Selling Automobiles Entered into Excel with Percent Change in Sales from 2010

DATA file
Top20CarsPercent

	A	B	C	D	E	F
1	Rank (by March 2011 Sales)	Manufacturer	Model	Sales (March 2011)	Sales (March 2010)	Percent Change in Sales from 2010
2	1	Honda	Accord	33616	29120	15.4%
3	2	Nissan	Altima	32289	24649	31.0%
4	3	Toyota	Camry	31464	36251	-13.2%
5	4	Honda	Civic	31213	22463	39.0%
6	5	Toyota	Corolla/Matrix	30234	29623	2.1%
7	6	Ford	Fusion	27566	22773	21.0%
8	7	Hyundai	Sonata	22894	18935	20.9%
9	8	Hyundai	Elantra	19255	8225	134.1%
10	9	Toyota	Prius	18605	11786	57.9%
11	10	Chevrolet	Cruze/Cobalt	18101	10316	75.5%
12	11	Chevrolet	Impala	18063	15594	15.8%
13	12	Nissan	Sentra	17851	8721	104.7%
14	13	Ford	Focus	17178	19500	-11.9%
15	14	Volkswagon	Jetta	16969	9196	84.5%
16	15	Chevrolet	Malibu	15551	17750	-12.4%
17	16	Mazda	3	12467	11353	9.8%
18	17	Nissan	Versa	11075	13811	-19.8%
19	18	Subaru	Outback	10498	7619	37.8%
20	19	Kia	Soul	10028	5106	96.4%
21	20	Ford	Fiesta	9787	0	----

Step 3. Click **Sort** in the **Sort & Filter** group

Step 4. Select the check box for **My data has headers**

Step 5. In the first **Sort by** dropdown menu, select **Sales (March 2010)**

Step 6. In the **Order** dropdown menu, select **Largest to Smallest** (see Figure 2.4)

Step 7. Click **OK**

The result of using Excel's Sort function for the March 2010 data is shown in Figure 2.5. Now we can easily see that, although the Honda Accord was the best-selling automobile in March 2011, both the Toyota Camry and the Toyota Corolla/Matrix outsold the Honda Accord in March 2010. Note that while we sorted on Sales (March 2010), which is in column E, the data in all other columns are adjusted accordingly.

Now let's suppose that we are interested only in seeing the sales of models made by Toyota. We can do this using Excel's Filter function:

Step 1. Select cells A1:F21

Step 2. Click the **Data** tab in the Ribbon

Step 3. Click **Filter** in the **Sort & Filter** group

Step 4. Click on the **Filter Arrow**  in column B, next to **Manufacturer**

Step 5. If all choices are checked, you can easily deselect all choices by unchecking (**Select All**). Then select only the check box for **Toyota**.

Step 6. Click **OK**

FIGURE 2.4 Using Excel's Sort Function to Sort the Top-Selling Automobiles Data

	A	B	C	D	E	F	G
1	Rank (by March 2011 Sales)	Manufacturer	Model	Sales (March 2011)	Sales (March 2010)	Percent Change in Sales from 2010	
2	1	Honda	Accord	33616	29120	15.4%	
3	2	Nissan	Altima	32289	24649	31.0%	
4	3	Toyota	Camry	31464	36251	-13.2%	
5	4	Honda	Civic	31213	22463	39.0%	
6	5	Toyota	Corolla/Matrix	30234	29623	2.1%	
7	6	Ford					
8	7	Hyundai					
9	8	Hyundai					
10	9	Toyota					
11	10	Chevrolet					
12	11	Chevrolet					
13	12	Nissan					
14	13	Ford					
15	14	Volkswagon					
16	15	Chevrolet					
17	16	Mazda					
18	17	Nissan					
19	18	Subaru	Outback	10498	7619	37.8%	
20	19	Kia	Soul	10028	5106	96.4%	
21	20	Ford	Fiesta	9787	0	----	

Sort

☒ My data has headers

Column	Sort On	Order
Sort by: Sales (March 2010)	Values	Largest to Smallest

OK Cancel

The result is a display of only the data for models made by Toyota (see Figure 2.6). We now see that of the 20 top-selling models in March 2011, Toyota made three of them. We can further filter the data by choosing the down arrows in the other columns. We can make all data visible again by clicking on the down arrow in column B and checking **(Select All)** and clicking **OK**, or by clicking **Filter** in the **Sort & Filter** Group again from the **Data** tab.

Conditional Formatting of Data in Excel

Conditional formatting in Excel can make it easy to identify data that satisfy certain conditions in a data set. For instance, suppose that we wanted to quickly identify the automobile models in Table 2.2 for which sales had decreased from March 2010 to March 2011. We can quickly highlight these models:

- Step 1.** Starting with the original data shown in Figure 2.3, select cells F1:F21
- Step 2.** Click the **Home** tab in the Ribbon
- Step 3.** Click **Conditional Formatting** in the **Styles** group
- Step 4.** Select **Highlight Cells Rules**, and click **Less Than** from the dropdown menu
- Step 5.** Enter *0%* in the **Format cells that are LESS THAN:** box
- Step 6.** Click **OK**

FIGURE 2.5 Top-Selling Automobiles Data Sorted by Sales in March 2010 Sales

	A	B	C	D	E	F
1	Rank (by March 2011 Sales)	Manufacturer	Model	Sales (March 2011)	Sales (March 2010)	Percent Change in Sales from 2010
2	3	Toyota	Camry	31464	36251	-13.2%
3	5	Toyota	Corolla/Matrix	30234	29623	2.1%
4	1	Honda	Accord	33616	29120	15.4%
5	2	Nissan	Altima	32289	24649	31.0%
6	6	Ford	Fusion	27566	22773	21.0%
7	4	Honda	Civic	31213	22463	39.0%
8	13	Ford	Focus	17178	19500	-11.9%
9	7	Hyundai	Sonata	22894	18935	20.9%
10	15	Chevrolet	Malibu	15551	17750	-12.4%
11	11	Chevrolet	Impala	18063	15594	15.8%
12	17	Nissan	Versa	11075	13811	-19.8%
13	9	Toyota	Prius	18605	11786	57.9%
14	16	Mazda	3	12467	11353	9.8%
15	10	Chevrolet	Cruze/Cobalt	18101	10316	75.5%
16	14	Volkswagon	Jetta	16969	9196	84.5%
17	12	Nissan	Sentra	17851	8721	104.7%
18	8	Hyundai	Elantra	19255	8225	134.1%
19	18	Subaru	Outback	10498	7619	37.8%
20	19	Kia	Soul	10028	5106	96.4%
21	20	Ford	Fiesta	9787	0	----

FIGURE 2.6 Top-Selling Automobiles Data Filtered to Show Only Automobiles Manufactured by Toyota

	A	B	C	D	E	F
1	Rank (by March 2011 Sales) ▼	Manufacturer ▼	Model ▼	Sales (March 2011) ▼	Sales (March 2010) ▼	Percent Change in Sales from 2010 ▼
2	3	Toyota	Camry	31464	36251	-13.2%
3	5	Toyota	Corolla/Matrix	30234	29623	2.1%
13	9	Toyota	Prius	18605	11786	57.9%

The results are shown in Figure 2.7. Here we see that the models with decreasing sales (Toyota Camry, Ford Focus, Chevrolet Malibu, and Nissan Versa) are now clearly visible. Note that Excel's Conditional Formatting function offers tremendous flexibility. Instead of highlighting only models with decreasing sales, we could instead choose **Data Bars**

FIGURE 2.7 Using Conditional Formatting in Excel to Highlight Automobiles with Declining Sales from March 2010

	A	B	C	D	E	F
1	Rank (by March 2011 Sales)	Manufacturer	Model	Sales (March 2011)	Sales (March 2010)	Percent Change in Sales from 2010
2	1	Honda	Accord	33616	29120	15.4%
3	2	Nissan	Altima	32289	24649	31.0%
4	3	Toyota	Camry	31464	36251	-13.2%
5	4	Honda	Civic	31213	22463	39.0%
6	5	Toyota	Corolla/Matrix	30234	29623	2.1%
7	6	Ford	Fusion	27566	22773	21.0%
8	7	Hyundai	Sonata	22894	18935	20.9%
9	8	Hyundai	Elantra	19255	8225	134.1%
10	9	Toyota	Prius	18605	11786	57.9%
11	10	Chevrolet	Cruze/Cobalt	18101	10316	75.5%
12	11	Chevrolet	Impala	18063	15594	15.8%
13	12	Nissan	Sentra	17851	8721	104.7%
14	13	Ford	Focus	17178	19500	-11.9%
15	14	Volkswagon	Jetta	16969	9196	84.5%
16	15	Chevrolet	Malibu	15551	17750	-12.4%
17	16	Mazda	3	12467	11353	9.8%
18	17	Nissan	Versa	11075	13811	-19.8%
19	18	Subaru	Outback	10498	7619	37.8%
20	19	Kia	Soul	10028	5106	96.4%
21	20	Ford	Fiesta	9787	0	----

Bar charts and other graphical presentations will be covered in detail in Chapter 3. We will see other uses for Conditional Formatting in Excel in Chapter 3.

from the **Conditional Formatting** dropdown menu in the **Styles Group** of the **Home** tab in the Ribbon. The result of using the **Blue Data Bar Gradient Fill**  option is shown in Figure 2.8. Data bars are essentially a bar chart input into the cells that shows the magnitude of the cell values. The widths of the bars in this display are comparable to the values of the variable for which the bars have been drawn; a value of 20 creates a bar twice as wide as that for a value of 10. Negative values are shown to the left side of the axis; positive values are shown to the right. Cells with negative values are shaded in red, and those with positive values are shaded in blue. Again, we can easily see which models had decreasing sales, but Data Bars also provide us with a visual representation of the magnitude of the change in sales. Many other Conditional Formatting options are available in Excel.

The **Quick Analysis** button  is a feature available in Excel 2013 and Excel 2016. The button appears just outside the bottom-right corner of a group of selected cells whenever you select multiple cells. Clicking the **Quick Analysis** button gives you shortcuts for Conditional Formatting, adding Data Bars, and other operations. Clicking on this button gives you the options shown in Figure 2.9 for **Formatting**. Note that there are also tabs for **Charts**, **Totals**, **Tables**, and **Sparklines**. Many of these functions will be covered in Chapter 3.