

Making Everything Easier!™

R

FOR

DUMMIES®

Learn to:

- Use R for data analysis and processing
- Write functions and scripts for repeatable analysis
- Create high-quality charts and graphics
- Perform statistical analysis and build models

Andrie de Vries
Joris Meys



Get More and Do More at Dummies.com®



Start with **FREE** Cheat Sheets

Cheat Sheets include

- Checklists
- Charts
- Common Instructions
- And Other Good Stuff!

To access the Cheat Sheet created specifically for this book, go to
www.dummies.com/cheatsheet/r

Get Smart at Dummies.com

Dummies.com makes your life easier with 1,000s of answers on everything from removing wallpaper to using the latest version of Windows.

Check out our

- Videos
- Illustrated Articles
- Step-by-Step Instructions

Plus, each month you can win valuable prizes by entering our Dummies.com sweepstakes. *

Want a weekly dose of Dummies? Sign up for Newsletters on

- Digital Photography
- Microsoft Windows & Office
- Personal Finance & Investing
- Health & Wellness
- Computing, iPods & Cell Phones
- eBay
- Internet
- Food, Home & Garden

Find out “HOW” at Dummies.com

*Sweepstakes not currently available in all countries; visit Dummies.com for official rules.



R FOR DUMMIES®

by Andrie de Vries and Joris Meys



WILEY

A John Wiley and Sons, Ltd, Publication

R For Dummies®

Published by
John Wiley & Sons, Ltd.
The Atrium
Southern Gate
Chichester
West Sussex
PO19 8SQ
England

Email (for orders and customer service enquires): cs-books@wiley.co.uk

Visit our home page on www.wiley.com

Copyright © 2012 John Wiley & Sons, Ltd, Chichester, West Sussex, England

Published by John Wiley & Sons Ltd, Chichester, West Sussex

All rights reserved. No part of this publication may be reproduced, stored in a retrieval system or transmitted in any form or by any means, electronic, mechanical, photocopying, recording, scanning or otherwise, except under the terms of the Copyright, Designs and Patents Act 1988 or under the terms of a licence issued by the Copyright Licensing Agency Ltd., Saffron House, 6-10 Kirby Street, London EC1N 8TS, UK, without the permission in writing of the Publisher. Requests to the Publisher for permission should be addressed to the Permissions Department, John Wiley & Sons, Ltd, The Atrium, Southern Gate, Chichester, West Sussex, PO19 8SQ, England, or emailed to permreq@wiley.co.uk, or faxed to (44) 1243 770620.

Trademarks: Wiley, the Wiley logo, For Dummies, the Dummies Man logo, A Reference for the Rest of Us!, The Dummies Way, Dummies Daily, The Fun and Easy Way, Dummies.com, Making Everything Easier, and related trade dress are trademarks or registered trademarks of John Wiley & Sons, Ltd. and/or its affiliates in the United States and other countries, and may not be used without written permission. All other trademarks are the property of their respective owners. John Wiley & Sons, Ltd. is not associated with any product or vendor mentioned in this book.

LIMIT OF LIABILITY/DISCLAIMER OF WARRANTY: THE PUBLISHER, THE AUTHOR, AND ANYONE ELSE IN PREPARING THIS WORK MAKE NO REPRESENTATIONS OR WARRANTIES WITH RESPECT TO THE ACCURACY OR COMPLETENESS OF THE CONTENTS OF THIS WORK AND SPECIFICALLY DISCLAIM ALL WARRANTIES, INCLUDING WITHOUT LIMITATION WARRANTIES OF FITNESS FOR A PARTICULAR PURPOSE. NO WARRANTY MAY BE CREATED OR EXTENDED BY SALES OR PROMOTIONAL MATERIALS. THE ADVICE AND STRATEGIES CONTAINED HEREIN MAY NOT BE SUITABLE FOR EVERY SITUATION. THIS WORK IS SOLD WITH THE UNDERSTANDING THAT THE PUBLISHER IS NOT ENGAGED IN RENDERING LEGAL, ACCOUNTING, OR OTHER PROFESSIONAL SERVICES. IF PROFESSIONAL ASSISTANCE IS REQUIRED, THE SERVICES OF A COMPETENT PROFESSIONAL PERSON SHOULD BE SOUGHT. NEITHER THE PUBLISHER NOR THE AUTHOR SHALL BE LIABLE FOR DAMAGES ARISING HEREFROM. THE FACT THAT AN ORGANIZATION OR WEBSITE IS REFERRED TO IN THIS WORK AS A CITATION AND/OR A POTENTIAL SOURCE OF FURTHER INFORMATION DOES NOT MEAN THAT THE AUTHOR OR THE PUBLISHER ENDORSES THE INFORMATION THE ORGANIZATION OR WEBSITE MAY PROVIDE OR RECOMMENDATIONS IT MAY MAKE. FURTHER, READERS SHOULD BE AWARE THAT INTERNET WEBSITES LISTED IN THIS WORK MAY HAVE CHANGED OR DISAPPEARED BETWEEN WHEN THIS WORK WAS WRITTEN AND WHEN IT IS READ.

For general information on our other products and services, please contact our Customer Care Department within the U.S. at 877-762-2974, outside the U.S. at 317-572-3993, or fax 317-572-4002.

For technical support, please visit www.wiley.com/techsupport.

Wiley also publishes its books in a variety of electronic formats and by print-on-demand. Some content that appears in standard print versions of this book may not be available in other formats. For more information about Wiley products, visit us at www.wiley.com.

British Library Cataloguing in Publication Data: A catalogue record for this book is available from the British Library.

ISBN 978-1-119-96284-7 (paperback); ISBN 978-1-119-96312-7 (ebook); ISBN 978-1-119-96313-4 (ebook); ISBN 978-1-119-96314-1 (ebook)

Printed and bound in the United States by Bind-Rite

10 9 8 7 6 5 4 3 2 1



WILEY

About the Authors

Andrie de Vries: Andrie started to use R in 2009 to analyze survey data, and he has been continually impressed by the ability of the open-source community to innovate and create phenomenal software. Andrie is director of PentaLibra Limited, a boutique market-research firm specializing in surveys and statistical analysis. He has contributed two R packages to CRAN and is developing several packages to make the analysis and reporting of survey data easier. He also is actively involved in the development of LimeSurvey, the open-source survey-management system. To maintain equilibrium in his life, Andrie is working toward a yoga teacher diploma at the Krishnamacharya Healing and Yoga Foundation.

Joris Meys, MSc: Joris is a statistical consultant and R programmer in the Department of Mathematical Modeling, Statistics, and Bioinformatics at Ghent University (Belgium). After earning a master's degree in biology, he worked for six years in environmental research and management before starting an advanced master's degree in statistical data analysis. Joris writes packages for both specific projects and general implementation of methods developed in his department, and he is the maintainer of several packages on R-Forge. He has co-authored a number of scientific papers as a statistical expert. To balance science with culture, Joris spends most of his spare time playing saxophone in a couple of local bands.

Dedication

This book is for my wife, Annemarie, because of her encouragement, support, and patience. And for my 9-year-old niece, Tanya, who is really good at math and kept reminding me that the deadline for this book was approaching!

—Andrie de Vries

For my mother, because she made me the man I am. And for Granny, because she rocks!

—Joris Meys

Authors' Acknowledgments

This book is possible only because of the tremendous support we had from our editorial team at Wiley. In particular, thank you to Elizabeth Kuball for her patient and detailed editing and gentle cajoling, and to Sara Shlaer for pretending not to hear the sound of missed deadlines. Thank you to Kathy Simpson for teaching us how to write in Dummies style, and to Chris Katsaropoulos for getting us started.

Thank you to our technical editor, Gavin Simpson, for his thorough reading and many helpful comments.

Thank you to the R core team for creating R, for maintaining CRAN, and for your dedication to the R community in the form of mailing lists, documentation, and seminars. And thank you to the R community for creating thousands of useful packages, writing blogs, and answering questions.

In this book, we use several packages by Hadley Wickham, whose contribution of ggplot graphics and other helpful packages like plyr continues to be an inspiration.

While writing this book we had very helpful support from a large number of contributors to the R tag at Stack Overflow. Thank you to James (JD) Long, David Winsemius, Ben Bolker, Joshua Ulrich, Kohske Takahashi, Barry Rowlingson, Roman Luštrik, David Purdy, Nick Sabbe, Joran Elias, Brandon Bertelsen, Michael Sumner, Ian Fellows, Dirk Eddelbuettel, Simon Urbanek, Gabor Grotendieck, and everybody else who continue to make Stack Overflow a great resource for the R community.

From Andrie: Thank you to the large number of people who have influenced my choices. Professor Bruce Hardie at London Business School made a throwaway comment during a lecture in 2004 that first made me aware of R; his pragmatic approach to making statistics relevant to marketing decision makers was inspirational. Gary Bennett at Logit Research has been very helpful with his advice and a joy to work with. Paul Hutchings at Kindle Research was very supportive in the early stages of setting up my business.

From Joris: Thank you to the professors and my colleagues at the Department of Mathematical Modeling, Statistics, and Bioinformatics at Ghent University (Belgium) for the insightful discussions and ongoing support I received while writing this book.

Publisher's Acknowledgments

We're proud of this book; please send us your comments at <http://dummies.custhelp.com>. For other comments, please contact our Customer Care Department within the U.S. at 877-762-2974, outside the U.S. at 317-572-3993, or fax 317-572-4002.

Some of the people who helped bring this book to market include the following:

Acquisitions, Editorial

Project Editor: Elizabeth Kuball
Executive Commissioning Editor: Birgit Gruber
Assistant Editor: Ellie Scott
Copy Editor: Elizabeth Kuball
Technical Editor: Gavin Simpson
Editorial Manager: Jodi Jensen
Senior Project Editor: Sara Shlaer
Editorial Assistant: Leslie Saxman
Special Editorial Help: Kathy Simpson
Cover Photo: ©iStockphoto.com/Jacobo Cortés
Cartoons: Rich Tennant
(www.the5thwave.com)

Composition Services

Sr. Senior Project Coordinator: Kristie Rees
Layout and Graphics: Sennett Vaughan
Johnson, Lavonne Roberts
Proofreaders: Lindsay Amones, Susan Hobbs
Indexer: Sharon Shock

Marketing

Associate Marketing Director: Louise Breinholt
Senior Marketing Executive: Kate Parrett

UK Tech Publishing

Michelle Leete, VP Consumer and Technology Publishing Director
Martin Tribe, Associate Director–Book Content Management
Chris Webb, Associate Publisher

Publishing and Editorial for Technology Dummies

Richard Swadley, Vice President and Executive Group Publisher
Andy Cummings, Vice President and Publisher
Mary Bednarek, Executive Acquisitions Director
Mary C. Corder, Editorial Director

Publishing for Consumer Dummies

Kathleen Nebenhaus, Vice President and Executive Publisher

Composition Services

Debbie Stailey, Director of Composition Services

Contents at a Glance



<i>Introduction</i>	<i>1</i>
<i>Part I: R You Ready?</i>	<i>7</i>
Chapter 1: Introducing R: The Big Picture.....	9
Chapter 2: Exploring R	15
Chapter 3: The Fundamentals of R	31
<i>Part II: Getting Down to Work in R</i>	<i>43</i>
Chapter 4: Getting Started with Arithmetic.....	45
Chapter 5: Getting Started with Reading and Writing	71
Chapter 6: Going on a Date with R.....	93
Chapter 7: Working in More Dimensions	103
<i>Part III: Coding in R</i>	<i>137</i>
Chapter 8: Putting the Fun in Functions	139
Chapter 9: Controlling the Logical Flow	159
Chapter 10: Debugging Your Code	179
Chapter 11: Getting Help.....	193
<i>Part IV: Making the Data Talk</i>	<i>203</i>
Chapter 12: Getting Data into and out of R.....	205
Chapter 13: Manipulating and Processing Data.....	219
Chapter 14: Summarizing Data.....	253
Chapter 15: Testing Differences and Relations	275
<i>Part V: Working with Graphics</i>	<i>299</i>
Chapter 16: Using Base Graphics.....	301
Chapter 17: Creating Faceted Graphics with Lattice.....	317
Chapter 18: Looking At ggplot2 Graphics.....	333

<i>Part VI: The Part of Tens</i>	<i>347</i>
Chapter 19: Ten Things You Can Do in R That You Would've Done in Microsoft Excel.....	349
Chapter 20: Ten Tips on Working with Packages	359
<i>Appendix: Installing R and RStudio.....</i>	<i>365</i>
<i>Index</i>	<i>371</i>

Table of Contents

***Introduction* 1**

About This Book	1
Conventions Used in This Book	2
What You're Not to Read	3
Foolish Assumptions	4
How This Book Is Organized	4
Part I: R You Ready?	4
Part II: Getting Down to Work in R	5
Part III: Coding in R	5
Part IV: Making the Data Talk	5
Part V: Working with Graphics	5
Part VI: The Part of Tens	6
Icons Used in This Book	6
Where to Go from Here	6

***Part 1: R You Ready?* 7**

Chapter 1: Introducing R: The Big Picture 9

Recognizing the Benefits of Using R	10
It comes as free, open-source code	10
It runs anywhere	11
It supports extensions	11
It provides an engaged community	11
It connects with other languages	12
Looking At Some of the Unique Features of R	12
Performing multiple calculations with vectors	12
Processing more than just statistics	13
Running code without a compiler	14

Chapter 2: Exploring R 15

Working with a Code Editor	16
Exploring RGui	16
Dressing up with RStudio	19
Starting Your First R Session	22
Saying hello to the world	22
Doing simple math	22
Using vectors	23
Storing and calculating values	23
Talking back to the user	25
Sourcing a Script	25



Navigating the Workspace.....	28
Manipulating the content of the workspace.....	28
Saving your work	28
Retrieving your work.....	29

Chapter 3: The Fundamentals of R 31

Using the Full Power of Functions.....	31
Vectorizing your functions	32
Putting the argument in a function.....	33
Making history.....	34
Keeping Your Code Readable	35
Following naming conventions	35
Structuring your code	38
Adding comments.....	40
Getting from Base R to More.....	40
Finding packages.....	40
Installing packages.....	41
Loading and unloading packages.....	41

***Part II: Getting Down to Work in R*..... 43**

Chapter 4: Getting Started with Arithmetic 45

Working with Numbers, Infinity, and Missing Values	45
Doing basic arithmetic	46
Using mathematical functions.....	48
Calculating whole vectors	51
To infinity and beyond	52
Organizing Data in Vectors.....	54
Discovering the properties of vectors	54
Creating vectors.....	56
Combining vectors.....	57
Repeating vectors	58
Getting Values in and out of Vectors	58
Understanding indexing in R.....	59
Extracting values from a vector	59
Changing values in a vector.....	60
Working with Logical Vectors.....	61
Comparing values	62
Using logical vectors as indices.....	63
Combining logical statements	64
Summarizing logical vectors	65
Powering Up Your Math with Vector Functions.....	66
Using arithmetic vector operations.....	66
Recycling arguments	69

Chapter 5: Getting Started with Reading and Writing 71

Using Character Vectors for Text Data	71
Assigning a value to a character vector	72
Creating a character vector with more than one element	72
Extracting a subset of a vector	73
Naming the values in your vectors	74
Manipulating Text	76
String theory: Combining and splitting strings	76
Sorting text	79
Finding text inside text	80
Substituting text	83
Revvig up with regular expressions	84
Factoring in Factors	86
Creating a factor	86
Converting a factor	87
Looking at levels	89
Distinguishing data types	90
Working with ordered factors	91

Chapter 6: Going on a Date with R 93

Working with Dates	93
Presenting Dates in Different Formats	95
Adding Time Information to Dates	97
Formatting Dates and Times	98
Performing Operations on Dates and Times	99
Addition and subtraction	99
Comparison of dates	100
Extraction	101

Chapter 7: Working in More Dimensions. 103

Adding a Second Dimension	103
Discovering a new dimension	103
Combining vectors into a matrix	106
Using the Indices	107
Extracting values from a matrix	108
Replacing values in a matrix	110
Naming Matrix Rows and Columns	111
Changing the row and column names	111
Using names as indices	112
Calculating with Matrices	113
Using standard operations with matrices	113
Calculating row and column summaries	114
Doing matrix arithmetic	115
Adding More Dimensions	117
Creating an array	117
Using dimensions to extract values	118

Combining Different Types of Values in a Data Frame	119
Creating a data frame from a matrix	119
Creating a data frame from scratch.....	121
Naming variables and observations	122
Manipulating Values in a Data Frame.....	123
Extracting variables, observations, and values	124
Adding observations to a data frame	125
Adding variables to a data frame.....	127
Combining Different Objects in a List	129
Creating a list.....	129
Extracting elements from lists	131
Changing the elements in lists	132
Reading the output of str() for lists.....	134
Seeing the forest through the trees.....	135

Part III: Coding in R..... 137

Chapter 8: Putting the Fun in Functions139

Moving from Scripts to Functions	139
Making the script	140
Transforming the script.....	140
Using the function.....	141
Reducing the number of lines	143
Using Arguments the Smart Way.....	145
Adding more arguments	145
Conjuring tricks with dots	147
Using functions as arguments	148
Coping with Scoping.....	150
Crossing the borders.....	150
Using internal functions	152
Dispatching to a Method	154
Finding the methods behind the function	154
Doing it yourself.....	156

Chapter 9: Controlling the Logical Flow.....159

Making Choices with if Statements	160
Doing Something Else with an if...else Statement.....	162
Vectorizing Choices	163
Looking at the problem.....	164
Choosing based on a logical vector.....	164
Making Multiple Choices	166
Chaining if...else statements.....	166
Switching between possibilities.....	167
Looping Through Values	168
Constructing a for loop	169
Calculating values in a for loop.....	169

Looping without Loops: Meeting the Apply Family	171
Looking at the family features	172
Meeting three of the members	173
Applying functions on rows and columns	173
Applying functions to listlike objects	175
Chapter 10: Debugging Your Code	179
Knowing What to Look For	179
Reading Errors and Warnings	180
Reading error messages	180
Caring about warnings (or not)	181
Going Bug Hunting	182
Calculating the logit	182
Knowing where an error comes from	183
Looking inside a function	184
Generating Your Own Messages	187
Creating errors	187
Creating warnings	188
Recognizing the Mistakes You're Sure to Make	189
Starting with the wrong data	189
Having your data in the wrong format	190
Chapter 11: Getting Help	193
Finding Information in the R Help Files	193
When you know exactly what you're looking for	193
When you don't know exactly what you're looking for	194
Searching the Web for Help with R	196
Getting Involved in the R Community	197
Using the R mailing lists	197
Discussing R on Stack Overflow and Stack Exchange	198
Tweeting about R	199
Making a Minimal Reproducible Example	199
Creating sample data with random values	199
Producing minimal code	201
Providing the necessary information	201
 Part IV: Making the Data Talk	 203
Chapter 12: Getting Data into and out of R	205
Getting Data into R	205
Entering data in the R text editor	205
Using the Clipboard to copy and paste	207
Reading data in CSV files	209
Reading data from Excel	211
Working with other data types	212
Getting Your Data out of R	214

Working with Files and Folders	215
Understanding the working directory.....	215
Manipulating files.....	216

Chapter 13: Manipulating and Processing Data219

Deciding on the Most Appropriate Data Structure	219
Creating Subsets of Your Data	221
Understanding the three subset operators	221
Understanding the five ways of specifying the subset.....	221
Subsetting data frames.....	222
Adding Calculated Fields to Data	227
Doing arithmetic on columns of a data frame.....	227
Using with and within to improve code readability	227
Creating subgroups or bins of data.....	228
Combining and Merging Data Sets	230
Creating sample data to illustrate merging	231
Using the merge() function.....	232
Working with lookup tables.....	234
Sorting and Ordering Data.....	236
Sorting vectors	237
Sorting data frames.....	237
Traversing Your Data with the Apply Functions	240
Using the apply() function to summarize arrays	241
Using lapply() and sapply() to traverse a list or data frame.....	242
Using tapply() to create tabular summaries	243
Getting to Know the Formula Interface.....	245
Whipping Your Data into Shape	246
Understanding data in long and wide format.....	247
Getting started with the reshape2 package.....	248
Melting data to long format	249
Casting data to wide format	250

Chapter 14: Summarizing Data253

Starting with the Right Data	253
Using factors or numeric data.....	254
Counting unique values.....	254
Preparing the data	255
Describing Continuous Variables	256
Talking about the center of your data.....	256
Describing the variation.....	257
Checking the quantiles.....	257
Describing Categories	258
Counting appearances.....	259
Calculating proportions	259
Finding the center	260
Describing Distributions.....	261
Plotting histograms	261
Using frequencies or densities.....	262

Describing Multiple Variables	264
Summarizing a complete dataset	265
Plotting quantiles for subgroups	266
Tracking correlations	268
Working with Tables	270
Creating a two-way table	271
Converting tables to a data frame	272
Looking at margins and proportions	273

Chapter 15: Testing Differences and Relations275

Taking a Closer Look at Distributions	276
Observing beavers	276
Testing normality graphically	276
Using quantile plots	277
Testing normality in a formal way	280
Comparing Two Samples	281
Testing differences	281
Comparing paired data	283
Testing Counts and Proportions	284
Checking out proportions	284
Analyzing tables	286
Extracting test results	287
Working with Models	288
Analyzing variances	288
Evaluating the differences	290
Modeling linear relations	292
Evaluating linear models	295
Predicting new values	297

Part V: Working with Graphics 299

Chapter 16: Using Base Graphics301

Creating Different Types of Plots	301
Getting an overview of plot	301
Adding points and lines to a plot	302
Different plot types	306
Controlling Plot Options and Arguments	308
Adding titles and axis labels	308
Changing plot options	309
Putting multiple plots on a single page	312
Saving Graphics to Image Files	314

Chapter 17: Creating Faceted Graphics with Lattice317

Creating a Lattice Plot	318
Loading the lattice package	319
Making a lattice scatterplot	319
Adding trend lines	320

Changing Plot Options	321
Adding titles and labels.....	322
Changing the font size of titles and labels.....	322
Using themes to modify plot options	324
Plotting Different Types.....	325
Making a bar chart.....	325
Making a box-and-whisker plot	326
Plotting Data in Groups.....	327
Using data in tall format.....	327
Creating a chart with groups.....	329
Adding a key	330
Printing and Saving a Lattice Plot.....	331
Assigning a lattice plot to an object	331
Printing a lattice plot in a script	331
Saving a lattice plot to file.....	332

Chapter 18: Looking At ggplot2 Graphics. 333

Installing and Loading ggplot2.....	333
Looking At Layers.....	334
Using Geoms and Stats	335
Defining what data to use	336
Mapping data to plot aesthetics	336
Getting geoms.....	337
Sussing Stats.....	340
Adding Facets, Scales, and Options	343
Adding facets.....	344
Changing options	345
Getting More Information	346

Part VI: The Part of Tens 347

Chapter 19: Ten Things You Can Do in R That You Would've Done in Microsoft Excel. 349

Adding Row and Column Totals	349
Formatting Numbers	350
Sorting Data.....	352
Making Choices with If.....	352
Calculating Conditional Totals.....	353
Transposing Columns or Rows	353
Finding Unique or Duplicated Values	354
Working with Lookup Tables	355
Working with Pivot Tables	355
Using the Goal Seek and Solver	356

Chapter 20: Ten Tips on Working with Packages 359

Poking Around the Nooks and Crannies of CRAN	359
Finding Interesting Packages	360

Installing Packages	360
Loading Packages	361
Reading the Package Manual and Vignette	361
Updating Packages	362
Unloading Packages	362
Forging Ahead with R-Forge	363
Conducting Installations from BioConductor	364
Reading the R Manual	364
 <i>Appendix: Installing R and RStudio</i>	 <i>365</i>
Installing and Configuring R	365
Installing R	365
Configuring R	366
Installing and Configuring RStudio	368
Installing RStudio	368
Configuring RStudio	368
 <i>Index</i>	 <i>371</i>

Introduction

Welcome to *R For Dummies*, the book that takes the steepness out of the learning curve for using R.

We can't guarantee that you'll be a guru if you read this book, but you should be able to do the following:

- ✓ Perform data analysis by using a variety of powerful tools
- ✓ Use the power of R to do statistical analysis and other data-processing tasks
- ✓ Appreciate the beauty of using vector-based operations rather than loops to do speedy calculations
- ✓ Appreciate the meaning of the following line of code:

```
knowledge <- apply(theory, 1, sum)
```
- ✓ Know how to find, download, and use code that has been contributed to R by its very active community of developers
- ✓ Know where to find extra help and resources to take your R coding skills to the next level
- ✓ Create beautiful graphs and visualizations of your data

About This Book

R For Dummies is an introduction to the statistical programming language known as R. We start by introducing the interface and work our way from the very basic concepts of the language through more sophisticated data manipulation and analysis.

We illustrate every step with easy-to-follow examples. This book contains numerous code snippets, several write-it-yourself functions you can use later on, and complete analysis scripts. All these are for you to try out yourself.

We don't attempt to give a technical description of how R is programmed internally, but we do focus as much on the why as on the how. R doesn't function as your average scripting language, and it has plenty of unique features that may seem surprising at first. Instead of just telling you how you have to talk to R, we believe it's important for us to explain how the R engine reads what you tell it to do. After reading this book, you should be able to manipulate your data in the form you want and understand how to use functions we *didn't* cover in the book (as well as the ones we do cover).

R and RStudio

R For Dummies can be used with any operating system that R runs on. Whether you use Mac, Linux, or Windows, this book will get you on your way with R.

R is more a programming language than an application. When you download R, you automatically download a console application that's suitable for your operating system. However, this application has only basic functionality, and it differs to some extent from one operating system to the next.

RStudio is a cross-platform application with some very neat features to support R. In this book, we don't assume you use any specific console application. But because RStudio provides a common user interface across the major operating systems, we think that you'll understand how to run it quite quickly. For this reason, we use RStudio to demonstrate some of the concepts rather than operating-system-specific editor.

This book is a reference. You don't have to read it from beginning to end. Instead, you can use the table of contents and index to find the information you need. In each chapter, we cross-reference other chapters where you can find more information.

Conventions Used in This Book

Code snippets appear like this example, where we simulate 1 million throws of two six-sided dice:

```
> set.seed(42)
> throws <- 1e6
> dice <- sapply(1:2,
+   function(x)sample(1:6, throws, replace=TRUE)
+ )
> table(rowSums(dice))
```

	2	3	4	5	6	7	8
28007	55443	83382	110359	138801	167130	138808	
	9	10	11	12			
110920	83389	55816	27945				

Each line of R code in this example is preceded by one of two symbols:

- ✓ **>:** The prompt symbol, **>**, is not part of your code, and you should not type this when you try the code yourself.
- ✓ **+:** The continuation symbol, **+**, indicates that this line of code still belongs to the previous line of code. In fact, you don't have to break a line of code into two, but we do this frequently, because it improves the readability of code and helps it fit into the pages of a book.

The lines that don't start with either the prompt symbol or the continuation symbol are the output produced by R. In this case, you get the total number of throws where the dice added up to the numbers 2 through 12. For example, out of 1 million throws of the dice, on 28,007 occasions, the numbers on the dice added to 2.

You can copy these code snippets and run them in R, but you have to type them exactly as shown. There are only three exceptions:

- ✓ Don't type the prompt symbol, `>`.
- ✓ Don't type the continuation symbol, `+`.
- ✓ Where you put spaces or tabs isn't critical, as long as it isn't in the middle of a keyword. Pay attention to new lines, though.

You get to write some R code in every chapter of the book. Because much of your interaction with the R interpreter is most likely going to be in interactive mode, you need a way of distinguishing between your code and the results of your code. When there is an instruction to type some text into the R console, you'll see a little `>` symbol to the left of the text, like this:

```
> print("Hello world!")
```

If you type this into a console and press Enter, R responds with the following:

```
[1] "Hello world!"
```

For convenience, we collapse these two events into a single block, like this:

```
> print("Hello world!")  
[1] "Hello world!"
```

This indicates that you type some code (`print("Hello world!")`) into the console and R responds with `[1] "Hello world!"`.

Finally, many R words are directly derived from English words. To avoid confusion in the text of this book, R functions, arguments, and keywords appear in `monofont`. For example, to create a plot, you use the `plot()` function in R. When talking about functions, the function name will always be followed by open and closed parentheses — for example, `plot()`. We refrain from adding arguments to the function names mentioned in the text, unless it's really important.

On some occasions we talk about menu commands, such as `File→Save`. This just means that you open the File menu and choose the Save option.

What You're Not to Read

You can use this book however works best for you, but if you're pressed for time (or just not interested in the nitty-gritty details), you can safely skip

anything marked with a Technical Stuff icon. You also can skip sidebars (text in gray boxes); they contain interesting information, but nothing critical to your understanding of the subject at hand.

Foolish Assumptions

This book makes the following assumptions about you and your computer:

- ✓ **You know your way around a computer.** You know how to download and install software. You know how to find information on the Internet and you have Internet access.
- ✓ **You're not necessarily a programmer.** If you are a programmer, and you're used to coding in other languages, you may want to read the notes marked by the Technical Stuff icon — there, we fill you in on how R is similar to, or different from, other common languages.
- ✓ **You're not a statistician, but you understand the very basics of statistics.** *R For Dummies* isn't a statistics book, although we do show you how to do some basic statistics using R. If you want to understand the statistical stuff in more depth, we recommend *Statistics For Dummies*, 2nd Edition, by Deborah J. Rumsey, PhD (Wiley).
- ✓ **You want to explore new stuff.** You like to solve problems and aren't afraid of trying things out in the R console.

How This Book Is Organized

The book is organized in six parts. Here's what each of the six parts covers.

Part I: R You Ready?

In this part, we introduce you to R and show you how to write your first script. You get to use the very powerful concept of vectors to make simultaneous calculations on many variables at once. You get to work with the R workspace (in other words, how to create, modify, or remove variables). You find out how save your work and retrieve and modify script files that you wrote in previous sessions. We also introduce some fundamentals of R (for example, how to extend functionality by installing packages).

Part II: Getting Down to Work in R

In this part, we fill you in on the three R's: reading, 'riting, and 'rithmetic — in other words, working with text and number (and dates for good measure).

You also get to use the very important data structures of lists and data frames.

Part III: Coding in R

R is a programming language, so you need to know how to write and understand functions. In this part, we show you how to do this, as well as how to control the logic flow of your scripts by making choices using `if` statements, as well as looping through your code to perform repetitive actions. We explain how to make sense of and deal with warnings and errors that you may experience in your code. Finally, we show you some tools to debug any issues that you may experience.

Part IV: Making the Data Talk

In this part, we introduce the different data structures that you can use in R, such as lists and data frames. You find out how to get your data in and out of R (for example, by reading data from files or the Clipboard). You also see how to interact with other applications, such as Microsoft Excel.

Then you discover how easy it is to do some advanced data reshaping and manipulation in R. We show you how to select a subset of your data and how to sort and order it. We explain how to merge different datasets based on columns they may have in common. Finally, we show you a very powerful generic strategy of splitting and combining data and applying functions over subsets of your data. When you understand this strategy, you can use it over and over again to do sophisticated data analyses in only a few small steps.

We're just itching to show you how to do some statistical analysis using R. This is the heritage of R, after all. But we promise to keep it simple. After reading this part, you'll know how to describe and summarize your variables and data using R. You'll be able to do some classical tests (for example, calculating a t-test). And you'll know how to use random numbers to simulate some distributions.

Finally, we show you some of the basics of using linear models (for example, linear regression and analysis of variance). We also show you how to use R to predict the values of new data using some models that you've fitted to your data.

Part V: Working with Graphics

They say that a picture is worth a thousand words. This is certainly the case when you want to share your results with other people. In this part, you discover how to create basic and more sophisticated plots to visualize your

data. We move on from bar charts and line charts, and show you how to present cuts of your data using facets.

Part VI: The Part of Tens

In this part, we show you how to do ten things in R that you probably use Microsoft Excel for at the moment (for example, how to do the equivalent of pivot tables and lookup tables). We also give you ten tips for working with packages that are not part of base R.

Icons Used in This Book



As you read this book, you'll find little pictures in the margins. These pictures, or *icons*, mark certain types of text:

When you see the Tip icon, you can be sure to find a way to do something more easily or quickly.



You don't have to memorize this book, but the Remember icon points out some useful things that you really should remember. Usually this indicates a design pattern or idiom that you'll encounter in more than one chapter.



When you see the Warning icon, listen up. It points out something you definitely don't want to do. Although it's really unlikely that using R will cause something disastrous to happen, we use the Warning icon to alert you if something is bound to lead to confusion.



The Technical Stuff icon indicates technical information you can merrily skip over. We do our best to make this information as interesting and relevant as possible, but if you're short on time or you just want the information you absolutely *need* to know, you can move on by.

Where to Go from Here

There's only one way to learn R: Use it! In this book, we try to make you familiar with the usage of R, but you'll have to sit down at your PC and start playing around with it yourself. Crack the book open so the pages don't flip by themselves, and start hitting the keyboard!

Part I

R You Ready?

The 5th Wave

By Rich Tennant



"Okay, ma'am, I'm going to ask you to walk a straight line, then I'm going to ask you to bisect that line with a perpendicular line that slopes to the equation $y = 3x + 5$."

In this part . . .

From financial headquarters to the dark cellars of small universities, people use R for data manipulation and statistical analysis. With R, you can extract stock prices and predict profits, discover beginning diseases in small blood samples, analyze the behavior of customers, or describe how the gray wolf recolonized European forests.

In this part, you discover the power hidden behind the 18th letter of the alphabet.

Chapter 1

Introducing R: The Big Picture

In This Chapter

- ▶ Discovering the benefits of R
 - ▶ Identifying some programming concepts that make R special
-

With an estimated worldwide user base of more than 2 million people, the R language has rapidly grown and extended since its origin as an academic demonstration language in the 1990s.

Some people would argue — and we think they're right — that R is much more than a statistical programming language. It's also:

- ✓ A very powerful tool for all kinds of data processing and manipulation
- ✓ A community of programmers, users, academics, and practitioners
- ✓ A tool that makes all kinds of publication-quality graphics and data visualizations
- ✓ A collection of freely distributed add-on packages
- ✓ A toolbox with tremendous versatility

In this chapter, we fill you in on the benefits of R, as well as its unique features and quirks.



You can download R at www.r-project.org. This website also provides more information on R and links to the online manuals, mailing lists, conferences and publications.

Tracing the history of R

Ross Ihaka and Robert Gentleman developed R as a free software environment for their teaching classes when they were colleagues at the University of Auckland in New Zealand. Because they were both familiar with S, a commercial programming language for statistics, it seemed natural to use similar syntax in their own work. After Ihaka and Gentleman announced their software on the S-news mailing list, several people became interested and started to collaborate with them, notably Martin Mächler.

Currently, a group of 18 people has rights to modify the central archive of source code. This group is referred to as the R Development Core Team. In addition, many other people have contributed new code and bug fixes to the project.

Here are some milestone dates in the development of R:

- ✓ **Early 1990s:** The development of R began.
- ✓ **August 1993:** The software was announced on the S-news mailing list. Since then, a set of active R mailing lists has been created.

The web page at www.r-project.org/mail.html provides descriptions of these lists and instructions for subscribing. (For more information, turn to “It provides an engaged community,” later in this chapter.)

- ✓ **June 1995:** After some persuasive arguments by Martin Mächler (among others) to make the code available as “free software,” the code was made available under the Free Software Foundation’s GNU General Public License (GPL), Version 2.
- ✓ **Mid-1997:** The initial R Development Core Team was formed (although, at the time, it was simply known as the core group).
- ✓ **February 2000:** The first version of R, version 1.0.0, was released.

Ross Ihaka wrote a comprehensive overview of the development of R. The web page <http://cran.r-project.org/doc/html/interface98-paper/paper.html> provides a fascinating history.

Recognizing the Benefits of Using R

Of the many attractive benefits of R, a few stand out: It’s actively maintained, it has good connectivity to various types of data and other systems, and it’s versatile enough to solve problems in many domains. Possibly best of all, it’s available for free, in more than one sense of the word.

It comes as free, open-source code

R is available under an open-source license, which means that anyone can download and modify the code. This freedom is often referred to as “free as in speech.” R is also available free of charge — a second kind of freedom, sometimes referred to as “free as in beer.” In practical terms, this means that you can download and use R free of charge.