A Brain-Friendly Guide

Head First Data Analysis

Predict your raise with linear regression

Experiment to discover who your customers *really* are

Load important statistical concepts directly into your brain

A learner's guide to big numbers, statistics, and good decisions

Sell more toys by optimizing your business model

Overcome your cognitive biases

Clean messy data for efficient analysis

O'REILLY®

Michael Milton

Head First Data Analysis

Information Theory/Data Analysis

What will you learn from this book?

There's a whole world of data out there, and it's your job to make sense of it all. Where to begin? *Head First Data Analysis* helps you organize your data in Excel or OpenOffice, take it further with R, find meaningful patterns with scatterplots and histograms, draw conclusions using heuristics, predict the future by experimenting and testing hypotheses, and display findings with clear visualizations.

Whether you're a product developer researching the viability of a new product, a marketing manager gauging the effectiveness of a campaign, a salesperson presenting data to clients, or a lone entrepreneur responsible for all these data-intensive functions and more, *Head First Data Analysis* is a complete learning experience for making data the most useful tool in your business.

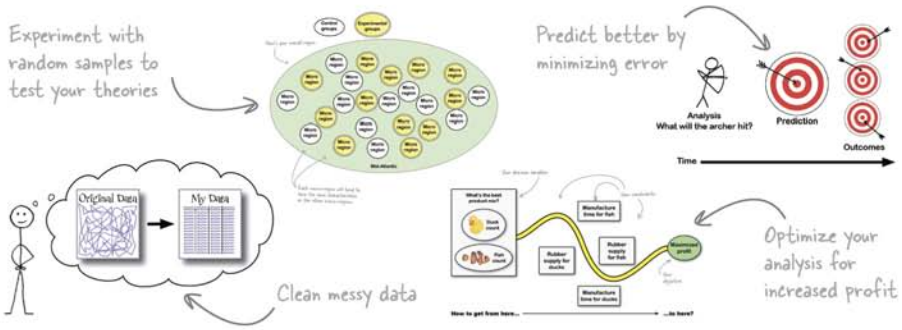

What's so special about this book?

We think your time is too valuable to waste struggling with new concepts. Using the latest research in cognitive science and learning theory to craft a multi-sensory learning experience, *Head First Data Analysis* uses a visually rich format designed for the way your brain works, not a text-heavy approach that puts you to sleep.

"It's about time a straightforward and comprehensive guide to analyzing data was written that makes learning the concepts simple and fun. Concepts are good in theory and even better in practicality."

— Anthony Rose, President, Support Analytics

"*Head First Data Analysis* shows how to find and unlock the power of data in everyday life and how systematic data analysis can improve decision making."

— Eric Heilman, Statistics teacher, Georgetown Preparatory School

"Buried under mountains of data? Fill your toolbox with the analytical skills that give you an edge and turn raw numbers into real knowledge."

— Bill Mietelski, Software engineer

US $49.99 CAN $62.99

ISBN: 978-0-596-15393-9

9 780596 153939 54999

O'REILLY®

oreilly.com
headfirstlabs.com

Safari Books Online

Free online edition for 45 days with purchase of this book. Details on last page.

Advance Praise for *Head First Data Analysis*

"It's about time a straightforward and comprehensive guide to analyzing data was written that makes learning the concepts simple and fun. It will change the way you think and approach problems using proven techniques and free tools. Concepts are good in theory and even better in practicality."

— Anthony Rose, President, Support Analytics

"*Head First Data Analysis* does a fantastic job of giving readers systematic methods to analyze real-world problems. From coffee, to rubber duckies, to asking for a raise, *Head First Data Analysis* shows the reader how to find and unlock the power of data in everyday life. Using everything from graphs and visual aides to computer programs like Excel and R, *Head First Data Analysis* gives readers at all levels accessible ways to understand how systematic data analysis can improve decision making both large and small."

— Eric Heilman, Statistics teacher, Georgetown Preparatory School

"Buried under mountains of data? Let Michael Milton be your guide as you fill your toolbox with the analytical skills that give you an edge. In *Head First Data Analysis*, you'll learn how to turn raw numbers into real knowledge. Put away your Ouija board and tarot cards; all you need to make good decisions is some software and a copy of this book."

— Bill Mietelski, Software engineer

Praise for other *Head First* books

"Kathy and Bert's *Head First Java* transforms the printed page into the closest thing to a GUI you've ever seen. In a wry, hip manner, the authors make learning Java an engaging 'what're they gonna do next?' experience."

—Warren Keuffel, Software Development Magazine

"Beyond the engaging style that drags you forward from know-nothing into exalted Java warrior status, *Head First Java* covers a huge amount of practical matters that other texts leave as the dreaded "exercise for the reader..." It's clever, wry, hip and practical—there aren't a lot of textbooks that can make that claim and live up to it while also teaching you about object serialization and network launch protocols."

—Dr. Dan Russell, Director of User Sciences and Experience Research IBM Almaden Research Center (and teacher of Artificial Intelligence at Stanford University)

"It's fast, irreverent, fun, and engaging. Be careful—you might actually learn something!"

—Ken Arnold, former Senior Engineer at Sun Microsystems Coauthor (with James Gosling, creator of Java), *The Java Programming Language*

"I feel like a thousand pounds of books have just been lifted off of my head."

—Ward Cunningham, inventor of the Wiki and founder of the Hillside Group

"Just the right tone for the geeked-out, casual-cool guru coder in all of us. The right reference for practical development strategies—gets my brain going without having to slog through a bunch of tired stale professor-speak."

—Travis Kalanick, Founder of Scour and Red Swoosh Member of the MIT TR100

"There are books you buy, books you keep, books you keep on your desk, and thanks to O'Reilly and the *Head First* crew, there is the ultimate category, *Head First* books. They're the ones that are dog-eared, mangled, and carried everywhere. *Head First SQL* is at the top of my stack. Heck, even the PDF I have for review is tattered and torn."

— Bill Sawyer, ATG Curriculum Manager, Oracle

"This book's admirable clarity, humor and substantial doses of clever make it the sort of book that helps even non-programmers think well about problem-solving."

— Cory Doctorow, co-editor of BoingBoing Author, *Down and Out in the Magic Kingdom* and *Someone Comes to Town, Someone Leaves Town*

"I received the book yesterday and started to read it...and I couldn't stop. This is definitely très 'cool.' It is fun, but they cover a lot of ground and they are right to the point. I'm really impressed."

— Erich Gamma, IBM Distinguished Engineer, and co-author of *Design Patterns*

"One of the funniest and smartest books on software design I've ever read."

— Aaron LaBerge, VP Technology, ESPN.com

"What used to be a long trial and error learning process has now been reduced neatly into an engaging paperback."

— Mike Davidson, CEO, Newsvine, Inc.

"Elegant design is at the core of every chapter here, each concept conveyed with equal doses of pragmatism and wit."

— Ken Goldstein, Executive Vice President, Disney Online

"I ♥ *Head First HTML with CSS & XHTML*—it teaches you everything you need to learn in a 'fun coated' format."

— Sally Applin, UI Designer and Artist

"Usually when reading through a book or article on design patterns, I'd have to occasionally stick myself in the eye with something just to make sure I was paying attention. Not with this book. Odd as it may sound, this book makes learning about design patterns fun.

"While other books on design patterns are saying 'Buehler… Buehler… Buehler…' this book is on the float belting out 'Shake it up, baby!'"

— Eric Wuehler

"I literally love this book. In fact, I kissed this book in front of my wife."

— Satish Kumar

Other related books from O'Reilly

- Analyzing Business Data with Excel
- Excel Scientific and Engineering Cookbook
- Access Data Analysis Cookbook

Other books in O'Reilly's *Head First* series

- Head First Java
- Head First Object-Oriented Analysis and Design (OOA&D)
- Head First HTML with CSS and XHTML
- Head First Design Patterns
- Head First Servlets and JSP
- Head First EJB
- Head First PMP
- Head First SQL
- Head First Software Development
- Head First JavaScript
- Head First Ajax
- Head First Physics
- Head First Statistics
- Head First Rails
- Head First PHP & MySQL
- Head First Algebra
- Head First Web Design
- Head First Networking

Head First Data Analysis

Michael Milton

O'REILLY®

Beijing • Cambridge • Farnham • Köln • Sebastopol • Taipei • Tokyo

Head First Data Analysis

by Michael Milton

Printed in the United States of America.

Published by O'Reilly Media, Inc., 1005 Gravenstein Highway North, Sebastopol, CA 95472.

O'Reilly Media books may be purchased for educational, business, or sales promotional use. Online editions are also available for most titles (*safari.oreilly.com*). For more information, contact our corporate/institutional sales department: (800) 998-9938 or *corporate@oreilly.com*.

Series Creators:	Kathy Sierra, Bert Bates
Series Editor:	Brett D. McLaughlin
Editor:	Brian Sawyer
Cover Designers:	Karen Montgomery
Production Editor:	Scott DeLugan
Proofreader:	Nancy Reinhardt
Indexer:	Jay Harward
Page Viewers:	Mandarin, the fam, and Preston

Printing History:

July 2009: First Edition.

No data was harmed in the making of this book.

RepKover™ This book uses RepKover,™ a durable and flexible lay-flat binding.

ISBN: 978-0-596-15393-9

[M]

Dedicated to the memory of my grandmother, Jane Reese Gibbs.

Author of Head First Data Analysis

Michael Milton

Michael Milton has spent most of his career helping nonprofit organizations improve their fundraising by interpreting and acting on the data they collect from their donors.

He has a degree in philosophy from New College of Florida and one in religious ethics from Yale University. He found reading *Head First* to be a revelation after spending years reading *boring* books filled with terribly important stuff and is grateful to have the opportunity to write an *exciting* book filled with terribly important stuff.

When he's not in the library or the bookstore, you can find him running, taking pictures, and brewing beer.

Table of Contents (Summary)

Table of Contents (the real thing)

Intro

Your brain on data analysis. Here *you* are trying to *learn* something, while here your *brain* is doing you a favor by making sure the learning doesn't *stick*. Your brain's thinking, "Better leave room for more important things, like which wild animals to avoid and whether naked snowboarding is a bad idea." So how *do* you trick your brain into thinking that your life depends on knowing data analysis?

introduction to data analysis

1 Break it down

Data is everywhere.

Nowadays, everyone has to deal with mounds of data, whether they call themselves "data analysts" or not. But people who possess a toolbox of data analysis skills have a **massive edge** on everyone else, because they understand what to ***do*** with all that stuff. They know how to translate raw numbers into intelligence that **drives real-world action**. They know how to **break down and structure** complex problems and data sets to get right to the heart of the problems in their business.

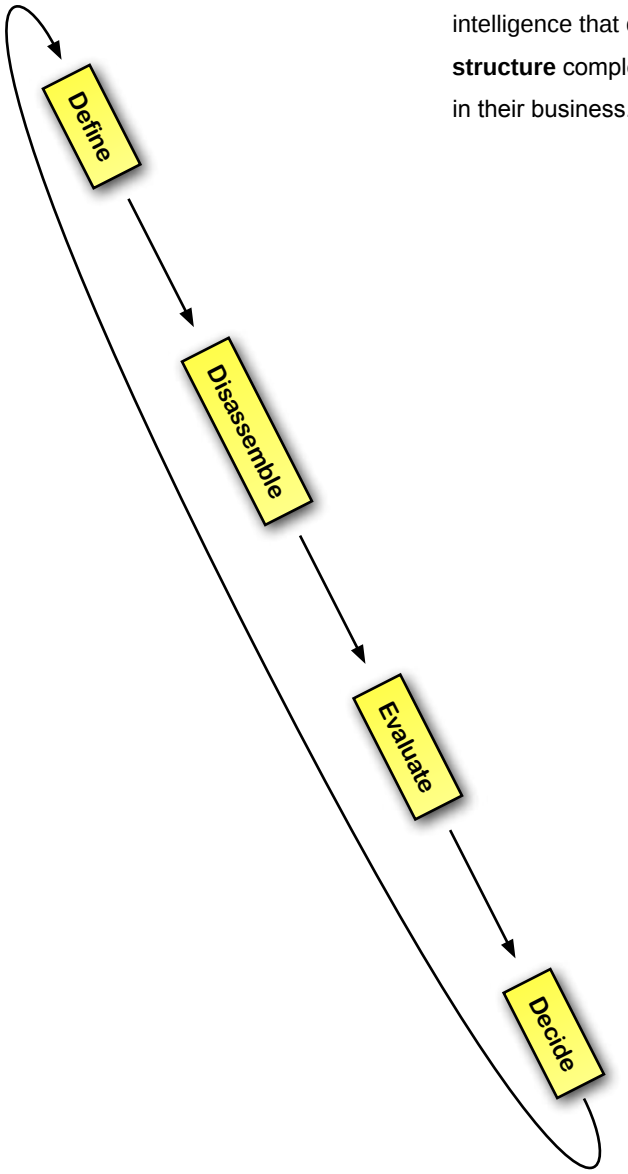

2 experiments
Test your theories

Can you show what you believe?

In a real **empirical** test? There's nothing like a good experiment to solve your problems and show you the way the world really works. Instead of having to rely exclusively on your **observational data**, a well-executed experiment can often help you make **causal connections**. Strong empirical data will make your analytical judgments all the more powerful.

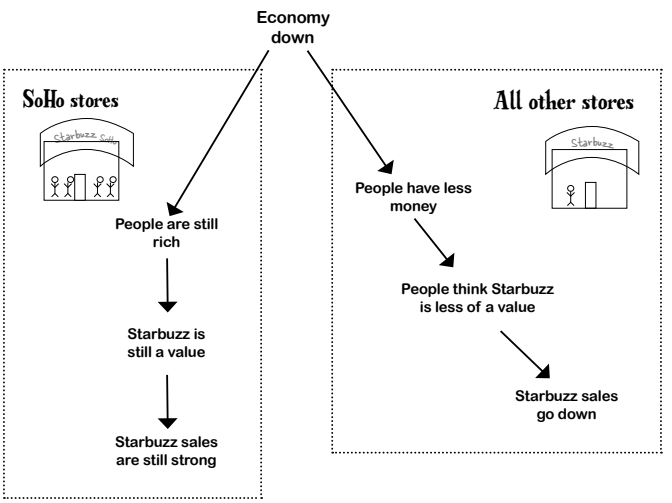

3 optimization

Take it to the max

We all want more of something.

And we're always trying to figure out how to get it. *If* the things we want more of—profit, money, efficiency, speed—can be represented numerically, then chances are, there's an tool of data analysis to help us tweak our *decision variables,* which will help us find the **solution** or *optimal point* where we get the most of what we want. In this chapter, you'll be using one of those tools and the powerful spreadsheet **Solver** package that implements it.

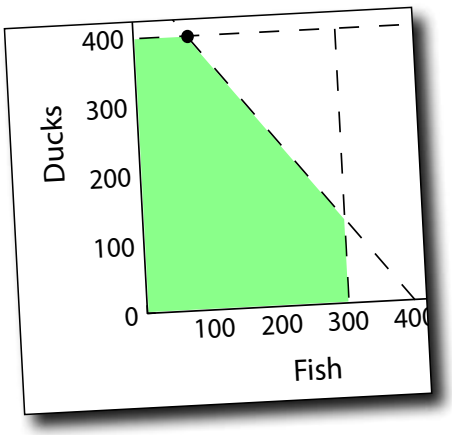

4 data visualization

Pictures make you smarter

You need more than a table of numbers.

Your data is brilliantly complex, with more variables than you can shake a stick at. Mulling over mounds and mounds of spreadsheets isn't just boring; it can actually be a waste of your time. A clear, highly multivariate visualization can, in a small space, show you the forest that you'd miss for the trees if you were just looking at spreadsheets all the time.

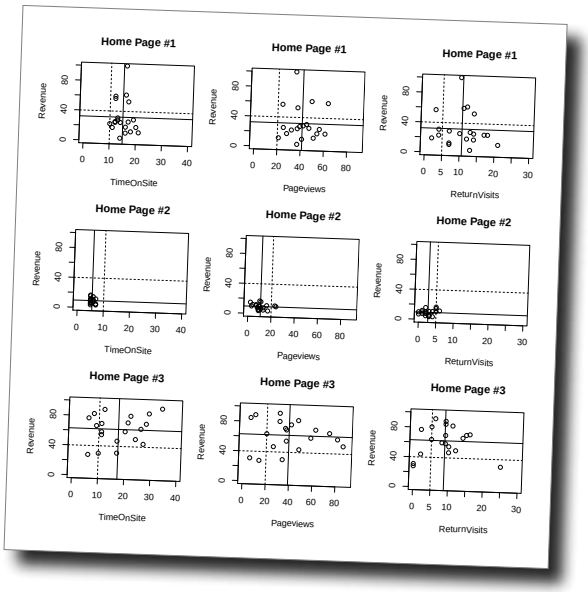

5 hypothesis testing

Say it ain't so

The world can be tricky to explain.

And it can be fiendishly difficult when you have to deal with complex, heterogeneous data to anticipate future events. This is why analysts don't just take the obvious explanations and assume them to be true: the careful reasoning of data analysis enables you to meticulously evaluate a bunch of options so that you can incorporate all the information you have into your models. You're about to learn about **falsification**, an unintuitive but powerful way to do just that.

6 bayesian statistics

Get past first base

You'll always be collecting new data.

And you need to make sure that every analysis you do incorporates the data you have that's relevant to your problem. You've learned how *falsification* can be used to deal with heterogeneous data sources, but what about **straight up probabilities**? The answer involves an extremely handy analytic tool called **Bayes' rule**, which will help you incorporate your **base rates** to uncover not-so-obvious insights with ever-changing data.

7 subjective probabilities

Numerical belief

Sometimes, it's a good idea to make up numbers.

Seriously. But only if those numbers describe your own mental states, expressing your beliefs. **Subjective probability** is a straightforward way of injecting some real *rigor* into your hunches, and you're about to see how. Along the way, you are going to learn how to evaluate the spread of data using **standard deviation** and enjoy a special guest appearance from one of the more powerful analytic tools you've learned.

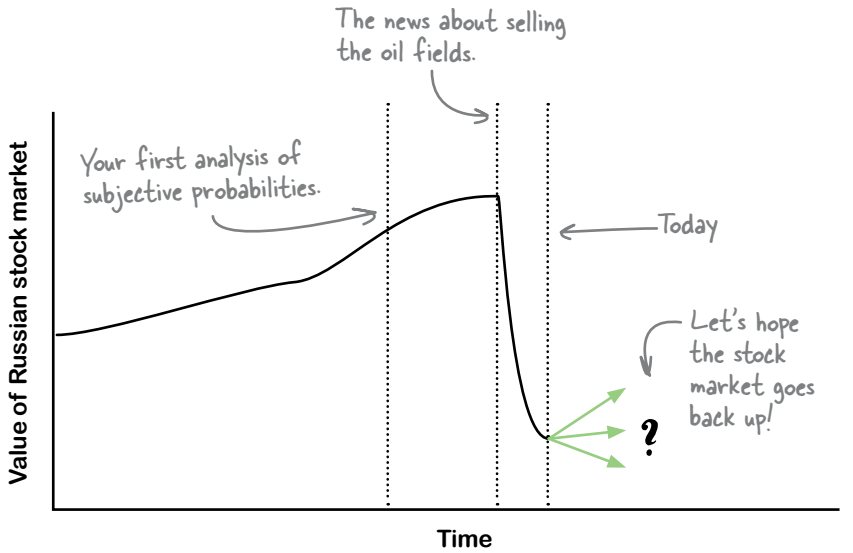

8 heuristics

Analyze like a human

The real world has more variables than you can handle.

There is always going to be data that you can't have. And even when you do have data on most of the things you want to understand, *optimizing* methods are often **elusive** and **time consuming**. Fortunately, most of the actual thinking you do in life is not "rational maximizing"—it's processing incomplete and uncertain information with rules of thumb so that you can make decisions quickly. What is really cool is that these rules can **actually work** and are important (and necessary) tools for data analysts.

9 histograms

The shape of numbers

How much can a bar graph tell you?

There are about a zillion ways of **showing data with pictures**, but one of them is special. **Histograms**, which are kind of similar to bar graphs, are a super-fast and easy way to summarize data. You're about to use these powerful little charts to measure your data's **spread, variability, central tendency**, and more. No matter how large your data set is, if you draw a histogram with it, you'll be able to "see" what's happening inside of it. And you're about to do it with a new, free, crazy-powerful **software tool**.

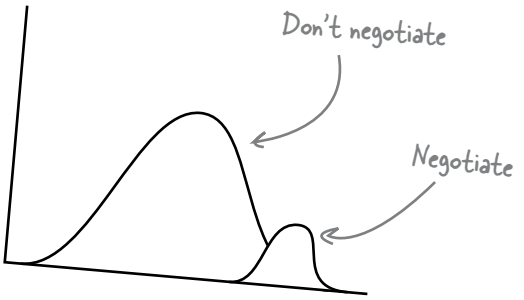

10 regression Prediction

Predict it.

Regression is an incredibly powerful statistical tool that, when used correctly, has the ability to help you predict certain values. When used with a controlled experiment, regression can actually help you predict the future. Businesses use it like crazy to help them build models to explain customer behavior. You're about to see that the judicious use of regression can be very profitable indeed.

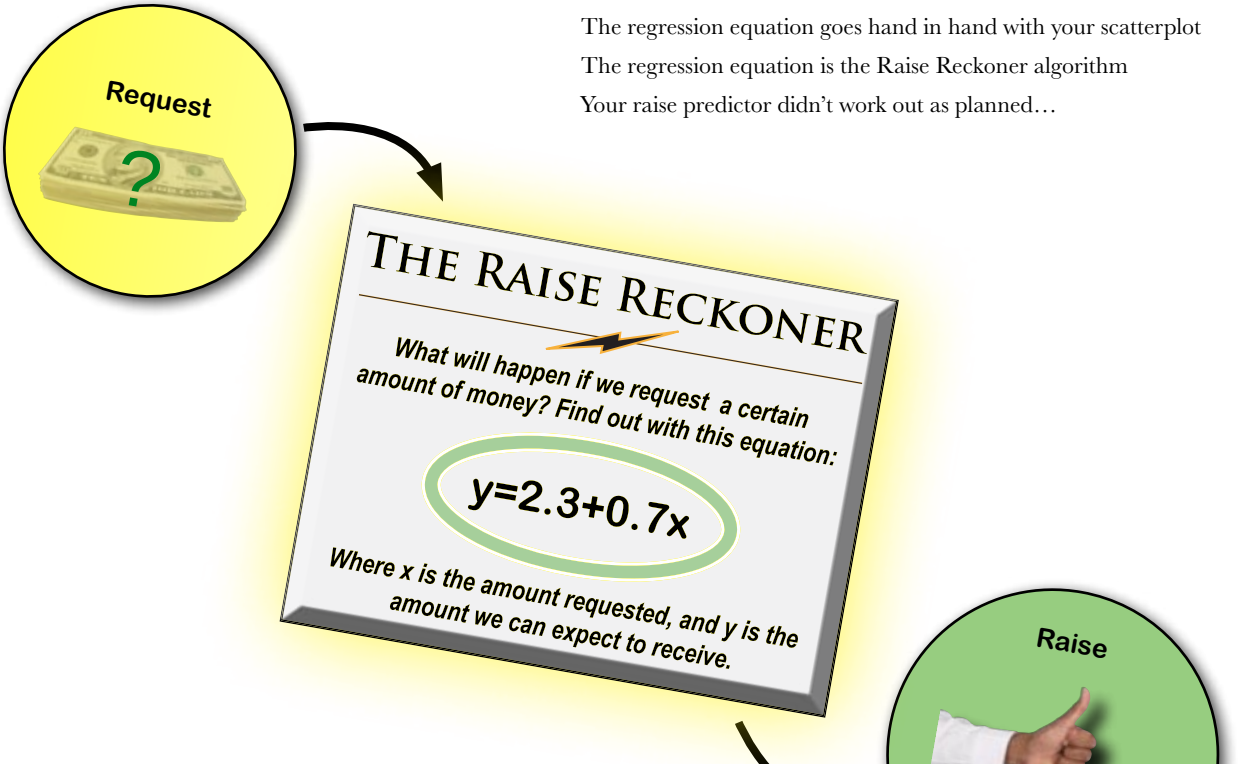

error

11 Err well

The world is messy.

So it should be no surprise that your predictions rarely hit the target squarely. But if you offer a prediction with an **error range**, you and your clients will know not only the average predicted value, but also how far you expect typical deviations from that error to be. Every time you express error, you offer a much richer perspective on your predictions and beliefs. And with the tools in this chapter, you'll also learn about how to get error under control, getting it as low as possible to increase confidence.

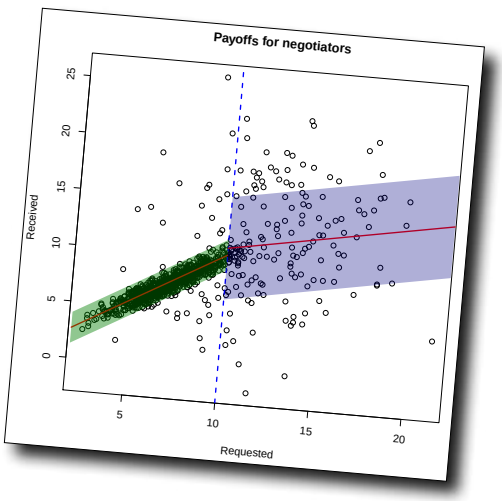

12 relational databases

Can you relate?

How do you structure really, really multivariate data?

A spreadsheet has only *two dimensions*: rows and columns. And if you have a bunch of dimensions of data, the **tabular format** gets old really quickly. In this chapter, you're about to see firsthand where spreadsheets make it really hard to manage multivariate data and learn **how relational database management systems** make it easy to store and retrieve countless permutations of multivariate data.

13 cleaning data

Impose order

Your data is useless…

…if it has messy structure. And a lot of people who ***collect*** data do a crummy job of maintaining a neat structure. If your data's not neat, you can't slice it or dice it, run formulas on it, or even really *see* it. You might as well just ignore it completely, right? Actually, you can do better. With a **clear vision** of how you need it to look and a few **text manipulation tools**, you can take the funkiest, craziest mess of data and **whip** it into something useful.

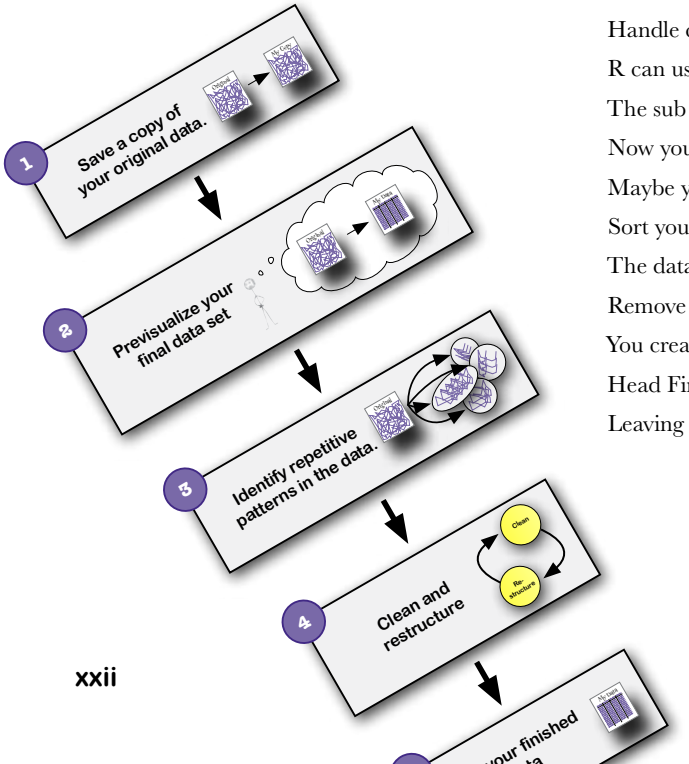

leftovers

The Top Ten Things (we didn't cover)

You've come a long way.

But data analysis is a vast and constantly evolving field, and there's so much left the learn. In this appendix, we'll go over ten items that there wasn't enough room to cover in this book but should be high on your list of topics to learn about next.

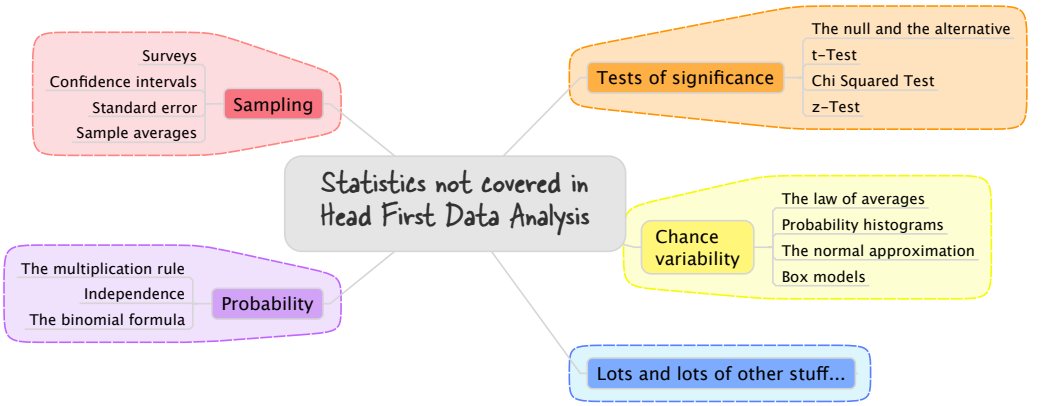

install r

Start R up!

Behind all that data-crunching power is enormous complexity.

But fortunately, getting R installed and ***started*** is something you can accomplish in just a few minutes, and this appendix is about to show you how to pull off your R install without a hitch.

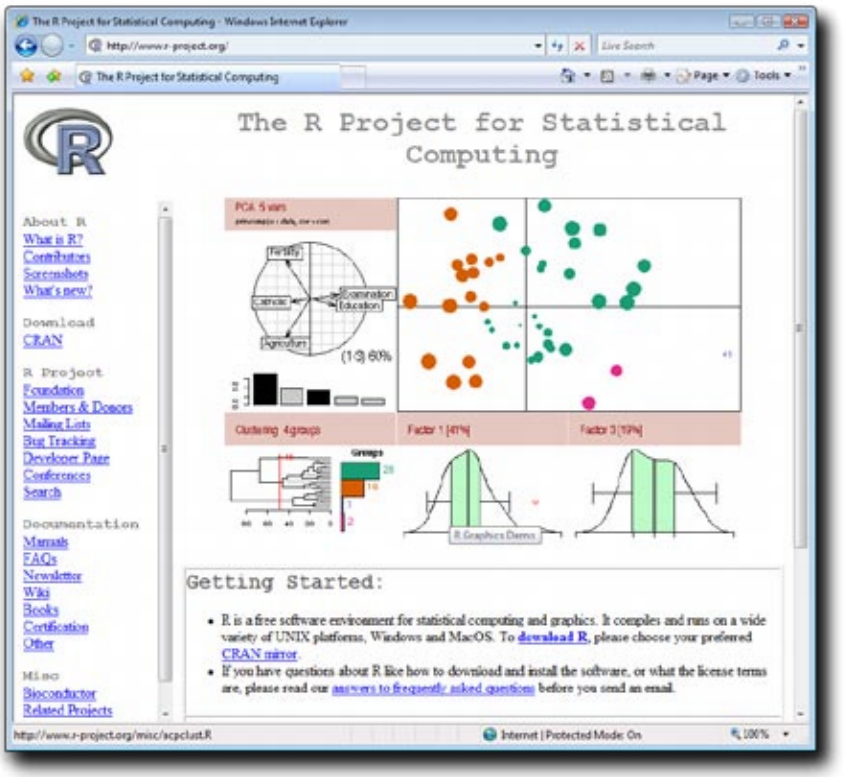

install excel analysis tools

The ToolPak

Some of the best features of Excel aren't installed by default.

That's right, in order to run the optimization from Chapter 3 and the histograms from Chapter 9, you need to activate the **Solver** and the **Analysis ToolPak**, two extensions that are included in Excel by default but not activated without your initiative.

Install the data analysis tools in Excel 432