

Math for Life



Ian Morrison

© 1998-2010 by Ian Morrison
All rights reserved.

<http://projectivepress.com/math4life/math4life.pdf>

Ian Morrison
Department of Mathematics
Fordham University
New York, New York

Draft of September 12, 2010.

The cover illustration shows E. M. Ward's painting of Exchange Alley, London during the **South Sea Bubble** of 1720. Among the many who were singed in this famous speculative blaze was Isaac Newton, who, after losing £20,000, bemoaned, "I can calculate the movement of the stars, but not the madness of men."

projective press, llc

Preface

This is an working draft of a set of lecture notes for a one-semester terminal course in mathematics aimed at intending humanities, business and social science majors. My goal is that even students with very weak mathematical preparations learn that mathematics *can* illuminate important questions in their everyday lives and that all but the weakest learn something about *how* it does so.

My first priority was finding questions that can capture the interest of such students. My second was giving complete explanations of the answers that combine rigor and accessibility. Together these have dictated a preference for depth over breadth. All the mathematics I develop deals with either finite probability and statistics or the time value of money, but these topics are treated thoroughly. Where my priorities were not mutually compatible, I have sacrificed completeness by employing mathematical black boxes but these are clearly identified when used.

This is a draft of September 12, 2010. The most recent draft of the hypertext version of the notes may be downloaded at:

<http://projectivepress.com/math4life/math4life.pdf>.

Comments, corrections and other suggestions for improving these notes are welcome. Please email them to me at morrison@fordham.edu.

IAN MORRISON
New York, NY



Contents

Preface	i
Contents	iii
Help! Navigating through MATH⁴LIFE	ix
Overview	ix
Using Acrobat or your browser	x
Choosing a view	xi
Using the MATH⁴LIFE navigation buttons	xii
Using links in MATH⁴LIFE documents	xiv
How MATH⁴LIFE is numbered	xvi
What MATH⁴LIFE 's type styles and colors mean	xvii
Using graphs and pictures	xix
Introduction	1
Goals of the course	1
Three questions that affect your life	2
How to learn mathematics effectively	6
1 Back to Basics	11
1.1 Order of operations	11
1.2 Handling numbers in various formats	25
Working with big numbers	25
Fractions versus decimals	28



	Rounding	30
1.3	Sums and series	35
	Summation notation	35
	Geometric summation and series formulae	38
1.4	Logarithms and exponentials	42
	Don't look down!	42
	The natural logarithm	50
	exp and exponentials	70
2	Dangerous misunderstandings	83
2.1	A stroll through the minefield	84
	“\$7'll get you \$12”	85
	Trolling for terrorists	87
	Lightning strikes twice	90
	AIG gives back: a fairy tale with a moral	93
	“We wuz robbed”	95
	He's on <i>Fire!</i>	101
	Chuck-a-luck	107
2.2	You want prions in that burger	122
3	When it really counts	136
3.1	Sets: the air we breathe	137
	Objects and oracles	138
	The order of a set	144
3.2	Subsets	145
3.3	Product sets	149
	Sequences	149
	Products of sets	156
3.4	Power sets	162
	Infinities and an argument from <i>The Book</i>	165
	Combinations	171
3.5	Lists and permutations	179
	Lists: sets with an order	179

	Lists, ordered subsets and sequences without repetition	181
	Permutations: counting lists	184
	Counting orderings	188
3.6	Shorthands for basic counts	193
	The two rules of guessing: Don't and "D'oh"	193
	Shorthand Questions	194
	The two question method	195
3.7	And, Or, and Not: the three hardest words	217
	andthen: extensive and	218
	andalso: restrictive and	220
	eitherorboth and orelse	226
	butnot	235
	The Four Quadrants and the Three Hard Words	238
	When conjunctions collide	244
3.8	Counting by the "divide and conquer" method	246
	The method	247
	First Examples and the Most Common Patterns	250
	More Challenging Examples	264
	Examples with Permutations and Powers	269
	A Classic Example: Poker Rankings	279
	The m&m's problem	283
4	Fat chance	298
4.1	Experiments, outcomes and sample spaces	299
4.2	Probability measures	305
	Probabilities of Outcomes and Events	305
	Basic Probability Formulae	308
4.3	Equally Likely Outcomes	315
	Basic Equally Likely Outcomes Formulae	316
	Setting up Equally Likely Outcomes Sample Spaces	324
	First steps in the minefield	334
4.4	Conditional Probability	337
	Basic Conditional Probabilities	337

	Distinguishing Conditional Probabilities	345
	Probability Measures for Compound Experiments	352
4.5	Organizing Related Probabilities	359
	Working with Probabilities in Tables	359
	Working with Probabilities in Tree Diagrams	374
	Analysing the game of Craps	384
4.6	Applications of Bayesian probability	392
	Filtering spam	395
	“\$7’ll get you \$12” Revisited	398
	The Hardy-Weinberg Principle	404
4.7	The dice don’t talk to each other	411
	Independence	411
	Applying independence: the binomial distribution	420
	BSE Testing	432
	Common misconceptions about independence	437
4.8	Who would have expected	440
	Random variables and expected values	441
	Relations between expected values	459
	Simpson’s Paradox	470
	Spreads: variance and standard deviation	475
4.9	Everything is back to normal	485
5	Time is money	511
5.1	Simple interest	511
5.2	Compound interest	527
5.3	Approximating compound interest	548
	The simple interest approximation	548
	The continuous approximation	551
5.4	Yields	560
	Annualized yields	563
	Continuous yields	568
	The rules of 69.3 and 72	574
5.5	Applications of compounding	581

Inflation 581

Population 591

Computation 598

Radiation 604

5.6 Amortization: saving 613

 The future amortization formula 614

 Working with the future amortization formula 620

5.7 Planning for future needs 627

5.8 Amortization: borrowing 640

 The Present amortization formula 641

 Working with the present amortization formula 643

5.9 Things everyone should know 650

 Equity 650

 “Money talks, nobody walks” 663

5.10 The hunt-and-peck method 672

 A lengthy example 672

 Linear Interpolation 682

 When good guess go bad 687

 The method 690

 More About Linear Interpolation 697

 Finding yields using the hunt-and-peck method 699

 Points 705

 Final comments 714

5.11 Saving to spend 715

 Fixed term annuities 716

 Life annuities 719

Index **725**

Help! Navigating through MATH⁴LIFE

Overview

This document will teach you how to navigate the **MATH⁴LIFE** website in various ways: via **Acrobat™** or your personal favorite pdf reader, or in your web browser, using a set of standard navigation buttons of all the pages, and using the network of internal links between the various section documents in the course. In addition, you'll learn how the course numbering system works and what the different type styles and colors mean so you can use these to help navigate as well. Links to the subsections which describe each are listed below.

Help! Table of Contents

OVERVIEW

USING **Acrobat** OR YOUR BROWSER

CHOOSING A VIEW

USING THE **MATH⁴LIFE** NAVIGATION BUTTONS

USING LINKS IN **MATH⁴LIFE** DOCUMENTS

HOW **MATH⁴LIFE** IS NUMBERED

WHAT **MATH⁴LIFE**'S TYPE STYLES AND COLORS MEAN

USING GRAPHS AND PICTURES



Using Acrobat or your browser

I'm going to assume that you are familiar with the operation of the **Acrobat** application which lets you view and work with PDF or Portable Document Format files. You've clearly managed to get hold of a copy of the application or the plug-in or you could not be reading this file. (If you need a copy for another computer that you use, you can get the latest version free from the [Acrobat Reader](#) web site.) If you have never worked with PDF files before, you might want to spend a few minutes looking over the manual which is also included in the standard **Acrobat** installation as a file named `Reader.pdf`. By the way, you can generally recognize PDF files by their names which virtually always end in the suffix ".pdf".

The **MATH⁴LIFE** web site is designed to be accessed interactively through a standard web browser: it works well with almost all versions (from this millennium) of standard browsers like Explorer, Firefox and Safari. However, the HTML format used by even the latest browsers doesn't allow you to view much mathematics: even something as simple as a fraction like $\frac{1}{x}$ is currently beyond it. (In the future a proposed extension of HTML called XML may solve this but there are currently no good tools for creating XML pages and no browsers which support displaying complex XML mathematics.) So, in **MATH⁴LIFE** this problem is solved by placing the math inside PDF files instead, and then using **Acrobat** to display this mathematical text. This approach means that to use the site you will need to configure your browser to use **Acrobat** or the PDFViewer plugin which is distributed with the Reader to display the PDF files in the course.

Exactly what you need to do to configure your browser varies considerably depending on what browser you use and what platform (Mac, Windows ...) you are working on. Fortunately, very detailed instructions for each of the common browser/platform combinations are available at the [Acrobat Reader Support Center](#) page. Once you



are correctly configured, you should be able to navigate through the **MATH⁴LIFE** site just like any other except that when you link to a PDF file you'll see it in an the **Acrobat** window. If you find that following links to other files causes your browser to begin downloading the file you requested instead of just displaying it, you are *not* correctly configured. Check that you have followed *all* the instructions for configuring your browser properly. If the problem persists, ask your instructor for help.

Choosing a view

You're probably used to using your browser to view HTML pages which are delivered at a fixed size. If the page you are viewing fits in the window you are using, there's not much point in enlarging the window. You'll just see more space. Conversely, if the current page is larger than your monitor, you can only see part of the page.

Not in **MATH⁴LIFE**. **Acrobat** can display a page from a PDF file at just about any reasonable size. In fancier terms, PDF images are scalable. You may have noticed that when you opened this file what you saw was the entire first page sized to just fit in your current window. This is how any document here will appear when you first open it.

What happens if you resize the window you are using? The page also resizes automatically! Now you'll see it at the largest size that fits into the *resized* window. If the window became bigger, so does the page and all the type on it. You can get the largest type and make the page easiest to read simply by maximizing the window it's in. The proportions of the pages on the site have been chosen so that when you do this the page will nearly fill the whole monitor both vertically and horizontally. (The actual aspect ratio—horizontal to vertical ratio—used is 2:3. The dimensions of pages have been chosen so that they'll fit nicely to the format of a standard trade book (which is 6 inches by 9 inches) and so that they'll be easy to read on


Using the MATH⁴LIFE navigation buttons

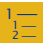
a standard early 21st century monitor. If you have a larger monitor, you can just magnify the pages (all the type scales and they'll be easier to read) or use a smaller window as you prefer.

The **Acrobat** reader also provides several other tools for adjusting the view which are described in detail in the manual. Its View menu lets you alter the standard view with which each new page is displayed and its magnifying glass tool lets you zoom in to see part of the page in more detail (and back out again). This last is particularly handy for looking closely at graphs and other pictures as described in [USING GRAPHS AND PICTURES](#).

Using the MATH⁴LIFE navigation buttons

All the pages in the **MATH⁴LIFE** website contain a standard set of buttons—located at the bottom left of each page—to assist you in navigating through the site and, ultimately, in learning the material in the course more rapidly and thoroughly. Most pages also contain more specific links: these are explained in the next subsection of this document. If you want to return from either kind of link to the page you came from, just press the Back button in your browser as usual. The navigation buttons below work just like those in the title bar if you want to experiment.

 This is the **HELP** button. Pressing it brings you to this document. Use it anytime you want to refresh your memory about how the site is organized and how to navigate it.

 The **TABLE OF CONTENTS** button takes you to a document which lists all the section documents in **MATH⁴LIFE** and various related files. The order in which the sections are listed is the one I recommend studying them in—the chapters and sections are also numbered in the this order. However, when you are working in later sections, you may find you want to go back to earlier ones to review

a topic which is giving you trouble. If the section you are working on does not have a link to the one you'd like to review, you can use this button to find the document you are after. Also, the **TABLE OF CONTENTS** is only *my* suggestion for how to present MATH⁴LIFE so do not be concerned if your instructor has omitted some topics and present others in a different order.



The **INDEX** button takes you to an alphabetical index to the important elements in the course. These include definitions, notations, formulas, methods and topics mentioned in examples and exercises. If you want to refresh your memory about any element and you don't see a link to it on the page you are reading, you can use this button to go to the index, find what you're interested in.

You can then jump to the spots in the course where it is discussed, because every page number in the index file is a link to the corresponding page. The course sections also contain many links to the index. Text **this color** or this **lighter shade** indicates that an important term is being used. Just clicking on the the term itself takes you to directly either to its definition or to the corresponding entry in the index. For more help on the using the links to and from the index, see [USING LINKS IN MATH⁴LIFE DOCUMENTS](#).



The **MAIL** button is located at the bottom right of each page rather than at the top. When you press it, it opens your web browser's email client to a blank letter addressed to me. If you have anything you'd like me to hear—things you like or dislike about MATH⁴LIFE, suggestions for improvements, questions the course didn't answer to your satisfaction or discussions you found confusing—please write to me. Your feedback will help make MATH⁴LIFE better for other students. I am especially interested if you think you have found a mathematical or factual error in the text. I have tried to root them all out but your help with any you think you have found will be greatly appreciated.



Using links in MATH⁴LIFE documents

Whenever one section refers to any other part of the course—like the elements found in the index but also to specific examples, problems and projects—it provides a link to the corresponding location. If you want to take another look at what is being referred to, just follow the link. This ability to keep yourself mentally up-to-date in any discussion is one of the most important innovations in MATH⁴LIFE. *Use it!* : math is a very sequential subject and it's usually impossible to understand a new concept well if you don't understand other concepts on which it depends. In the MATH⁴LIFE website, I've tried to make sure that you are never more than a click away from those other concepts. Please take advantage of all the work that went into making this possible.

Links are colored, not underlined. Four colors are commonly used. I'll first discuss the different kinds and then give some examples of each.

[This color](#) is used for links within a single document. Following these links is pretty much instantaneous.

[This color](#) is used for links for a *different* document in the MATH⁴LIFE site. If you have already opened this document in the current session, you'll jump to it right away. However, if you have not worked with the file to which you are linking, your browser will have to contact the MATH⁴LIFE website and download the file for you to work with. This may cause a small delay—usually no more than a couple of seconds—in bringing up the material you linked to. The lighter color warns you that there may be a short delay in making the link.

Let me make one other point about these cross-document links. In most cases, links to other files take you to a section you've already covered in the course. That's what I mean when I say math is sequential. If you find yourself linking back to a previous section often, you'd be wise to take a break and review it separately. In the long

run, you'll find that it's much faster to get the prior material straight once than to have to keep going back and reviewing it every time it's needed later on. If you don't, you'll soon find yourself jumping back 2 or 3 times in a row. Very inefficient.

Links to and from the **INDEX** use [this color](#) or this [lighter shade](#). The [deeper shade](#) is used to flag the primary occurrence of an important element—the definition of a new term or the statement of a formula. So this color marks something that we'll be referring to frequently later on and that you should make careful note of. Clicking on the term will take you directly to its entry in the index.

The [lighter shade](#) is used to indicate secondary appearances of an important term. In such cases, clicking on the term will take you, not to the index, but directly to its primary occurrence so you can see what it means immediately. The [lighter shade](#) is also used to mark terms that are indexed but that are not defined in the course—things like proper names, historical events and so on. You can recognize that you're looking at an entry with no definition or primary occurrence because when you click on it, you'll be taken to its index entry.

In the opposite direction, every [page number](#) in the **INDEX** links you to the corresponding page in the course. Bold type and arrowheads are used to indicate what you can expect to find at the end of each link. Let look at a hypothetical example (not actually in the index) dealing with the quadratic formula.

quadratic formula, [22](#)➤, [26](#), [31](#), [35](#), [44](#)

The page containing the primary or defining use of this term—in this case, it would be the statement of the quadratic formula—is indicated by the boldface page number [26](#). The pages marked with arrowheads—[22](#)➤ and [31](#)—indicate the start and end of a section of the course dealing with the quadratic formula. Typically, this section would contain an introduction to the formula, its derivation, the statement (on page [26](#)) examples of its use, exercises and so on. If you want to really master the formula you should read this entire

How MATH⁴LIFE is numbered

range of pages and it's much easier to just show the start and end than to include an entry for every page in the range.

Finally, the pages like [35](#) and [44](#) point to later uses of the formula outside the primary discussion, usually applications to something we're covering later in the course.

Finally, links to files outside the MATH⁴LIFE web site [have this color](#). Like the cross-document links they'll take a while to load, possibly quite a while depending on how busy the external site is. Fortunately, there are only a few of these links mostly to help you in researching projects in the course.

Here are some examples of the three kinds of links if you want to experiment. First an internal or in-document link to the [OVERVIEW](#) at the top of this file. Next a couple of links to another document in the course: one to the top of [SECTION 5.1](#) and one to the main formula of that section, the [SIMPLE INTEREST FORMULA 5.1.6](#). Third, links to the definition of a [periodic rate](#) in the main text and to the index entry on [intermediate balances](#). Finally links to files outside the site, to the [Acrobat](#) web site mentioned above and to the author's [home page](#).

How MATH⁴LIFE is numbered

The MATH⁴LIFE course is not only a web site. It's a book too. Since the book can't use links for navigation it has to provide some other way of letting you refer to important ideas. Mathematicians like to do this by numbering such ideas. So all important ideas in the course have a number which is used when the idea is stated—either in the book or online—and also whenever the idea is referred to in the book or linked to online.

These numbers aren't so important online where you can just link directly to the statement of an idea from a reference to it. But in the



What MATH⁴LIFE's type styles and colors mean

book they are your only way of tracking down references. To make this as simple as possible a single numbering system is used for everything—formulas, figures tables etc. This has the big advantage that to find a concept with a bigger number you always flip towards the back of the book and to find one with a smaller number you flip towards the front. The same applies within sections online: bigger numbers are forward, smaller ones back.

The numbers used have three parts, the number of the chapter where the concept is found, the number of the section of that chapter, and a sequence number within that section. The parts are separated by periods. Let's look, for example, at some links to the section on yields. This is [SECTION 5.4](#): that is the *fourth* section of the chapter on interest and the time value of money. The main formula of that section is the [ANNUALIZED YIELD FORMULA 5.4.4](#): this is the *fourth* element in that section. I know that the [CONTINUOUS YIELD FORMULA 5.4.17](#) comes later in this section because it is element number 17. Likewise, I know that [SIMPLE INTEREST FORMULA 5.1.6](#) comes in an earlier section, the *first* section of this chapter.

What MATH⁴LIFE's type styles and colors mean

Colored type is used throughout MATH⁴LIFE to tell you what kind of material you are looking at. We've already seen the four colors used for [in-document links](#), [cross-document links](#), [index links](#) and [external site links](#). Here is what the other colors you'll see signify.

Most of the book is black like this paragraph. This is the color used for informal discussions, the part of the book you'll generally be reading. It's used in most of the book: when we are introducing a topic, deriving important ideas, making secondary points about formulas or problems, or dealing with the interactions between the real world and the subjects in MATH⁴LIFE.



What MATH⁴LIFE's type styles and colors mean

When you see this crimson color, *pay special attention*. It signals that an important concept, definition, or formula is being stated. To continue reading the section and do the problems it contains, you'll need to have a good grasp of the material in crimson. I'll go further: you should try to *memorize* the crimson material before trying the problems that use it. Learning math is a lot like learning a foreign language. The crimson concepts are our vocabulary and grammar. Working the problems is speaking the language. You can't expect to speak the language if you don't know the grammar and vocabulary. "Haste makes waste." You may save a few minutes by trying to skim the crimson material but then you'll just wind up wasting hours and getting frustrated trying to do the problems. To help you check whether you are ready to proceed, we plan to add frequent self-tests to the next version of MATH⁴LIFE.

This color is used for worked examples, usually problems just like those you'll be asked to do later. These illustrate how to use various formulas and methods. If you are having difficulty getting started on a problem, page back. The most recent example will usually provide a model for solving a very similar problem.

This color means homework. It is used for exercises, problems and projects that you are asked to work. Usually, I provide solutions to a few of the problems—in addition to the examples mentioned above—as models for you to use in working them.


Finally, a few words about type styles. Only one style—called small-caps—is important to recognize. It is used whenever numbered concepts are used. When a numbered element is introduced, its name and number are given in bold small-caps like **SIMPLE INTEREST FORMULA (2.1.4)** or **PROBLEM (2.2.23)**. When such a concept is referred to, ordinary small-caps are used as in SIMPLE INTEREST FORMULA (2.2.4) or PROBLEM (2.2.23). Bold type is also used in the index to distinguish the pages on which the primary or defining occurrence of a term occurs: see [USING LINKS IN MATH⁴LIFE DOCUMENTS](#) for more

details.

Here are the other type styles used in **MATH⁴LIFE** and what they signal. Don't worry if the type terms are unfamiliar; it's the signals that matter. We use italic type *like this* both for emphasis and to signal the informal introduction of a numbered concept which will be defined carefully later on. As you can see from this file, titles are given in bold sans-serif type **which looks like this**. References to computer files (including web addresses and other URLs) are given in a monospaced sanserif font as in `http://www.fordham.edu` or `readme.txt`.

Using graphs and pictures

“A picture is worth a thousand words” is a saying which applies often in mathematics. Throughout the course, you'll find graphs and pictures designed to illustrate important points. Often, you'll want to look at these in more detail than is available at the relatively low resolution of a computer monitor. In most web pages you'd be stuck because graphics have a single fixed resolution, but in **MATH⁴LIFE** all the graphics are *scalable* and can be drawn at any resolution. Just select the magnifying glass tool from the **Acrobat** toolbar and click on the spot you'd like to zoom in on. To zoom in further, click again. To zoom out, hold the **alt** or **option** key down—you'll see the plus sign in the magnifying glass change to a minus—and click. You can jump back to the standard full page view by choosing the **Fit** page command from the **Acrobat** View menu. Try it out on the shrunken graphic below—you'll see this graph at full size [SECTION 5.9](#).



Introduction

Goals of the course

This text is aimed at college freshmen, mostly intending to major in the humanities and social sciences, who are required to take a semester of college mathematics. For most of you, this is the last math course you'll ever take. And probably wouldn't be taking *this* course, if you weren't required to. As I often put it to my own students, I know that given a choice between taking this course and having a root canal with no anesthetic, you'd be in the dentist's office tomorrow morning. Most of you find math boring and irrelevant, and most of you find it difficult and frustrating.

That bothers *me*, but not because I think everyone should be a math major. Rather, it's because mathematics provides powerful and practical tools for dealing with important problems that are sure to arise in your everyday life, and because there's no good non-mathematical substitute for these tools in dealing with these problems. If you don't have a basic appreciation of these tools, you'll make poor decisions that will affect your personal finances, your health, and a host other areas of your future life.

So the goal of **MATH⁴LIFE** is to at least introduce you to some of the most important—and simplest—of these tools. I don't expect most of you to become producers of mathematical solutions. though I hope a few of you will be inspired by the course to pursue math-



Three questions that affect your life

ematics further. I don't even expect that you will all learn to be intelligent operators of mathematical tools, though I hope that many of you will. What I do want every student to take away from this course is an appreciation for the applicability of mathematical tools to real-life problems, and a readiness to ask others to help you use tools to address your problems. In short, I want you all to become intelligent future consumers of mathematics.

Three questions that affect your life

The hardest task I face is to capture your interest and the best way I can think of to achieve this is to focus the course on answering questions you'll have to face in your everyday lives—questions whose answers are important to everyone. Hence the name of the course. Here I'll just give three examples to whet your appetite.

Our first example is a financial one. You're probably going to have to be responsible for accumulating almost all of your retirement savings—I know that as freshmen, your retirement seems like the remotest of all possible concerns, but one of the most important lessons we'll learn is that not confronting this problem when you are young is a recipe for a world of pain. Few of you will have employers who provide pension benefits, and the Social Security system will have run out of money long before you ever qualify. Question one is: How much will you need to save for your retirement and how should you do so?

Let's assume that you hope to earn \$100,000 a year in your peak earning years. Making standard assumptions about retirement income (that it should be about 80% of your peak salary) and optimistic assumptions about the returns you can hope for on investments (ask your parents how well they slept during the fall of 2008, if you want to understand why these are very open to question), you'd need a nest egg of at least \$1,500,000 if you retired *today*. But you won't



Three questions that affect your life

retire for another 40–45 years. If we again are optimistic and assume inflation runs at historical averages of 3–4% a year, prices will quadruple before you retire. That means that *you* will need to have saved \$6,000,000.¹

If you work for 40 years, that means that you'll need to save—not earn but *save*—\$150,000 a year! How can you ever do this? This just doesn't seem possible. Fortunately it is, but only if you adhere to certain rules of savings. We'll learn these rules, and how to meet such a retirement goal, in [CHAPTER 5](#) where we'll also study other universal financial problems like how to effectively manage consumer credit and when it's better to buy and when to rent your home.

Our second question is a medical one. Mad cow disease or BSE (**bovine spongiform encephalopathy**) causes the gradual degeneration of the brain (encephalopathy) and spinal cord into a spongy (spongiform) mass—hence the scientific name. People usually acquire the disease by eating meat from infected cattle, hence the common name The disease often takes a decade or more to declare itself by progressive dementia (memory loss and hallucinations) and physical impairment (jerky movements, rigidity, ataxia and seizures) and another before death ensues. There is currently no cure or effective treatment for BSE and it's invariably fatal. Question two is: What are the chances you'll catch mad cow disease from the next burger you eat?²

We know the probability is not *very* high, but I expect that's a bit vague for most of you. Would you eat that burger if you knew the chance was 1 in a 1000? 1 in a 1,000,000? We need to quantify (that is give a definite numerical value to) the probability of a burger being infected before we can make an informed decision about whether or not we want to risk eating it.

¹If your parents are independently wealthy, you can skip to question two.

²Vegetarians can skip to question three.



Three questions that affect your life

It's the job of the U.S.D.A. to ensure the safety of your food supply so you'd expect them to be testing cattle for BSE. They are. But I claim that their testing program is deliberately designed *not* to answer the question of what fraction of cows have BSE, other than to say it's not *too* high. Feel safe now? Why I think the U.S.D.A. is deliberately not trying to answer question two is explained in [SECTION 2.2](#)—in fact, [CHAPTER 2](#) introduces many other common misconceptions that we'll try to clarify our study of probability in [CHAPTER 4](#) where we'll work out do the mathematical details.

The final example deals with global warming and speaks to our desire to be responsible citizens and use our votes to influence public policy. If *every* year were hotter than the one before for **anthropogenic** reasons (that is, because people were adding greenhouse gases to the atmosphere), we'd see this trend easily and it would be hard to claim that temperatures weren't rising. But the climate signal is influenced by many factors other than the anthropogenic ones and is very noisy. We observe a mix of hotter and cooler years and any given year is likely to be warmer in some places and cooler in others. Even in a single place, we'll see periods that are hotter than average and others that are cooler. How can we separate out all this noise and test whether there's a real trend towards higher temperatures?

One idea that has been extensively studied is to simply count the total number of record high and low temperatures reported from all places throughout the entire year. This approach eliminates a lot of the non-anthropogenic fluctuations by aggregating all locations and seasons. If average temperatures are not changing, then we'd expect roughly as many new lows as highs³ But if temperatures are really rising, we'd expect to see more new highs than new lows. In fact, the ratio of the number of new highs to the number of new lows would

³The total number of *each* decreases, because the longer we keep weather records the harder it becomes to surpass all our observations.

give us a rough measure of the strength of the warming trend.

There's one big catch here. If average temperatures are not changing, then we can view each new extreme as a coin toss with outcomes of high or low instead of head or tail. If we toss a coin, say, 100 times, we expect *about* 50 heads and 50 tails. But we don't expect *exactly* 50 of each because of random variation. In one series of tosses we might see 45 heads, and then in the next 57, and in the next 52 and so on. Likewise, even if average temperatures are not changing, there'll be years when we see more new highs than lows (and vice versa). So we might observe such years just due to random deviation and not because of a warming trend. Suppose we do see a preponderance of new highs over new lows in one year, or even in a run of consecutive years. How do we tell if this is due to chance or to a real warming trend?

The analogous question for coin tossing would be: "How many heads do I have to see in 100 tosses before I am convinced the coin comes up heads more often than tails?". Pretty clearly, we need to see more than, say, 52 or even 57. But if we saw 100 heads, or even 80, we'd be pretty convinced the chance of heads on each toss must be bigger than 50%. So what is the point between observations like 52 and 57 that we attribute to chance and those like 80 and 100 that we think indicate a coin biased towards heads. An answer to this question can be applied to highs and lows. If the observed preponderance of highs or lows is small enough, we attribute it to random effects. But if it's big enough, we are convinced there's trend. The tough question is then "How big is big enough?".

This is the kind of question that is answered by statistics. This is a hugely important subject. Problems of distinguishing between random and real effects arise in almost every area of modern research—think of testing drugs and other medical procedures, or analyzing surveys and market research to name just two. Really commanding the ideas of statistics calls for more math than we'll do here but we

will cover just enough to answer the questions above dealing with coin tosses and global warming (and many others) in the final section of [CHAPTER 4](#).

How to learn mathematics effectively

Even if, as I hope, the preceding questions have convinced that learning mathematics need not be a total waste of time, I'm sure that most of you want to get a good grade with the *least possible effort*. I'll close this introduction with a bit of advice for doing so.

The key idea was perhaps best put by Albert Einstein when he said, "Everything should be as simple as possible, but *not* simpler than possible." Here, what this means is that you must approach learning mathematics in a fairly disciplined way. In fact, the big difference between students who find math easy and those who find it hard is *not* in their intelligence: it's in the discipline with which they approach learning the subject.

You'll frequently be tempted to save time by cutting corners in the "simple as possible" process outlined below. If you do so, it'll become "simpler than possible" and the result will be that you leave gaps and won't learn *effectively*. Even in the short run, you'll wind up spending *more* time than necessary, and you'll find the subject much more frustrating. As you proceed, those little gaps add up, and pretty soon, even topics that should be easy start to seem difficult. In the long run, you learn less and less efficiently, spending more and more time, and absorbing less and less. So here's my four step plan for learning mathematics as painlessly as possible.

1. SUMMARIZE AND MEMORIZE First, read the text, summarize in your notebook the new definitions, formulas and theorems you have encountered, and memorize this summary. You probably will not understand everything you are memorizing at this stage. Understanding comes through practice in steps two and three. However, if you

do not have the basic notions at your fingertips you will get stuck in these steps. As a result you will waste time, become frustrated and likely never achieve understanding. Most problems in learning mathematics are due to skipping or short-cutting this step. Remember Step 1 does not *take* a lot of time but skipping it can *cost* a lot of time.

2. PRACTICE AND ASK Now begin to understand the ideas you have seen in Step 1 by working through, first, the examples then, the solved exercises that use these ideas. You should expect to spend the largest portion of your time on this step. Because the amount of practice needed varies greatly from student to student, and—even for the same student—from concept to concept, it's hard to say exactly how much time you should spend here. One key point: *Start early!*, because, as we'll see below you may want to break and seek help.

Begin by reworking examples in the text, checking your answers against those in the solutions. While doing this you may need to refer to your summary from Step 1, and to return to the text or class notes to see how similar problems are worked. Do you find yourself flipping back in the text a lot? That's a sure sign that you have cheated on Step 1. Stop right there and go back and review your summary of the concepts you're trying to master.

Sometimes you will get stuck: either you cannot see how to do the problem at all or you cannot get the correct answer. If so, do not waste a lot of time poring over the problem on which you are stuck. Give yourself 5-10 minutes to try the problem. Then, if you are still stuck, make a note to ask me, a friend, or a Help Room tutor about the problem which you were unable to solve. Taking advantage of this kind of help is the main advantage of studying mathematics in a classroom setting. Spending a long time staring at a problem is usually a waste of your time and a recipe for frustration.

Of course, if you have an assignment due tomorrow—or worse, later



today—you won't have time to seek help. That's why it's so important to start early.

3. SELF-TEST When you have got the hang of these problems, you should test yourself by trying the solved problems *without peeking at the solutions*. You are ready to move on when you can do these problems and you “know” your answers are right. If you are not sure about your answers then go back and look again at the worked examples trying to look only at the questions and develop our own answers.

Now go back to the solved problems and try to write your own solutions without peeking. Then check your answers. If you are wrong, you need more practice in Step 2 (and maybe, again Step 1).

If you are getting the solved problems right, try the unsolved problems—in real life, there are no answers in the back of the book. Step 3 is complete when you feel confident about your answers to these unsolved problems too.

4. VALIDATE This is the point of your class homework assignments. These tell you whether or not you have really learned a given topic. If you have, you will know how to work the problems on these assignments. If not, you should ask yourself which of the first 3 steps you failed to complete properly. If in doubt, ask your instructor for guidance. Don't just let a topic you did poorly on slide: next week's homework will often require you to use ideas from this week's. Therefore, it is important to make sure you do not let gaps in your mastery accumulate as the course goes on.

To summarize, let me emphasize two points. First, Don't cheat on the four steps—you'll only cheat yourself. When learning mathematics, it is essential to go through all these steps and to do so *in the order shown above*. Leave any one out and you will only find attempting the next one to be frustrating. What's worse, instead of saving time, you will *waste* it.



How to learn mathematics effectively

Second, when you do get stuck, *Ask!* Ask your instructor or go to your department's Help Room, or ask a friend when you are having difficulties. Otherwise, you are missing out on the greatest benefit of studying math in a organized university setting and you are not making the best use of the tuition you have paid.



Chapter 1

Back to Basics

What the topics in this chapter have in common is that they involve basic notions and skills that we'll need at many points in the rest of the course. They form a tool chest we'll use in our more specialized projects in the course. You'll probably have seen most of these topics in high school and what I'll say here will just be a review.

One topic merits a special mention. Mistakes with [SECTION 1.1](#) are by far the biggest source of dope slapping, "How could I be so stupid?" errors throughout the course. Even if you think you know this material well, *please* read this section over carefully, try to absorb the advice it gives in avoiding such errors and work the problems in it.

1.1 Order of operations

Let's begin this section with very easy multiple choice question.

BRUMER'S PROBLEM 1.1.1: Is the quantity -3^2 equal to:

- i) $+9$; or,
- ii) -9 .



1.1 Order of operations

Pencils down please. Are you sure of your answer? The problem is named after Armand Brumer, an emeritus professor in my department. I ask it at the first meeting of every freshman class I teach. Most years, at least two-thirds of my students answer it incorrectly. I've never had class where two-thirds gave the right answer. Still sure of your answer? Good.

What's the issue here? It's that the expression -3^2 is somewhat ambiguous. It's clear that it involves two operations, a minus and a square. What's not clear is the order these should be performed in. And this order matters!. If we minus first and then square, we get $(-3)^2 = (-3)(-3) = +9$. If we square first, we get $-(3^2) = -(9) = -9$.

Get the order of operations wrong and, even if you do all the right operations, you get the answer wrong. Let me repeat what I said at the start of this chapter. Mistakes with **order of operations** are by far the biggest source of dope slapping, "How could I be so stupid?" errors in freshman mathematics courses of all types.

How can you avoid such errors? The two little calculations above hide the answer, the most important point of this section. We can *always* clarify the meaning of *any* potentially ambiguous expression by adding parentheses. In the first calculation, writing $(-3)^2$ made it clear that I wanted to minus before squaring. In the second, writing $-(3^2)$ made it clear that I wanted to square before minusing.

PARENTHESES RULE! 1.1.2: *Any operation(s) enclosed in a pair of parentheses must be completed before performing any operation(s) outside that pair of parentheses. So when in doubt, add parentheses!*

The moral is clear. Make every effort to train yourself to add parentheses whenever you're the least bit unsure whether an expression you are writing down is completely unambiguous. I cannot urge this strongly enough. Parentheses are free. They're a renewable, biodegradable resource and emit no greenhouse gases. Learn to use



1.1 Order of operations

them freely and you've eliminated the biggest source of errors. Here are some problems for practice.

PROBLEM 1.1.3: Each of the following expressions can be interpreted in two ways corresponding to two different orders of performing the indicated operations. Show how to obtain each of the two interpretations unambiguously by adding parentheses to the expression and calculate the value each gives. I have worked the first few to get you started.

i) 3^{2^3}

Solution

This can mean either $3^{(2^3)} = 3^8 = 6561$ or $(3^2)^3 = 9^3 = 729$.

ii) $14/7 + 7$

Solution

This can mean either $(14/7) + 7 = 2 + 7 = 9$ or $14/(7 + 7) = 14/14 = 1$.

iii) $\frac{\frac{72}{6}}{3}$

Solution

This can mean either $\frac{\left(\frac{72}{6}\right)}{3} = \frac{12}{3} = 4$ or $\frac{72}{\left(\frac{6}{3}\right)} = \frac{72}{2} = 36$.

iv) $23 - 17 + 5$

v) $16/8 \times a$

vi) $-6 + 4/2$

vii) $2x^2$

viii) $27^2/3$

ix) $6^4 - 2$

x) $30 - 8 \cdot 4$

xi) $1 + p^T$

xii) $8 - 4 - 4$

By the way the answer to [BRUMER'S PROBLEM 1.1.1](#) is -9 . I *know* most of you said $+9$ and I hope you'll draw the lesson. No order of operations error is too easy or stupid to fall into. It's just so easy to say,



1.1 Order of operations

“OK, minus 3, now square, 9.”. It just happens to be wrong! If you said -9 , I hope you’ll draw the same lesson. You were right this time but the only reliable defense against order of operations errors in the future is to use parentheses liberally.

In a better world, this section would end here. Not using parentheses would be a major felony, recidivists would be rapidly reassigned to careers in the postal service, and I’d never have to worry about order of operations errors again. The problem is that even if you resolve to start putting parentheses under your pillow when you go to bed, you’re going to run into many ambiguous expressions written by other people. How do you decide what they mean?

The answer is dangerously simple. For centuries, people who work with mathematical expressions have used a set of rules that resolve ambiguities by specifying a preferred order of operations in each ambiguous case. It’s a simple answer because the rules are simple: a couple of acronyms **POGEMDAS** and **BTDLTR** are all you need to remember. It’s a dangerous solution because it’s very easy to ignore these conventions even if you know them. For example, I’m sure most of you learned the **PEMDAS** rule in high school. From asking my classes over the years, I know that includes most of those who violated the E-before-M convention and gave the *wrong* answer $+9$ to **BRUMER’S PROBLEM 1.1.1**. Knowing the conventions is only effective if you are ceaselessly vigilant about applying them. To make a somewhat vivid but very apposite analogy, **POGEMDAS** is to using parentheses as abstinence is to using a condom. The Math Help Rooms of America are filled with students who promised God they weren’t going to disorder operations before the final.

But we do need to know these conventions so we can correctly evaluate other people’s ambiguous expressions. So that I won’t have to digress once we start to list the precedence rules, I’m first going to review some **OTHER GROUPING** operations that we’ll be using and that are simply ignored in the **PEMDAS** rule that you may



1.1 Order of operations

have learned in high school. These operations are function like honorary parentheses: they group together other “interior” operations and they should be done before anything else—except, of course, operations inside ordinary parentheses.

The four GROUPING OPERATIONS we’ll see in this course are functions, horizontal bar fractions, superscript exponents and radicals. Let’s quickly review how each works. As we do so, I’ll point out some common ways order-of-operations errors arise when working with them. Sometimes these errors occur when you’re trying to simplify algebraic expressions on paper. But for most of you, they’ll generally bite when you are simply using your calculator to evaluate these expressions.

The function notation $f(\text{---})$ gives us a way to work with any operation that we might find useful by simply giving the operation a name. In this course, we’ll work most with the natural logarithm and exponential functions which take a single argument, as in $\ln(\frac{S}{B})$ or $\exp(0.01r \cdot y)$. But we’ll also work with functions having more than one argument like the combination function $C(n, r)$ and the permutation function $P(n, r)$. And you’ve probably seen lots of other examples that we won’t need (the trigonometric functions and their inverses, financial functions, ...).

A function f really denotes a rule for calculating a function value $f(x)$ given an argument or input value x , but, as in the examples above (where x was $\frac{S}{B}$ or $\frac{r}{100}y$), you may need to perform other operations to compute this argument. But you need both pieces—the f and the x —to get the function value. If you separate the parts of a function expression what you get is nonsense.

A very common mistake of this type when simplifying expressions to incorrectly write $\ln(x + y) = \ln(x) + \ln(y)$ by “applying” the distributive law. Another mistake that’s even easier to slip into occurs when the rule for f involves algebraic operations. Suppose $f(x) = x^2 - x$. If you’re smart, you’ll write this $f(x) = (x^2 - x)$.



1.1 Order of operations

What possible point can the parentheses have, since the entire value of f lies inside them? The answer comes when you try compute an expression like $f(5) - f(2)$ using a calculator. Let's try it. A very common answer is to enter $5^{\wedge}2 - 5 - 2^{\wedge}2 - 2 = 25 - 5 - 4 - 2 = 14$. Looks pretty good, but it's wrong. Adding the parentheses shows why: now we get $(5^{\wedge}2 - 5) - (2^{\wedge}2 - 2) = 25 - 5 - 4 + 2 = 18$. In this case, the parentheses reminded us that we need to distribute the minus sign in $-f(2)$ across the *entire* value of f . Putting the minus and the formula for $f(2)$ side-by-side, as in the first version, doesn't achieve this.

Radicals function in much the same way: you must perform any operations inside or under the radical before you can evaluate it. The most common is the square root radical, familiar from such expressions as the $\sqrt{a^2 + b^2}$ that appears in Pythagoras' theorem. As with a function, you can't separate the radical from its argument. A common simplification error of this type is to write $\sqrt{a^2 + b^2} = \sqrt{a^2} + \sqrt{b^2} = a + b$. We'll also work with higher radicals. A superscript y at the left of the radical, as in $\sqrt[y]{\frac{S}{B}}$, indicates a y^{th} root.

Next, fractions come in two flavors. There's the numerator-beside-denominator or slash form, as in $2/3$, you use to enter a fraction in your calculator and the numerator-over-denominator or horizontal bar form, as in $\frac{2}{3}$. In this book, I *only* use horizontal bar fractions (except when illustrating how to enter formulae into a calculator) and I strongly encourage you to cultivate the same habit. The reason is that horizontal bar fractions prevent you from making a lot of errors by forcing you to group any operations that go into the numerator or into the denominator.

In slash fractions, you need to use parentheses to group these operations or your value will often be wrong. Typical examples is $\frac{20-10}{5} = \frac{10}{5} = 2$ and $\frac{24}{6+2} = \frac{24}{8} = 3$. You get the right answer by writing $(20 - 10)/5 = 10/5 = 2$ or $24/(6 + 2) = 24/8 = 3$ but you must always parenthesize aggressively or you won't get the answer



1.1 Order of operations

you want. Since the only way to enter fractions into a calculator is in slash form, you need to be constantly vigilant that you have added any necessary parentheses when doing so.

Decades of experience have taught me that if I ask questions that contain the two calculations above, I will see several answers of 18 and 6. These come from students who used their calculator to do the calculation and entered $20 - 10/5$ or $24/6 + 2$. The calculator did what they said not what they meant: it knows PEMDAS (see [POGEMDAS 1.1.5](#) below if you don't) and does the division before the addition or subtraction. For more, see [EXAMPLE 1.1.8](#).

Finally, exponentials also come in two flavors. There's the vertical or superscript r^u form with the exponent above the base that we all use automatically in handwritten mathematics—for me, at the board and for you, in your notes—and there's the base-beside-exponent form r^u that you use to enter exponentials in your calculator. As with fractions, it's in the transition from the superscript form that we naturally use to the base-beside-exponent form required by your calculator that most errors creep in.

In this course, I *always* write exponentials in superscript form in this book. Once again, the reason is that in this form any operations that take place inside the exponent are automatically visually grouped together. In base-beside-exponent form, any operations that take place inside the exponent must be parenthesized or your value will virtually *always* be wrong. Typical examples are r^{u+1} and $e^{20 \cdot y}$. These are *not* the same as $r^u + 1$ or $e^{20 * y}$ —for the difference see [EXAMPLE 1.1.8](#)

CALCULATOR PARENTHESES RULE 1.1.4: *When entering an expression into a calculator:*

- i) *Always include all parentheses present in the formula you are working with.*
- ii) *Always add parentheses around any function value if they are not already present.*



1.1 Order of operations

iii) *Always add parentheses around the numerator or any fraction and around the denominator of any fraction if they are not already present.*

iv) *Always add parentheses around any exponent if they are not already present.*

Remember, never omit parentheses. When in doubt, add parentheses!

Make these rules a habit early and you'll save yourself from making many mistakes like those discussed above. One last time, make parentheses your friends. They're the only reliable protection against order of operations errors. We're ready to learn the rules for interpreting mathematics written by others who like to live dangerously.

POGEMDAS 1.1.5: *POGEMDAS is an acronym for PARENTHESES, OTHER GROUPINGS, EXPONENTIATION, MULTIPLICATION, DIVISION, ADDITION, SUBTRACTION. The order in which the words appear indicates the precedence of the corresponding operations (that is, the order in which they should be performed).*

So first, PARENTHESES RULE! 1.1.2: perform any operations in parentheses before any outside.

Now come the OTHER GROUPINGS. Evaluate any functions as soon as the value of the argument inside the function's parentheses—the "(—)" above—has been computed. Perform any operation inside a RADICAL before taking the RADICAL. Compute the numerator and denominator of a HORIZONTAL bar fraction and then take their quotient.

Next, perform any EXPONENTIATIONS. After these, perform any MULTIPLICATIONS and DIVISIONS. Finally, carry out any ADDITIONS and SUBTRACTIONS.

But there's a catch. POGEMDAS allows ties of two kinds. First, there are ties when two operations represented the same letter in POGEMDAS are adjacent: which do we perform first? Second, there are two pairs of different operations that need to be viewed as ties when they are adjacent to each other. MULTIPLICATION and DIVISION have the



1.1 Order of operations

same precedence; and, ADDITION and SUBTRACTION have the same precedence.

We need some convention to break the ties. Once again, the only safe way is to eliminate the ties by adding parentheses. But when others have ignored this suggestion, the BTDLTR rule comes in.

BTDLTR 1.1.6: *BTDLTR is an acronym for BREAK TIES TOP DOWN LEFT TO RIGHT. Ties between operations that are written vertically—HORIZONTALS and EXPONENTIALS—are broken using a conventional top-down order. Ties between operations that are written horizontally—between MULTIPLICATIONS and DIVISIONS or between ADDITIONS and SUBTRACTIONS are broken using a conventional left-to-right order. An informal way of expressing these tie-breakers is that we break ties by performing the operations in the order in which we encounter them (at least if we're using a top-down, left-right script like the Latin alphabet). This makes breaking ties with BTDLTR easy, since it asks us to break them in natural reading—or, on your calculator, keying—order.*

A few comments. These conventions are a bit intricate. We'll soon work a lot of examples that I hope will make everything clear. But, isn't using parentheses really a lot easier than using these conventions? You may object that my conventions are a lot more complicated than the ones you were taught in high school. You're probably right, but the reason is that what you taught in high school leaves out many cases. Whole classes of operations (the OTHER GROUPINGS) are ignored and breaking ties is treated incompletely or not at all.

Conventions exist so we'll be able to assign a preferred order to *any* set of operations and it's not a satisfactory solution to resolve only the easiest or most common conflicts. In my survey of math help websites, I found none aimed at high school students that give a complete set of solutions. In surveying, the top 50 or so Google™ hits on “order of operations” the only really complete set of rules I found was (surprise) in a [Wikipedia article](#).

1.1 Order of operations

While I am ranting, would somebody please explain to me why it's easier to remember "Please excuse my dear Aunt Sally" than PEMDAS? I won't bother trying to work Sally into POGEMDAS.

OK. On to some examples. Lots of them.

PROBLEM 1.1.7: In [PROBLEM 1.1.3](#), you interpreted each of the following expression in two different ways. Which interpretation was the "right" one according to the conventions [POGEMDAS 1.1.5](#) and [BTDLTR 1.1.6](#). I have worked the first few to get you started.

i) 3^{2^3}

Solution

Recall that this might mean $3^{(2^3)} = 3^8 = 6561$ or $(3^2)^3 = 9^3 = 729$. Here we have a tie between two EXPONENTIALS which we break TOP DOWN since these are written vertically. So $3^{(2^3)} = 3^8 = 6561$ is the right interpretation here.

ii) $14/7 + 7$

Solution

Here we have a DIVISION and an ADDITION. Since the former has precedence over the latter, $(14/7) + 7 = 2 + 7 = 9$ is the conventional interpretation here.

iii) $\frac{\frac{72}{6}}{3}$

Solution

Here we have a tie between two horizontal bars which we break TOP DOWN so $\frac{\left(\frac{72}{6}\right)}{3} = \frac{12}{3} = 4$ is the conventional answer.

iv) $23 - 17 + 5$

Solution

Here we have a tie between a SUBTRACTION and an ADDITION which we break LEFT TO RIGHT so $(23 - 17) + 5 = 6 + 5 = 11$ is "right" here.

v) $16/8 \times a$

vi) $-6 + 4/2$



1.1 Order of operations

vii) $2x^2$

viii) $27^2/3$

ix) $6^4 - 2$

x) $30 - 8 \cdot 4$

xi) $1 + p^T$

xii) $8 - 4 - 4$

Let's look at some examples that illustrate typical cases where the [CALCULATOR PARENTHESES RULE 1.1.4](#) is needed (and where it's easy to leave out the extra parentheses needed). Each of the formulae in [EXAMPLE 1.1.8](#) below comes up later in the course. In this example, we'll see how parentheses that are not needed when formulae are written with horizontal bar fractions and superscript exponentials become necessary when we rewrite the formulae using slash fractions and base-beside-exponent exponentials. In particular, you very often need to add parentheses when making such a conversion.

I'll choose one pair of parentheses in each formula and decide whether deleting this pair of parentheses from the formula changes its meaning or leaves it the same, using the conventions [POGEMDAS 1.1.5](#) and [BTDLTR 1.1.6](#).

EXAMPLE 1.1.8:

i) $\frac{(1-r^{(u+1)})}{(1-r)}$

Solution

Removing the parentheses in the numerator gives $\frac{1-r^{u+1}}{(1-r)}$. This means the same thing because a horizontal bar fraction always says to calculate both numerator and denominator *before* the division. The effect is the same as if the numerator were parenthesized, as in the original formula. For the same reason, removing the parentheses in the denominator leaves the meaning unchanged.

ii) $(1 - r^{u+1}) / (1 - r)$

Solution



1.1 Order of operations

This is the same formula as in [i\)](#) but written with the quotient in slash form. Removing the parentheses in the numerator gives $1 - r^{u+1}/(1 - r)$ and this means $1 - (r^{u+1}/(1 - r))$ because now the DIVISION must precede the SUBTRACTION to its left. With the parentheses, the subtraction came first. So these are different, as we can see by putting everything over a common denominator and getting

$$1 - r^{u+1}/(1 - r) = 1 - (r^{u+1}/(1 - r)) = \frac{(1-r) \cdot 1 - r^{u+1}}{(1-r)} = \frac{(1-r-r^{u+1})}{(1-r)}.$$

Likewise, removing the parentheses from the denominator changes the meaning. Now the subtraction on the right must come after the division instead of before. In fact, $(1 - r^{u+1})/1 - r = 1 - r^{u+1} - r$ has no denominator at all!

iii) $r^{(u+1)}$

Solution

Removing these parentheses has no effect because in r^{u+1} the superscripted exponent groups the addition so that the $u + 1$ comes before the exponential base r .

iv) $r \wedge (u + 1)$

Solution

This is the same formula as in [iii\)](#) but written in base-beside-exponent form. Now removing the parentheses does change the meaning. By POGEMDAS, the EXPONENTIAL precedes the ADDITION so $r \wedge u + 1 = (r \wedge u) + 1$. Going back to superscript form— $r^u + 1$ makes the difference clear as does plugging in almost any values for r and u . For example, if $r = u = 2$, then $r \wedge (u + 1) = 2 \wedge (2 + 1) = 2^3 = 8$ while $r \wedge u + 1 = (2 \wedge 2) + 1 = 4 + 1 = 5$.

v) $e^{(.05 \cdot y)}$

Solution

Removing these parentheses has no effect because in $e^{(.05 \cdot y)}$ the superscripted exponent groups the product so that the $.05 \cdot y$ comes before the exponential base e .

vi) $e \wedge (.05 * y)$

Solution



1.1 Order of operations

This is the same formula as in **v)** but written in base-beside-exponent form. Now removing the parentheses does change the meaning. By POGEMDAS, the EXPONENTIAL precedes the MULTIPLICATION so $e^{.05} * y = (e^{.05}) * y$. Going back to superscript form— $e^{.05} \cdot y$ makes the difference clear—the y is no longer part of the exponent as does plugging in almost any value for y and u . For example, if $y = 20$, then $e^{.05} * y = e^{.05} * 20 = 1.051271096 * 20 = 21.02542192$ while $e^{(.05 * y)} = e^{(.05 * 20)} = e^1 = 2.718281828$.

Moral: when entering displayed formulae containing function values, horizontal bar fractions and superscript exponents into calculators, always pay careful attention to the **CALCULATOR PARENTHESES RULE 1.1.4**. When in doubt, add parentheses!

Before I give you some calculator problems, there's one more procedure we need to recall. What do we do when we have a fraction in which the numerator or the denominator is itself a fraction—or both are? There's are two easy procedures to handle all such cases. One is to invert the denominator (even if its not a fraction!) and then multiply by it. The other is to clear any denominators by multiplying by them both above and below. Here are some practice problem with the first few cases worked as models.

PROBLEM 1.1.9: Convert each of the fractions of fractions below to a simple fraction (with no divisions in either the numerator or denominator) by both the methods outlined above.

i) $\frac{\frac{6}{3}}{\frac{2}{4}}$

Solution

This one is a typical worst case. First, you can see that the middle bar here is both wider and thicker than the bars above and below it. So it groups it's numerator and denominator each of which just happens to be a fraction.

- a. Dividing by $\frac{2}{4}$ is the same as multiplying by its inverse $\frac{4}{2}$ so $\frac{6}{3} \cdot \frac{4}{2} = \frac{24}{6} = 4$.



1.1 Order of operations

b. To clear two denominators—3 and 4—we multiply by them above and below to get $\frac{\frac{6}{3}}{\frac{4}{4}} = \frac{3 \cdot 4}{3 \cdot 4} = 1$. Cancelling the 3s above and the 4s below, we again get $\frac{4 \cdot 6}{3 \cdot 2} = \frac{24}{6} = 4$.

ii) $\frac{2x+1}{\frac{1}{x}}$

Solution

Again, the larger bar tells us to group the $2x + 1$ and the $\frac{1}{x}$. This looks easier as there's only a fraction in the denominator, but we need to remember to ensure that the numerator $2x + 1$ is calculated before any other operations and the only way to ensure this is to parenthesize it.

a. Dividing by $\frac{1}{x}$ is the same as multiplying by its inverse $\frac{x}{1}$ so $\frac{2x+1}{\frac{1}{x}} = (2x + 1) \frac{x}{1} = (2x + 1)x = 2x^2 + x$. Note that if we do not parenthesize, then we get $2x + 1 \frac{x}{1} = 2x + x = 3x$ which is wrong.

b. To clear the x in the denominator, we multiply by it above and below to get $\frac{2x+1}{\frac{1}{x}} = \frac{x(2x+1)}{x \frac{1}{x}}$. Cancelling the x s below, we again get $\frac{x(2x+1)}{1} = (2x + 1)x = 2x^2 + x$. Again, not parenthesizing gives an incorrect answer: $\frac{x2x+1}{x \frac{1}{x}} = \frac{2x^2+1}{1} = 2x^2 + 1$.

iii) $\frac{\frac{1}{30}}{\frac{n-1}{1}}$

iv) $\frac{\frac{S}{B}}{\frac{T-i}{1}}$

PROBLEM 1.1.10:

i) Each of the following expressions is the displayed form of a formula we'll see later in the course with values inserted for the variables. Translate each expression into a form in which you can enter it into your calculator, being sure to add any parentheses necessary to avoid changing its value.

a. $\$2400(1 + \frac{7.23}{100 \cdot 365})^{730}$

b. $100 \left(\left(\frac{5000}{4000} \right)^{\frac{1}{3}} - 1 \right)$

c. $\frac{6.25}{100 \cdot 4} \cdot \$1256.74 \cdot 16$

d. $\frac{10}{1} \frac{9}{2} \frac{8}{3} \frac{7}{4}$



1.2 Handling numbers in various formats

ii) Use your calculator to evaluate your version of each expression in i). If you don't get the value indicated below, go back and compare the original expression and your version to see where they differ. Correct your version and reevaluate it.

a. \$2773.345291.

b. 7.721734500

c. \$314.1850000.

d. 210.

1.2 Handling numbers in various formats

This section gives recommendations for dealing with complicated numerical answers. In particular, we'll see how to answer three questions involving such answers that arise frequently in the course:

- i) How do we handle whole number answers when they get really big—that is, too big even for our calculators to deal with?
- ii) If an computation leads to a fractional answer and the numerator and denominator of the fraction are big, when should we live with the complicated fraction and when is it better to get our calculator to give its decimal form?
- iii) When do we want to leave decimal answers unrounded, when do we want to round them, and when we *do* round, how do we decide how many decimals to keep in the rounded answer?

Working with big numbers

The answer to the question, “What's a *big* whole number?” is relative. When you were in grade school, you probably thought that any number greater than, say, 1000 was big. Today you might think that numbers in the millions or billions are big. The answer is also



1.2 Handling numbers in various formats

relative because, whatever size *you* may think makes a number big, there are numbers a lot bigger. And in this course, we'll see lots of operations that lead to really big numbers, even starting from small numbers. When we deal with ordering sets in [COUNTING ORDERINGS](#), we'll want to consider the number of ways of ordering a set of 100 elements, which is denoted $P(100,100)$ or $100!$ and is the product of the whole numbers from 1 to 100. I'm sure none of you would say that the number 100 is big. But I think you'll also all agree that

$100! = 933262154439441526816992388562667004907159682643816214685929638952175999932299156089414639761565182862536979208272237582511852109168640000000000000000000000$

is very big (even if the type I have had to use is very small). For the record, it has 158 digits.

Numbers this big won't come up too often, not because the topics we cover don't lead to problems involving such big numbers, but because the problems you'll be asked to work have been carefully chosen to avoid them. When they do come up, however, they pose one major problem. They're too big even for your calculator. A typical calculator today displays numbers to between 10 and 12 digits. That may seem like a lot but it's insufficient even to handle many of the real-world problems we'll come across in studying mathematics of finance. For example, when I looked it up on December 19, 2008, the [U.S. National Debt](#) was \$10598195081084.56—that 16 digit number is a little over ten trillion dollars (and your share is about \$35,000).

Your calculator can sleaze out on a number like the National Debt by ignoring the last few digits—what’s a few thousand dollars between citizens?—and representing the number in scientific notation as

$$1.05981950811 \times 10^{13} \quad \text{or} \quad 1.05981950811\text{e}13.$$

SCIENTIFIC NOTATION 1.2.1: A number is represented in *scientific notation* to d significant digits by giving a non-zero leading digit, a decimal point, $(d - 1)$ digits to the right of the decimal (making, with the leading digit, a total of d digits) and a power of 10 either in the usual form or in the “calculator” form with an ‘e’ followed by the exponent.



1.2 Handling numbers in various formats

The examples preceding the definition are given to 12 places (a leading digit and 11 decimals). Notice that the last decimal is not the 0 that was the 12th digit of the National Debt, but a 1. That's because I rounded at this place and, because the next digit is an 8, I rounded *up* to get the 1. (We'll review rounding in more detail in a moment.) Of course, I can only correctly round in this way if I somehow am given at least 1 more digit of the number I'm converting to scientific form. For the numbers in scientific notation that you'll get from your calculator, you won't need to worry about this because virtually all calculators today carry a couple of "spare" significant digits in their internal calculations exactly to ensure that the digits they display are correctly rounded.

But your calculator has limitations. We'll see examples where it gives the wrong answer to apparently straightforward computations when we look at [THE CONTINUOUS APPROXIMATION](#). And most calculators allow only two digits for the exponent in a scientific form, so they have no way at all of representing 100!. For the record, it's $9.33262154 \times 10^{157}$.

What's really called for, when it's necessary to handle big numbers like this, is a better calculator. Such calculators are universally implemented as applications that run on your PC. They go by the name of computer algebra systems or symbolic computation packages: common ones are Mathematica™, Maple™ and MathLab™. Even though the CPUs on most PCs can only handle about 27 digits at a time, these systems can compute with very large numbers by using software techniques to break up calculations into pieces that the CPU can handle and then reassembling a final answer. I used Maple™ to compute the value of 100! above, and it's how I computed the other very large numbers that will occasionally come up to illustrate ideas in the course.

I should warn you that it's easy to choke even these algebra systems. The size of the numbers they can deal with is, in principle, limited



1.2 Handling numbers in various formats

not by the CPU but only by the amount of memory on your PC. In principle, because when the numbers get really big—in the millions of digits, for example—it can take an impractically long time to find the answer. And sometimes all the memory, even your entire hard drive, is just too small. None of these systems can find $22!!$ (the factorial of $22!$) for the very good reason that this number has about 10^{21} *digits* and even to list these digits, let alone compute them, all the memory and hard disk storage on *all the computers on earth* would not suffice.

Trust me. I'll warn you when we're going to look at a number that might choke your calculator. Otherwise, in this course, you can pretty much relax. If you ever do need to work with really big numbers in another course, there are two options. Invest in a computer algebra system: Maple has student editions for about the price of a graphing calculator. Or, if you just need a quick scientific answer that's too big for your calculator, you can always just `googl` it. Try googling $100!$ now—just enter it in the search box, or use the link I have provided—to check the scientific form above. But again a warning is in order as you'll see if you compare `google`'s answers to $170!$ and $171!$.

PROBLEM 1.2.2: Write down each of the numbers below in scientific notation as requested:

- i) 2^{30} to 8 significant digits.
- ii) $\frac{28374948973}{37647637}$ to 6 significant digits.
- iii) 12345678 to 7 significant digits.

Fractions versus decimals

In a lot of the formulae we'll deal with later, the answer we're after is expressed as a fraction with a whole number numerator and a whole number denominator. If you're like most students today, your first move on seeing such a fraction is to divide the numerator by the



1.2 Handling numbers in various formats

denominator in your calculator to convert it to a decimal approximation. In this subsection, I'd just like to point out that you often lose information in the fraction-to-decimal conversion process. This loss of information is of two kinds.

The first kind of loss is a loss of meaning. Very often the numerator of a fraction comes to us as the answer to one question and the denominator comes to us as the answer to a second question. The vast majority of the probabilities we'll compute in [CHAPTER 4](#) will arise in this way with the both numerator and denominator arising as the answer to a counting or "How many?" question. For example, the probability of being dealt a full house at poker is $\frac{3744}{2598960}$. Here 2598960 is the number of possible **poker** hands and 3744 is the number of those hands that are **full houses**. (We'll see how to derive these numbers in [A CLASSIC EXAMPLE: POKER RANKINGS](#). Just take them on faith now.) The fraction $\frac{3744}{2598960}$ preserves both these answers. As soon as I convert it to a decimal approximation— $\frac{3744}{2598960} \simeq 0.0014405762304922$ —they vanish. Even if I cancel common factors and rewrite the fraction in lowest terms— $\frac{3744}{2598960} = \frac{6}{4165}$ —they vanish.

CONVERTING FRACTIONS TO DECIMALS 1.2.3: *When you have arrived at the numerator and denominator of a fraction separately, as the answers to 2 sub-questions in a problem, don't convert the fraction to a decimal.*

The goal of this recommendation is to preserve the meanings of the numerator and denominator as long as possible. However, not converting to decimal form often has practical benefits. It's also very common for problems to contain many such two-answer fractions with the *same* denominator. You can probably guess why: because these denominators are all the answer to the *same* question.

For example, the number of poker hands that are flushes is 5148 and the number that are straights is 10240. The probability of being dealt a flush is $\frac{5148}{2598960}$, the number of flushes over the num-



1.2 Handling numbers in various formats

ber of poker hands which has decimal form 0.0019807923169268. What's the probability of being dealt a straight? Right, $\frac{10240}{2598960}$ or 0.0039400375534829. The probability of any being dealt any kind of poker hand is a fraction with denominator 2598960 for the same reason. If we later need to do arithmetic with these probabilities, we can just work with the numerators since the denominators are all the same. It's a lot easier to work with the numbers 3744, 5148 and 10240 than with the decimals

The second kind of loss is a loss of accuracy. Whole-number-ratio fractions are *exact* representations of an answer. What we get when we convert to decimal form is just an *approximation* to this answer except in rather special cases. This loss of accuracy is usually harmless but it can become important, especially if, as is all too common, intermediate answers are rounded to just a few places.

Rounding

This brings us to one of the simplest, but also one of the most important, rules in the course.

FIRST RULE OF ROUNDING 1.2.4: *Don't!*

More precisely, never round an intermediate answer.

The **SECOND RULE OF ROUNDING 1.2.6** tells us how to round *final* answers, but this is much less important. Before, we turn to it, here are a few examples to illustrate why heeding the **FIRST RULE OF ROUNDING 1.2.4** is so important.

PROBLEM 1.2.5: Carry out each calculation below in four ways:

- keeping the answer in fractional form at all times, then converting the final answer to decimal form.
- converting the fraction on the left to decimal form and then completing the calculation using the all the decimals that your calculator gives.



1.2 Handling numbers in various formats

- c. converting the fraction on the left to decimal form and then completing the calculation using only the first 6 decimals that your calculator gives.
- d. converting the fraction on the left to decimal form and then completing the calculation using only the first 4 decimals that your calculator gives.
- i) $67500 - \left(\frac{285}{7}\right)^3$

Solution

- a. To keep the calculation in fractional form is a bit of work. First the cube is $\frac{285 \cdot 285 \cdot 285}{7 \cdot 7 \cdot 7} = \frac{23149125}{343}$. Then to do the subtraction, we need to put everything over the common denominator 343 getting $\frac{67500 \cdot 343 - 23149125}{343} = \frac{3375}{343}$. In decimal form, this is 9.839650146.
- b. I kept everything here to 14 places to simulate a typical graphing calculator. First, $\frac{285}{7} = 40.714285714286$, then we cube this to get 67490.160349856, and finally

$$67500 - 67490.160349856 = 9.839650144.$$

Notice that even though we kept our intermediate answers to 14 places, our final answer was off by 2 in the 14th place. This kind of rounding error (it wasn't we who rounded but our calculator, behind our backs so to speak) is usually not important. But not always, as we'll see later in [PROBLEM 5.3.8](#). And we have no way of knowing how badly we've been bitten by an error of this type other than to check against an answer found in some other way (above, by using fractions rather than decimals).

- c. Now I round $\frac{285}{7} = 40.714285714286$ to 6 places getting 40.7143. That seems like lots of places in this problem. Cubing this gives 67490.231392153 and finally

$$67500 - 67490.231392153 = 9.768607847.$$

Now our answer is off by about 0.07, meaning we can only expect 1 or 2 place accuracy, even though we kept 6 places. This much



1.2 Handling numbers in various formats

error will frequently be enough to render an answer useless (or even downright misleading) in an application.

d. Now I round $\frac{285}{7} = 40.714285714286$ to 4 places getting 40.71. Cubing this gives 67468.849911 and an answer of

$$67500 - 67468.849911 = 31.150089.$$

This answer is total garbage, but there's again no way to tell because all those (incorrect) decimals make it look pretty good.

ii) $\left(\frac{10000}{1111} - 9\right) * 102.478$

All approximate calculations run the risk of losing accuracy. That's what happened in part [i\)b](#) above. There's an entire branch of mathematics, numerical analysis, devoted to understanding and eliminating or at least controlling the sources of such errors. That said, I have tried hard to ensure that the questions I ask are ones your calculator will be able to answer accurately.

Let me emphasize, however, that it was the decision to round an intermediate result—not our calculator—that was solely responsible caused the much more serious inaccuracies in the answers in [i\)c](#) and [i\)d](#). In those parts, the calculator was keeping 14 places as best it could. But once we'd rounded the intermediate answer most of those places became useless. Rounding essentially hits your calculator on the back of the head with a frying pan, dazing it for the rest of the calculation. Notice also how rounding aggravated the tendency to leak accuracy during an approximate calculation. Instead of losing 1 place as happened when we did not round, we lost 4 or 5 when we did.

In some sections of the course, in particular, throughout the chapter [CHAPTER 5](#), failing to heed the [FIRST RULE OF ROUNDING 1.2.4](#) and rounding intermediate answers is one of the primary causes of incorrect final answers.

What about those final answers? Why not just let the [FIRST RULE OF ROUNDING 1.2.4](#) cover them and leave them unrounded? The issue

1.2 Handling numbers in various formats

here is not one of mathematical accuracy, but rather one of providing accurate information to someone reading your answer. There's a universally agreed convention about how this should be done.

SECOND RULE OF ROUNDING 1.2.6: *In a final answer, give all significant digits of whose correctness you are reasonably confident, but no more.*

In other words, readers are entitled to trust the accuracy of all the significant digits you include in a final answer. And they're entitled to assume that you are *not* sure of any further digits. The final answer you show should tell the reader exactly how accurate the answer being offered is.

Remember that to be sure of the number in one place you need to know that the actual error is less than 5 units in the next place. For example, when I write that the products of the reaction weighed 1.34kg, I am implying that I measured them as between 1.335kg and 1.345kg. When I estimate that there are 2.70 million students in parochial schools in the United States, I am implying that my data assures me that there are at least 2695000 such students and at most 2705000. Notice also that, for this reason, I'm saying something different if I say there are 2.7 million students in parochial schools. Even though both "expand" to the number 2700000 the latter figure has a bigger margin of error—there are between 2650000 and 2750000 students. By writing 2.7 million, I have warned you that my information about the students is less accurate.

Most violations of this rule involve giving more digits in an answer than you really know. This is a pretty cheap lie, but unfortunately it often achieves its goal, that of making the reader think that you have a more accurate understanding than you really do. You should always be on the lookout for this kind of dishonesty in the media, and sadly, you'll often see it even in scholarly publications.

It can be easy to catch. When a pollster announces that 37.8% of Americans favor the death penalty, it's pretty safe to be a bit skept-

1.2 Handling numbers in various formats

tical. To keep the cost reasonable, most polls are based on samples of at most a few thousand heads. As we'll see when we look at statistical tests later on, with such a there's typically an uncertainty of about 2% either way, even if the survey methodology is perfect. In actual surveys, accuracy of $\pm 4\%$ is considered good.

But catching such lies can also be hard. Using unjustified extra places makes any author seem to know more. The hope is often to snow you into accepting a weak argument with a flurry of decimals. Your only defense is to read the work in detail to assess both whether the arguments are sound and whether the data justifies basing such accurate conclusions on them.

EXAMPLE 1.2.7: In each of the calculations below, find an unrounded answer. Then give a suitable rounded final answer.

- i) If a rod that is 3.3 feet long is made of material that weighs 1.2 pounds per foot, how much does the rod weigh?

Solution

Finding the answer is easy: $3.3 \times 1.2 = 3.96$. What's not so clear is how to round this. Why not just leave the answer as 3.96 pounds? Well, notice that both the ingredients that went into the calculation were given to us to 2 place accuracy. If, as careful readers, we believe that accuracy, then we should expect our answer to be accurate to two places. So we should round the answer as 4.0 pounds.

The rule here is one worth noting although we won't need to formalize it. Answers usually have at most the accuracy of their least accurate component. That "at most" is really a necessary proviso. In our example, the worst cases for our calculation are $3.25 \times 1.15 = 3.7375$ and $3.35 \times 1.25 = 4.1875$. So we really *know* only that the weight is 4 pounds to the nearest pound. These are the kind of things those numerical analysts get paid to worry about. I won't ask you to take this much care.

- ii) If 1 gallon of paint will cover 220 square feet of wall, how many



1.3 Sums and series

gallons will I need to paint a house with 2140 square feet of wall space?

Solution

Again, finding the answer is easy: $\frac{2140}{220} = 9.7272727272727$. Here the rounding is dictated by even more common sense considerations. We don't have much choice because it's hard to know how accurate to think those areas are. It "smells" like we have 2 place accuracy in the figure 220 (although 3 is possible) and 3 in the figure 2140. If so, we can expect 2 figure accuracy in the answer so we'll need 9.7 gallons. But since paint is sold in whole gallons, all this is moot: we'll need 10. Note that I could have helped the reader by using scientific notations. If I had written 2.2×10^2 and 2.14×10^3 , I'd have made it clear that these were 2 and 3 place figures.

Probably the best answer to this problem, is the common sense one, "You'll need a bit less than 10 gallons of paint".

PROBLEM 1.2.8: In each of the calculations below, find an unrounded answer. Then give a suitable rounded final answer.

- If 2.34×10^{16} atoms of hydrogen react with 1×10^{16} atoms of oxygen, how many molecules of water (H_2O) will be produced?
- If 46% of American voters are Republicans and 52% of Republicans are women, what percent of American voters are Republican women?

1.3 Sums and series

Summation notation

One type of computation that comes up in many contexts is that of adding up or *summing* a sequence of **terms**—values obtained from



1.3 Sums and series

a single formula by plugging in a range of different values of the variable. The range may be finite in which case we have a summation or infinite—these are usually called series. The variable in such a sequence of terms is called an **index of summation** and the lowest and highest values it takes on in the sum are called the lower and upper limits of summation. Rather than write out all the terms (boring when there are a large number and impossible when there are an infinite number), we save time and energy by using a shorthand summation notation for such sums.

SUMMATION NOTATION 1.3.1: *The summation notation $\sum_{i=l}^u f(i)$ is a shorthand for the sum for the formula $f(i)$, starting with lower limit $i = l$ and increasing i successively by 1 until we reach the upper limit $i = u$ —or forever, if $u = \infty$. That is,*

$$\sum_{i=l}^u f(i) = f(l) + f(l+1) + f(l+2) + \cdots + f(u-1) + f(u),$$

and

$$\sum_{i=l}^{\infty} f(i) = f(l) + f(l+1) + f(l+2) + f(l+3) + \cdots.$$

EXAMPLE 1.3.2: Since all the sums we'll need for here have lower index l equal to either 1 or 0, these are the only kind of examples I'll give now. Later, when we deal with expected values in probability, we'll look at more general cases.

i) (Arithmetic summations) Sums in which the formula $f(i)$ is linear—a constant times i plus another constant—are called **arithmetic**. As a first example, consider a case when $f(i) = i$:

$$\sum_{i=1}^{10} i = 1 + 2 + 3 + 4 + 5 + 6 + 7 + 8 + 9 + 10 = 55.$$

As a second, with $f(i) = 3i - 2$, consider $\sum_{i=0}^5 3i - 2 = (3 \cdot 0 - 2) + (3 \cdot 1 - 2) + (3 \cdot 2 - 2) + (3 \cdot 3 - 2) + (3 \cdot 4 - 2) + (3 \cdot 5 - 2) = -2 + 1 + 4 + 7 + 10 + 13 = 33$. Note that there's no real need to mention the function f : we just write the formula for the terms to

the right of the summation sign \sum . Infinite arithmetic series almost always total infinity so are seldom seen.

ii) (Finite geometric summations) Sums in which the formula $f(i)$ is exponential—a constant times the i^{th} power of a fixed base r called, for historic reasons, the **ratio**—are called **geometric**. You might think *power summation* might be a more mnemonic term, but, while this is sometimes used, the term geometric is the standard one. As a first example, consider $\sum_{i=0}^6 3^i = 3^0 + 3^1 + 3^2 + 3^3 + 3^4 + 3^5 + 3^6 = 1 + 3 + 9 + 27 + 81 + 243 + 729 = 1093$. As a second $\sum_{i=1}^5 \left(\frac{1}{2}\right)^i = \left(\frac{1}{2}\right)^1 + \left(\frac{1}{2}\right)^2 + \left(\frac{1}{2}\right)^3 + \left(\frac{1}{2}\right)^4 + \left(\frac{1}{2}\right)^5 = \frac{1}{2} + \frac{1}{4} + \frac{1}{8} + \frac{1}{16} + \frac{1}{32} = \frac{31}{32}$.

iii) (Infinite geometric series) Infinite geometric series may or may not give a finite total. The series $\sum_{i=1}^{\infty} 3^i$ does not because, as you can see from the summation above, the individual terms get bigger and bigger, and so does their sum. But the terms in the series $\sum_{i=1}^{\infty} \left(\frac{1}{2}\right)^i$ are getting small quite quickly. In a moment, we'll have a formula that will tell us that this series sums to exactly 1 but it's not hard to guess this answer. Notice that the sum $\frac{31}{32}$ of the first 5 terms in ii) can be rewritten $1 - \frac{1}{32}$. The next term after those summed in ii) is $\frac{1}{64}$ corresponding to $i = 6$, and if we add this to the summation we get $\frac{31}{32} + \frac{1}{64} = \frac{63}{64} = 1 - \frac{1}{64}$. You can check that the next term is $\frac{1}{128}$ and that adding it gives the summation $1 - \frac{1}{128}$. Each summation is less than 1 by the last term added, and since these terms go to 0 when i goes to ∞ , the summations go to 1.

When mathematicians are faced with making long computations they are usually just as unhappy as most of you. The difference is that, instead of giving up or plowing away, a mathematician's reaction is to ask: "If I think hard and come up with the right clever idea, can't I find some way to get the answer to this calculation without doing all the work?" The power of mathematics is that when you do think hard the answer is often, if not always, "Yes!". In computing summations and series, the way to get an answer without all the arithmetic is to find a summation formula, that is, a formula for the total of all the

1.3 Sums and series

terms that eliminates the index of summation and gives an answer only involving the *limits* of summation. Such a formula is also often called **closed form** for the summation.

Finding a closed form formula usually calls for more machinery—mostly from calculus—than we want to include here. Fortunately, geometric summations and series are the only kind that comes up frequently in this course and a little algebra is all it takes to find a closed form for these.

Geometric summation and series formulae

Let's start with finite geometric summations. In this case, the clever idea is an easy one based on difference of squares factorization, $(1 + r)(1 - r) = 1 - r^2$. Why aren't there more terms? If we expand out, we see that $(1 + r)(1 - r) = 1(1 - r) + r(1 - r) = 1 - r + r - r^2 = 1 - r^2$. The two r 's cancel because of the opposite signs. The clever idea has two parts. First, let's think of $(1 + r)$ as $r^0 + r^1$ as we did above and view the factorization as

$$(r^0 + r^1)(1 - r) = 1 - r^2.$$

Second, let's ask what happens when we add an r^2 . We find

$$\begin{aligned} & (r^0 + r^1 + r^2)(1 - r) \\ &= r^0(1 - r) + r^1(1 - r) + r^2(1 - r) \\ &= r^0 - r^1 + r^1 - r^2 + r^2 - r^3 \\ &= 1 - r^3. \end{aligned}$$

There are 2 cancellations this time and we still have only two terms in the final right hand side. Adding a r^3 , we find

$$\begin{aligned} & (r^0 + r^1 + r^2 + r^3)(1 - r) \\ &= r^0(1 - r) + r^1(1 - r) + r^2(1 - r) + r^3(1 - r) \\ &= r^0 - r^1 + r^1 - r^2 + r^2 - r^3 + r^3 - r^4 \\ &= 1 - r^4. \end{aligned}$$

1.3 Sums and series

This time there are 3 cancellations this time and we still have only two terms in the final right hand side. Once again, a pattern is clear. The only thing that changes when we add another power is that the exponent of r on the final right hand side goes up by 1. If we go up to the u^{th} power what we find is:

$$\begin{aligned}\sum_{i=0}^u r^i &= (r^0 + r^1 + \dots + r^{u-1} + r^u)(1 - r) \\ &= r^0(1 - r) + r^1(1 - r) + \dots + r^{u-1}(1 - r) + r^u(1 - r) \\ &= r^0 - r^1 + r^1 - r^2 + \dots + r^{u-1} - r^u + r^u - r^{u+1} \\ &= 1 - r^{u+1}.\end{aligned}$$

This is the basis of the **GEOMETRIC SUMMATION FORMULA 1.3.3** which I hope you have already seen. To get the formula, we just divide both sides by $(1 - r)$ leaving just the series on the left side.

GEOMETRIC SUMMATION FORMULA 1.3.3: *For any integer $u \geq 0$,*

$$\sum_{i=0}^u r^i = r^0 + r^1 + \dots + r^{u-1} + r^u = \frac{1 - r^{u+1}}{1 - r}.$$

PROBLEM 1.3.4: Use the **GEOMETRIC SUMMATION FORMULA 1.3.3** to evaluate the sums $\sum_{i=0}^6 3^i$ and $\sum_{i=1}^5 \left(\frac{1}{2}\right)^i$ from ii) of **EXAMPLE 1.3.2**.

Partial Solution: I'll do the second sum. There's one small point to overcome here: this sum has lower limit 1 and the **GEOMETRIC SUMMATION FORMULA 1.3.3** gives a closed form for sums with lower limit 0. There are several potential solutions. Perhaps the first that comes to mind is to find closed form when the lower limit is 1. What's wrong with this approach? If we try to write down a new formula every time we have a slightly different problem, pretty soon we'll have as many formulae as problems and it'll be impossible to remember them all. A key art in mathematics is to find a way to understand lots of problems with a very few ideas: the possibility of doing this is what makes mathematics such a powerful tool.

Here's an easy solution that has this flavor. Instead of adjusting the formula to fit the sum, adjust the sum to fit the formula. We do know



1.3 Sums and series

what $\sum_{i=0}^5 \left(\frac{1}{2}\right)^i$ is. The **GEOMETRIC SUMMATION FORMULA 1.3.3** gives $\frac{1-r^{n+1}}{1-r} = \frac{1-\frac{1}{2}^{5+1}}{1-\frac{1}{2}} = \frac{\frac{63}{64}}{\frac{1}{2}} = \frac{63}{32}$. But this sum differs from $\sum_{i=1}^5 \left(\frac{1}{2}\right)^i$ only by the 0th term $\left(\frac{1}{2}\right)^0 = 1$. Subtracting this term, we get the total $\frac{31}{32}$ found in **EXAMPLE 1.3.2**.

PROBLEM 1.3.5: Use the **GEOMETRIC SUMMATION FORMULA 1.3.3** to evaluate the sums by relating each to a summation covered by the formula:

- i) $\sum_{i=2}^7 2^i$
- ii) $\sum_{i=0}^6 5 \cdot 10^i$

Hint: Since *every* term in the last sum has the common factor 5 in it, this is the same as $5 \sum_{i=0}^6 10^i$. This idea comes up a often so make note of it.

Finding closed forms for series is usually even harder than finding them for summations. For one thing, it's often tricky just to decide whether it makes sense to total an infinite summation or series—if so, we say the series converges; if not, it diverges. The terms had better approach 0 or the sum can never settle down but when the terms do approach 0, the total may or may not make sense. Consider, for example, the series

$$\sum_{i=1}^{\infty} \frac{1}{i} = \frac{1}{1} + \frac{1}{2} + \frac{1}{3} + \cdots + \frac{1}{i} + \cdots$$

This series, so famous it has a name—the harmonic series—is divergent (cannot be summed). More precisely, if you add *enough* terms, you can make the sum as big as you like. But this sum gets big v-e-e-r-y slowly, so slowly that, even with a computer, you can't add up *enough* terms to make it *clear* that the sum really does go off to infinity. For example, the sum of the first 1,000,000,000 terms is only about 21.30. Even when you can sum a series strange things can happen. The alternating harmonic series

$$\sum_{i=1}^{\infty} \frac{(-1)^{i-1}}{i} = \frac{1}{1} - \frac{1}{2} + \frac{1}{3} - \frac{1}{4} + \cdots$$

1.3 Sums and series

converged to the sum $\ln(2) \simeq 0.693 \dots$. But, you can also make this converge to *any number you like* if you're allowed to choose the order in which you add up the terms. If you don't believe this, try asking your instructor to rearrange it to sum to your Social Security number.

Geometric series $\sum_{i=0}^{\infty} r^i$ are one important exception. The size of the ratio r determines, in an easy way, whether or not the series converges or diverges. If $|r| \geq 1$ then $|r^i| \geq 1$ for any i . In other words, the terms stay big and the summations never settle down. If $|r| < 1$, the $|r^i|$ goes to 0 as i goes to ∞ . In fact, the terms r^i go to 0 quite fast and the series converges. We can see this more directly, and get a formula for the sum of such series by looking at what happens to the summations in the [GEOMETRIC SUMMATION FORMULA 1.3.3](#) for $\sum_{i=0}^u r^i$ when the upper limit u goes to ∞ . The only place u appears in the formula $\frac{1-r^{u+1}}{1-r}$ is in the power r^{u+1} . If $|r| < 1$, then as u goes to ∞ , this power goes to 0. The resulting formula for geometric series is thus, surprisingly, even simpler than for geometric summations.

GEOMETRIC SERIES FORMULA 1.3.6: *If $|r| < 1$, then the sum of the geometric series with ratio r is given by the formula $\sum_{i=0}^{\infty} r^i = \frac{1}{1-r}$.*

EXAMPLE 1.3.7: We can now check the guess we made in [EXAMPLE 1.3.2](#), that $\sum_{i=1}^{\infty} \left(\frac{1}{2}\right)^i = 1$. We use the same trick as in [ii\) of PROBLEM 1.3.4](#). The series $\sum_{i=0}^{\infty} \left(\frac{1}{2}\right)^i = \frac{1}{1-\frac{1}{2}} = \frac{1}{\frac{1}{2}} = 2$ and the series $\sum_{i=1}^{\infty} \left(\frac{1}{2}\right)^i$ differs from it by the term $\left(\frac{1}{2}\right)^0 = 1$.

PROBLEM 1.3.8: Use the geometric series formula to find the sums:

- i) $\sum_{i=0}^{\infty} \left(\frac{5}{11}\right)^i$
- ii) $\sum_{i=0}^{\infty} \left(\frac{5}{36}\right)^2 \left(\frac{25}{36}\right)^i$ Hint: Use the same idea as in [PROBLEM 1.3.5](#).
- iii) $\sum_{i=0}^{\infty} \left(\frac{4}{36}\right)^2 \left(\frac{26}{36}\right)^i$

The last two sums will come up when we look at the odds of winning at the game of craps in [ANALYSING THE GAME OF CRAPS](#).

1.4 Logarithms and exponentials

CHALLENGE 1.3.9: Use the [GEOMETRIC SUMMATION FORMULA 1.3.3](#) to evaluate the sum $s = \sum_{i=0}^{\infty} i \cdot \left(\frac{1}{2}\right)^i$.

Hint: This is a challenge because that factor of i in the terms changes with each term so there's no way to factor it out as in ii) of [PROBLEM 1.3.5](#). Write $\frac{1}{2}s$ as a series using the series for s . Then take the difference $s - \frac{1}{2}s$ as a difference of series. To see why this helps, try writing out the first few terms of each without computing the values of the powers of $\frac{1}{2}$.

There's much more to be said, whole theories of summations and series. Entire books have been devoted to such subjects as finding formulae for different types of series, and finding ways to squeeze sums so that they fit formulae. Fortunately, you won't need very much of this technology. We'll say a bit more about summations when we discuss probabilities and I'll wait until then to introduce the further ideas we'll need.

1.4 Logarithms and exponentials

Don't look down!

First, an explanation of the title. Watch the segment of [Zoom and Bored](#)—a classic Road Runner™ cartoon from 1957—that runs from about 0:45 seconds in 1:20 in. From about 0:55 seconds in, Wile E. Coyote, is suspended in mid-air after having run off the edge of a cliff while in a cloud of dust. He remains suspended even after he suspects that he's in mid-air and starts feeling with his paw for the ground. It's only when he looks down and, as the dust clears, *knows* that he's in mid-air that, at about the 1:15 mark, he actually falls.

In this section, I'd like to run you off a cliff with a simple question. Then we'll feel around for the ground for a while while I convince

1.4 Logarithms and exponentials

you we're suspended in midair. Don't worry, you'll be fine. Just don't look down! Then in the remainder of this section I'll walk you back onto terra firma and, in the process, we'll learn the basic facts about logarithms and exponentials that we're going to need in the rest of the course.

Surprisingly, the path back depends, not on a lot of algebra, but on a few easy pictures. The story is one of the most elegant in basic mathematics, usually reserved for advanced calculus courses, but, as we'll see, it doesn't even require any algebra: you just have to draw the *right* few easy pictures. Eventually we'll find out the simple answer to the simple question. Let's begin with the question.

SIMPLE QUESTION 1.4.1: *What is $10^{\sqrt{2}}$?*

If I put this on a midterm, I expect that most of you would just reach for your calculators, type in $10^{\wedge} 2\text{nd } x^2 2 \text{) ENTER}$ —at least on a standard TI-8x—get back 25.95455351947, or however many of the digits of that decimal your calculator gives you, and go on to the next question. From one point of view—that of giving a good decimal approximation to the number $10^{\sqrt{2}}$, this answers my question.

But it does *not* answer it in the sense in which I want to pose it: What number are we talking about when we write down the expression $10^{\sqrt{2}}$? To get a feel for the issue, let's ask

EVEN SIMPLER QUESTION 1.4.2: *What is $\sqrt{2}$?*

This question is simpler for a couple of reasons. Once again, I'm not interested in answers like 1.4142135623731 that your calculator might give you. But now it *is* easy to answer the question as I intend it: What number are we talking about when we write down the expression $\sqrt{2}$? By $\sqrt{2}$, we mean *the* positive number whose square is 2. The number $-\sqrt{2}$ also has square 2 so how do I know I can write that italic *the*: How do I know there's only *one* positive number whose square is 2? Well, because if $a < b$ then $a^2 < b^2$. If x is any positive number Then either $x < \sqrt{2}$ and then $x^2 < 2$ or $x > \sqrt{2}$ and then $x^2 > 2$.

1.4 Logarithms and exponentials

If you're really alert, you may be asking a less simple question. How do we know that there's *any* such number? Why couldn't every positive number either have square less than 2 or square greater than 2? For the number $\sqrt{2}$, I can sleaze out by appealing to the

THE PYTHAGOREAN THEOREM 1.4.3: *If c is the length of the hypotenuse of a right triangle and a and b and c are the lengths of the other two sides, then $a^2 + b^2 = c^2$.*




FIGURE 1.4.4: A right triangle with sides 1, 1, and $\sqrt{2}$.

So, in [FIGURE 1.4.4](#), where $a = b = 1$, $c^2 = 1^2 + 1^2 = 2$ —so $c = \sqrt{2}$. This is the first, but not the last time a picture will come to our rescue.

Why did I say that I was sleazing out with this explanation? Well, what if a asked

NOT QUITE SO SIMPLE, QUESTION 1.4.5: *What is $\sqrt[3]{2}$?*

PROBLEM 1.4.6: Show that there is at most one “ $\sqrt[3]{2}$ ”—that is, at most one number whose cube equals 2. Hint: Just replace the squares with cubes in the argument that there's only one $\sqrt{2}$.

Now, however, Pythagorus deserts us when we ask why such a number has to exist. The reason one *does* exist intuitively easy: “The real number line \mathbb{R} has no holes”. Let's take this on faith, and not much faith is needed, since our mental picture of the real number line as the x -axis in the plane is that it's a barrier dividing the (x, y) -plane into upper and lower halves. You can't go from below the x -axis to above it without either crossing it or “jumping” over it.




FIGURE 1.4.7: Graph of $x^3 - 2$ showing the root at $\sqrt[3]{2}$

Look at the picture above of the part of the graph of the function $x^3 - 2$ for x between 1 and 2. The graph starts below the x -axis at $(1, -1)$ and ends above it at $(2, 6)$. It doesn't jump over the axis anywhere so there must be a point c (close to 1.25992104989487 if you must know) where the graph crosses. At this point, $c^3 - 2 = 0$ so $c^3 = 2$ and $c = \sqrt[3]{2}$.

That's all completely clear, isn't it? Well, yes and no. Everything I have said is intuitively obvious—clear to our geometric intuition. Mathematicians soon realized, however, that some further argument was required to justify the claim that a function like $x^3 - 2$ has no jumps. These arguments are straightforward—they're mentioned (if not actually taught) in the first weeks of any calculus course—and I'll just take them for granted here.

What's astonishing is that 200 years went by before any mathematician said, "Wait a minute. How do we really know there are no holes in the real line? Why couldn't there be a infinitesimal pinhole, too small to see even with an electron microscope, right at $\sqrt[3]{2}$?" Wile E. Coyote could have told them what a big mistake asking this was: they looked down! Well, once they had asked the question and looked down, it turned out to be quite tricky to get back to the edge of the

1.4 Logarithms and exponentials

cliff. Only math majors ever learn how, and not all of them. You, however, can relax: there *aren't* any holes in the real line. And we're going to be satisfied with our intuition that this is so, without worrying about any of the technicalities needed to confirm this intuition.

Whew! OK, let's recap. We know what number $\sqrt{2}$ is: it's the unique positive number whose square is 2 (and we know that such a number does exist). Why is this answer better than the answer 1.4142135623731? Because the answer 1.4142135623731 is *wrong*! You can check that

$$1.4142135623731^2 = 2.000000000000001400410360361.$$

So 1.4142135623731 may be *close* to the number whose square is 2 but it is *not* the number whose square is 2: it's the number whose square is 2.000000000000001400410360361.

Picky, picky, picky. No! Not at all. There's a right answer—one and only one right answer—to the [EVEN SIMPLER QUESTION 1.4.2](#): $\sqrt{2}$ is the number whose square is 2 and that number is not 1.4142135623731.

OK, we were off by a 1 in the 14th decimal place. Probably all we need to do is compute $\sqrt{2}$ to a few more places. Not so! The number 1.4142135623731 is a rational number. It's the ratio of whole numbers $\frac{14142135623731}{10000000000000}$. Any more accurate decimal approximation to $\sqrt{2}$ like say 1.41421356237309504880168872 is also a rational number. You just multiply above and below by a power of 10 that clears the decimal point: in this case, $\frac{141421356237309504880168872}{10000000000000000000000000}$. And, no matter *how* you do this, you won't get a number whose square is 2 because

$\sqrt{2}$ IS IRRATIONAL 1.4.8:

It must be Greek day today because this insight goes back to Euclid. What does saying that $\sqrt{2}$ is **irrational** mean? Just what I claimed above. No rational number—that is, *no* number that can be written as a fraction $\frac{a}{b}$ with a and b whole numbers



1.4 Logarithms and exponentials

can be the number $\sqrt{2}$. Not just that the better approximation 1.41421356237309504880168872 is not $\sqrt{2}$. Though let's note that it isn't:

[illegible]

No, it means that no matter what a and b we try we'll *never* have $\left(\frac{a}{b}\right)^2 = 2$. Why not? Computing with examples will show nothing. The fact that the two rational numbers above do not square to 2 tell us nothing about what might happen with the next rational number we try. No matter how many examples we try we can never be sure that we haven't simply failed to check the *right* rational number amongst the infinitely many possibilities.

Actually, Euclid's argument is much simpler, depending on nothing more than the fact that a whole number cannot be both even and odd. More precisely, we can only have a solution of $\left(\frac{a}{b}\right)^2 = 2$ if there's a whole number that's both even and odd, and since we can't have the latter, we can't have the former either.

Here's how it goes. If a and b are both even we can cancel a factor 2 from both, reducing the size of both, without changing the ratio $\frac{a}{b}$ —and hence preserving the equation $\left(\frac{a}{b}\right)^2 = 2$. If they're still both even, we can cancel another 2. Since each time we do this a and b get smaller, we must eventually wind up with one, or perhaps both, odd.

Now multiply the equation $\left(\frac{a}{b}\right)^2 = 2$ by b^2 getting $a^2 = 2b^2$. The right side of this equation is even, hence so is the left side. But the square of an odd number is odd, so a must be even. This means two things. First, since we know a and b are not *both* even, b must be odd. Second, since a is even, we can write $a = 2c$ with c another whole number.

Plugging in $a = 2c$ our equation becomes $4c^2 = 2b^2$, or cancelling a 2, $2c^2 = b^2$. Now the left side is even (because of the 2) and the right side is odd (because b is odd and hence so is b^2). Our rational form



1.4 Logarithms and exponentials

for $\sqrt{2}$ has led to the impossible whole number that is both even and odd.

OK, so $\sqrt{2}$ is irrational: so what? So the *only* way we have of answering the **EVEN SIMPLER QUESTION 1.4.2** is the way I answered it above: $\sqrt{2}$ is the unique positive number whose square is 2. Because we can't nail $\sqrt{2}$ with any rational number, we cannot nail it with any decimal. The *only* way to describe it is by what it does—turn into 2 when squared.

Now we can start to think about how we might answer the **SIMPLE QUESTION 1.4.1**. If we're lucky, $10^{\sqrt{2}}$ might be rational. No such luck. But we can still hope to pin down $10^{\sqrt{2}}$ as we did $\sqrt{2}$, by describing what it (and only it) does. To get a feel for how we might do this, let's ask

SOME REALLY SIMPLE QUESTIONS 1.4.9: *What are*

- i) 10^4 ,
- ii) $10^{\frac{1}{3}}$, and
- iii) $10^{\frac{4}{3}}$?

Answering these involves nothing more than recalling what we mean by an exponential like 10^x . Such an exponential tells us to multiply the base—in this case the 10—by itself x times. For example, we can say immediately that $10^4 = 10 \cdot 10 \cdot 10 \cdot 10 = 10000$. Let's note right away that the rules of exponents all follow from this. For example, $10^x \cdot 10^y$ tell us to multiply 10 by itself x times and y times, or $x + y$ times in all. But so does 10^{x+y} . So $10^x \cdot 10^y = 10^{x+y}$; likewise, $10^x \cdot 10^y \cdot 10^z = 10^{x+y+z}$ and so on.

We need to pause for a moment over $10^{\frac{1}{3}}$. How do I multiply 10 by itself $\left(\frac{1}{3}\right)^{\text{rd}}$ of a time? But if we remember that $\sqrt[3]{2} = 2^{\frac{1}{3}}$, then we find the answer to this question above. For $10^{\frac{1}{3}} \cdot 10^{\frac{1}{3}} \cdot 10^{\frac{1}{3}} = 10^{\frac{1}{3} + \frac{1}{3} + \frac{1}{3}} = 10^1 = 10$. Thus $10^{\frac{1}{3}}$ is the number whose cube is 10 just as $\sqrt[3]{2}$ was the number whose cube was 2. Even the argument that there's only one such number is the same.



1.4 Logarithms and exponentials

We can see another familiar rule of exponents in this last calculation $(10^x)^z = 10^{x \cdot z}$: there are x 10s in each 10^x and z 10^x s in $(10^x)^z$ for a total of $x \cdot z$ 10s in all: this is just $10^{x \cdot z}$. Turning this around lets us identify $10^{\frac{4}{3}}$ as $(10^{\frac{1}{3}})^4$.

There is, of course nothing special about 3 and 4. If we understand x^{th} powers 10^x and y^{th} roots $10^{\frac{1}{y}}$, we understand all rational powers $10^{\frac{x}{y}}$ since $10^{\frac{x}{y}} = (10^{\frac{1}{y}})^x$. In fact, there's nothing special about 10 either. Let's just record the properties we have checked.

RULES OF EXPONENTS 1.4.10: *If b is any positive base and x, y and z are any rational exponents, then the following rules of exponentials hold*

- i) $b^x \cdot b^y = b^{x+y}$ and $\frac{b^u}{b^v} = b^{u-v}$.
- ii) $(b^x)^z = b^{x \cdot z}$, and

PROBLEM 1.4.11: OK, so I didn't check that $\frac{b^u}{b^v} = b^{u-v}$. You check it. Hint: If you multiply both sides by b^v , you get b^u on the left side. To see what you get on the right, try setting $x = u - v$ and $y = v$ in RULES OF EXPONENTS 1.4.10.i).

Fine, so now we have a head of steam where does that leave us with the SIMPLE QUESTION 1.4.1? In a cloud of dust, standing on nothing more than thin air, having just run off the side of a cliff. We understand all rational powers $10^{\frac{a}{b}}$ perfectly. But those are the *only* powers 10^x where there's any hope of "multiplying the base by itself x times". All the rules involve powers in which the exponents are related by arithmetic operations like addition and multiplication. From rational pieces, those recipes can only turn out other rational exponents. Since the number $\sqrt{2}$ is *not* rational, we have no way to make sense out of $10^{\sqrt{2}}$! No, we're not missing something.

It gets worse. Now let me ask

ANOTHER SIMPLE QUESTION 1.4.12: *What is the meaning of the function $\log_{10}(x)$?*



1.4 Logarithms and exponentials

I hope you remember that the standard answer is that $\log_{10}(x)$ is the number y such that $10^y = x$. What does that mean if y is not rational? We no longer know since we've realized that we standing on thin air with exponentials having irrational exponents.

In [INFINITIES AND AN ARGUMENT FROM *The Book*](#), we'll see that most real numbers—almost all, in a precise sense—are not rational. The upshot is that we know what a few logarithms and exponentials mean—those where we get (very) lucky and y is rational—but have no idea what most of them are.

The dust has cleared and we are standing on thin air. It's OK. Take a deep breath. Whatever you do, don't look down until you've read the rest of this section! Once you do you'll be safe.

The natural logarithm

Before we can answer the [SIMPLE QUESTION 1.4.1](#), we'll need to understand logarithms. And since we now realize that we don't understand what most exponentials mean, we need to find some other way to think about logarithms. Our solution applies the same technique we used to understand [EVEN SIMPLER QUESTION 1.4.2](#), “What is $\sqrt{2}$?”: draw a picture that represents this number. The pictures we'll eventually draw to understand logarithms involve areas rather than lengths but are in some sense even simpler; we don't need any fancy hardware like [THE PYTHAGOREAN THEOREM 1.4.3](#). In the final subsection, we'll make sense of exponentials by simply reflecting the graph of a logarithm function. To get warmed up, let's go back to the first applications of logarithms. These had little to do with their relation to exponentials.

Logarithms were probably discovered by a Swiss clockmaker names Joost Bürgi in about 1588. Bürgi was man of wide interests—he dabbled in mathematics and astronomy, and his horological inventions enabled him to greatly improve the accuracy of mechanical clocks,



FIGURE 1.4.13: The First Books on Logarithms

allowing them to be used for the first time for accurate astronomical observations. Constructing these clocks involved making extensive calculations and it was to simplify the multiplications these necessitated that he devised logarithms. As we’ll see in a moment, you can multiply numbers by taking logarithms—until recently tables of values were used to do this—and adding. The great astronomers Tycho Brahe and Johannes Kepler, with whom Bürgi worked, also realized the enormous benefits of logarithms. They speed up calculation so much that the great French mathematician and astronomer, Pierre Simon Laplace is said to have remarked that logarithms, “by shortening the labors, doubled the life of the astronomer”.

Credit for discovering logarithms commonly goes to the Scottish mathematician, physicist and astronomer John Napier who published the first book on logarithms in 1614. His approach was a bit clumsy but Henry Briggs and he devised an improved version which was described in books published by Briggs in 1617 and 1624, the

1.4 Logarithms and exponentials

latter with extensive tables.

After the appearance of Napier's book, Kepler convinced Bürgi to publish his ideas—he appears to have done a considerable amount of the writing as well—in 1620. In 1624, Kepler published his own work on logarithms, giving the first completely transparent explanation of their properties and he included his own 8 place tables of logarithms in his *Rudolphine Tables*. This compilation of the astronomical observations of Brahe and himself accurately predicted the future positions of the planets and, by doing so, was perhaps the biggest factor in gaining general acceptance for the Copernicus' heliocentric theory.

The key property on which all this depends is the following:

LOGARITHM PROPERTY 1.4.14: *A function \ln has the logarithm property and is called a logarithm if the \ln of a product is the sum of the \ln s of the factors:*

$$\ln(a \cdot b) = \ln(a) + \ln(b).$$

In other words, a logarithm is a function that converts products of inputs into sums of outputs.

Let's suppose, for a moment, that we can find such a function and compute its values. We produce two tables, a forward table in which we can look up $\ln(a)$ given a and a reverse table in which we can look up c given $\ln(c)$. The strategy for finding the product $c = a \cdot b$ is then clear. Look up $\ln(a)$ and $\ln(b)$ in our forward table. Find the sum $d = \ln(a) + \ln(b)$. The **LOGARITHM PROPERTY 1.4.14** then says $d = \ln(c)$ so we can find c by looking up d in our reverse table.

EXAMPLE 1.4.15: Let's work two simple examples. I'll just tell you the necessary table values—to 8 places—but soon we'll see how to get more accurate ones from our calculator.

i) First, we multiply 2 times 3. We look up $\ln(2) = 0.69314718$ and $\ln(3) = 1.0986123$ and sum to get 1.7917595. Then we find that the number whose \ln is 1.7917595 is 6.0000002. Basically, OK, but



1.4 Logarithms and exponentials

why is there that 2 at the end. It's because the function $\ln(d)$ grows much more slowly than d . In fact, every d between 5.9999999 and 6.0000004 has 8-place \ln equal to 1.7917595. So we should just be prepared to lose some accuracy in the reverse lookup phase, as we did here.

ii) Next, let's show that $\sqrt{3}$ is about 1.73 by finding $1.73 \cdot 1.73$. We look up $\ln(1.73) = 0.54812141$, add it to itself to get 1.0962428, and perform a reverse lookup, finding that, if $\ln(d) = 1.0962428$, then $d = 2.9928999$. Here we lose no accuracy, since the right answer is $1.73 \cdot 1.73 = 2.9929 \approx 2.99$ to 3 places. This is as close as we can get to $\sqrt{3}$ with a root to 3-places because $1.74 \cdot 1.74 = 3.02$.

In the author's youth, this method was implemented without tables in quaint analog calculating devices called **slide rules**. Instead of tables, a slide rule used ruler-like scales to perform a variety of calculations. For performing multiplications, it's only necessary to look at the D and C scales in the pictures below. When the slides are aligned, the number on the C scale opposite a on the D scale is $\ln(a)$. So to look up logarithms, we just read across from the D scale to the C scale. To go in the other direction, we lookup $\ln(d)$ on the C scale and read across to find d on the D scale. The use of such slides has particularly nice feature: we do not even need to do any addition to multiply two numbers. We just move the 1 on the C scale opposite one of the factors on the D scale—this is 2 in the top picture—and then move down the C scale to the other factor—3 in the top picture. The number on opposite this on the D scale is the product—here 6.




FIGURE 1.4.16: Multiplying 2 by 3 on a slide rule

1.4 Logarithms and exponentials

handle whole numbers m , then Egyptian fractions—those that are of the form $\frac{1}{m}$ —and finally general rational fractions $c = \frac{m}{n}$.

First, $\ln(b^2) = \ln(b \cdot b) = \ln(b) + \ln(b) = 2 \cdot \ln(b)$. Next, $\ln(b^3) = \ln(b^2 \cdot b) = 2 \cdot \ln(b) + \ln(b) = 3 \cdot \ln(b)$. Continuing in this way shows that $\ln(b^m) = m \cdot \ln(b)$ for m a positive whole number. But, $b^{-m} = \frac{1}{b^m}$ and $\ln(\frac{1}{b^m}) = -\ln(b^m)$ by [ii](#)). Putting these together $\ln(b^{-m}) = -\ln(b^m) = -m \ln(b)$ which is [iv](#)) for negative integers.

Using [RULES OF EXPONENTS 1.4.10.ii](#)), $(a^{\frac{1}{m}})^m = a^{\frac{1}{m} \cdot m} = a^1 = a$. Taking \ln of both sides gives and using whole number case from the preceding paragraph, $\ln((a^{\frac{1}{m}})^m) = m \cdot \ln(a^{\frac{1}{m}}) = \ln(a)$. Dividing by m now gives $\ln(a^{\frac{1}{m}}) = \frac{1}{m} \ln(a)$. In other words, part [iv](#)) is correct for inverses of whole numbers too.

Finally, if $c = \frac{m}{n}$, then $b^c = b^{\frac{m}{n}} = (b^{\frac{1}{n}})^m$. So by the “whole number” and “inverse of whole number” cases of [iv](#)):

$$\ln(b^c) = \ln((b^{\frac{1}{n}})^m) = m \cdot \ln(b^{\frac{1}{n}}) = m \cdot \frac{1}{n} \cdot \ln(b) = \frac{m}{n} \cdot \ln(b) = c \cdot \ln(b).$$

This handles rational c and, for the moment, those are the only c we understand how to use as exponents. But [iv](#)) will also apply to general real exponents once we understand what such exponentials mean.

To sum up, the whole package of familiar properties of logarithms really just boils down to the single [LOGARITHM PROPERTY 1.4.14](#): \ln converts multiplication of inputs into addition of outputs. But we’re still no closer to understanding why any function with this key property exists.

In fact, it was half a century before the way we’re now going to show that a logarithm function actually exists was discovered. I’ll explain this part of the history at the end of this section, but, since we define $\ln(a)$ not as a number but as an *area*, our definition is clearly one that calls for a fair bit of hindsight. Not that’s its complicated; quite the contrary.



1.4 Logarithms and exponentials




FIGURE 1.4.20: The area defining $\ln(a)$

AREA DEFINITION OF \ln 1.4.21: For $a > 1$, we define $\ln(a)$ to be the area between the curve $y = \frac{1}{x}$ and the x -axis $y = 0$ and between the vertical lines $x = 1$ and $x = a$, as shown in [FIGURE 1.4.20](#).

For $0 < a < 1$, we define $\ln(a)$ to be the negative of this area.

The function \ln is called the **natural logarithm** and we read the value $\ln(a)$ as “natural log of a ”, or, more concisely, as “ell-en of a ”.

The area in [FIGURE 1.4.20](#) is certainly a perfectly good, if somewhat unusual, way to define the number $\ln(a)$. But when you first see it, it seems crazy. Let me respond to the two most obvious objections to it. First, where are the exponents? Nowhere. Let me emphasize that $\ln(a)$ is an *area* but that this area must be the *right* definition of a logarithm if we can show that it has the [LOGARITHM PROPERTY 1.4.14](#) that $\ln(a \cdot b) = \ln(a) + \ln(b)$. That’s remarkably easy to see from a few simple variations on [FIGURE 1.4.20](#). Let’s stick for the moment to pictures where $a > 1$.

Let’s start with an example, and try to see that $\ln(2) + \ln(3) = \ln(2 \cdot 3) = \ln(6)$. Using the [AREA DEFINITION OF \$\ln\$ 1.4.21](#), this means we are asking the:

KEY QUESTION 1.4.22: Do the areas in [FIGURE 1.4.23](#) and [FIGURE 1.4.24](#) sum to the area in [FIGURE 1.4.25](#)?

The first thing to notice is that the region in [FIGURE 1.4.24](#) whose area equals $\ln(3)$ is the left end of the region in [FIGURE 1.4.25](#) whose area equals $\ln(6)$. This is shown in [FIGURE 1.4.26](#). So if we cancel this common area, we can restate the [KEY QUESTION 1.4.22](#): does the

1.4 Logarithms and exponentials




FIGURE 1.4.23: The area defining $\ln(2)$




FIGURE 1.4.24: The area defining $\ln(3)$




FIGURE 1.4.25: The area defining $\ln(6)$

area on the left of [FIGURE 1.4.27](#) (copied from the picture of $\ln(2)$ in [FIGURE 1.4.23](#)), equal the area on the right of [FIGURE 1.4.27](#) (copied from the right side of the picture of $\ln(6)$ in [FIGURE 1.4.25](#))?




FIGURE 1.4.26: Showing $\ln(3)$ inside $\ln(2 \cdot 3)$




FIGURE 1.4.27: Checking that $\ln(2 \cdot 3) - \ln(3)$ equals $\ln(2)$

In [FIGURE 1.4.27](#), the lengths of the straight sides of the two regions

1.4 Logarithms and exponentials

are shown in red. To each point (x, y) on the left, there corresponds a point $(3x, \frac{y}{3})$ on the right, and vice versa. For example, the corner $(1, 1)$ goes to the corner $(3, \frac{1}{3})$, the corner $(2, \frac{1}{2})$ on the left goes to the corner $(6, \frac{1}{6})$. What's more, $y = \frac{1}{x}$ is the same as $xy = 1$ and this, in turn, is the same as $(3x)(\frac{y}{3}) = 1$ so points on the left portion of the graph get sent to points on the right portion.

Sending (x, y) to $(3x, \frac{y}{3})$ stretches the region on the left of [FIGURE 1.4.27](#) horizontally by a factor of 3 and simultaneously squishes or compresses it vertically by a factor of 3. All the horizontal lengths on the right are 3-times the corresponding horizontal lengths on the left, and all vertical lengths on the right are $\frac{1}{3}$ of the corresponding vertical lengths on the left. The horizontal stretch multiplies areas by 3 and the vertical squish divides them by 3 so the overall effect is to *leave areas unchanged*. So the areas on the left and right of [FIGURE 1.4.27](#) are indeed equal!

So much for the [LOGARITHM PROPERTY 1.4.14](#) in the special case $a = 2$ and $b = 3$. But there's absolutely nothing special about this case! We can see this by redrawing figures [FIGURE 1.4.26](#) and [FIGURE 1.4.27](#) with each 2 replaced by an a and each 3 replaced by a b as shown in [FIGURE 1.4.28](#) and [FIGURE 1.4.29](#).




FIGURE 1.4.28: Showing $\ln(b)$ inside $\ln(a \cdot b)$




FIGURE 1.4.29: Checking that $\ln(a \cdot b) - \ln(b)$ equals $\ln(a)$



1.4 Logarithms and exponentials

Once again “stretching and squishing”—but this time by a factor of b instead of 3—takes the left region in [FIGURE 1.4.29](#) to the right region and leaves the area unchanged.

Now’s a good time to tackle values of $\ln(a)$ with $0 < a < 1$. Let’s put $b = \frac{1}{a}$ so $b > 1$ and recall that in [AREA DEFINITION OF \$\ln\$ 1.4.21](#), we set $\ln(a)$ or $\ln(\frac{1}{b})$ equal to the *negative* of the area under $y = \frac{1}{x}$ for x between $a = \frac{1}{b}$ and 1. First, note that a negative is forced on us by the rule $\ln(\frac{1}{b}) = -\ln(b)$ from [OTHER LOGARITHM PROPERTIES 1.4.19.ii](#)). Since we have the signs straight, what we need to check, to verify this property, is that the area under $y = \frac{1}{x}$ for x between $\frac{1}{b}$ and 1 is equal to the area under $y = \frac{1}{x}$ for x between 1 and b . This follows by the same “stretch and squish” argument that we used to show the [LOGARITHM PROPERTY 1.4.14](#). I’ve made it a problem for you to practice with.

PROBLEM 1.4.30:

i) Show that the area on the left of [FIGURE 1.4.31](#) equals the area on the right of [FIGURE 1.4.31](#) by “stretching and squishing” by a factor of 2.




FIGURE 1.4.31: Checking that $\ln(\frac{1}{2}) = -\ln(2)$

ii) Once again, there’s nothing special about 2. Show that the area on the left of [FIGURE 1.4.32](#) equals the area on the right of [FIGURE 1.4.32](#) by “stretching and squishing” by a factor of b .




FIGURE 1.4.32: Checking that $\ln\left(\frac{1}{b}\right) = -\ln(b)$

OK, so the **AREA DEFINITION OF \ln 1.4.21** has the **LOGARITHM PROPERTY 1.4.14**. Are there others?

UNIQUENESS OF \ln 1.4.33: *If $\log(x)$ is any function that has the **LOGARITHM PROPERTY 1.4.14**, then $\log(x) = \ell \cdot \ln(x)$ for some constant ℓ . In other words, up to taking constant multiples, the only function with the **LOGARITHM PROPERTY 1.4.14** is \ln .*

Showing this takes some work, and, since we won't need this fact, I won't digress to give any details. I just want to note that it tells us that, strange as the **AREA DEFINITION OF \ln 1.4.21** may appear, it is essentially the *only* way to define a logarithm function. But we can use it to answer **ANOTHER SIMPLE QUESTION 1.4.12**: “What do we mean by the function $\log_{10}(x)$?”

PROBLEM 1.4.34: Why do we expect $\log_{10}(x)$ to have each of the following properties?

- The **LOGARITHM PROPERTY 1.4.14**: $\log_{10}(a \cdot b) = \log_{10}(a) + \log_{10}(b)$.
- $\log_{10}(10) = 1$.

By **UNIQUENESS OF \ln 1.4.33**, property i) tells you that $\log(x) = \ell \cdot \ln(x)$ for some ℓ . Use property ii) to show that $\ell = \frac{1}{\ln(10)}$ and conclude that $\log_{10}(x) = \frac{1}{\ln(10)} \ln(x) \simeq 0.4342944819 \ln(x)$.

1.4 Logarithms and exponentials

In other words, even if we only wanted to use logarithms base 10, we'd still need to understand the natural logarithm because \log_{10} is just a multiple of \ln . Fine, but why *not* take such a multiple like \log_{10} ? Why do we call \ln the *natural* logarithm? Because, while we *could* have studied $\log_{10}(x) \simeq 0.4342944819 \ln(x)$, to do so we'd have had to draw pictures of areas under $y = \frac{0.4342944819}{x}$. Aside from the nuisance of that 0.4342944819, this would make it much harder to see the “stretch-and-squish” pictures behind the [LOGARITHM PROPERTY 1.4.14](#). We'll see another example of why it's just easier to work with \ln in [BERNOULLI'S LIMIT FOR \$\ln\$ 1.4.42](#).

Now let's turn to a question that I hope many of you have been wondering about. What does checking that the [AREA DEFINITION OF \$\ln\$ 1.4.21](#) has the [LOGARITHM PROPERTY 1.4.14](#) buy us, if, as it appears, there's no way to *calculate* the number $\ln(a)$. Once again, first impressions are misleading. It's actually quite easy to use the [AREA DEFINITION OF \$\ln\$ 1.4.21](#) to calculate approximate values of $\ln(x)$. The area definition makes it clear that $\ln(1) = 0$: just look at the picture in [FIGURE 1.4.35](#).



FIGURE 1.4.35: The “area” defining $\ln(1)$

But we can go a lot further. We can estimate $\ln(b)$ for any b by simply “counting the squares”. By this, I mean that we plot the region whose area gives $\ln(b)$ on some “graph paper” ruled with squares of a fixed size. I've done this for $b = 2$ in [FIGURE 1.4.36](#) using squares of side $\frac{1}{4}$ (and area $\frac{1}{16}$).

We then just outline, as I have also done, those squares that lie completely under the graph and those that lie completely above the graph, getting a region underneath whose area— $\frac{9}{16}$ since it contains 9 of the squares—is definitely *less* than $\ln(2)$ and another above






FIGURE 1.4.36: Estimating $\ln(2)$ with squares of side $\frac{1}{4}$

whose area— $\frac{14}{16}$ —is definitely *greater* than $\ln(2)$. In other words we find that $0.56 \approx \frac{9}{16} < \ln(2) < \frac{14}{16} \approx 0.88$.

PROBLEM 1.4.37: By counting squares in [FIGURE 1.4.38](#), where each



FIGURE 1.4.38: Estimating $\ln(2)$ with squares of side $\frac{1}{10}$

square has side $\frac{1}{10} = 0.1$ and area $\frac{1}{100} = 0.01$, show that $0.63 < \ln(2) < 0.77$.

It's clear that, by taking small enough squares, we can get lower and upper bounds for $\ln(2)$ (or any other logarithm value) that are as accurate as we like. For example, in [SECTION 5.4](#), we'll use the approximate value $\ln(2) \approx 0.693$ a lot, and if we ever needed something very accurate, we could compute the better approximation 0.6931471805599453 . But it's also clear that we would not want to have to get these answers by counting the squares, because to get such accuracy you need too many squares. Even with 100 squares, we weren't able to completely nail the first place in $\ln(2)$ and it turns out that to get the value $\ln(2) \approx 0.693$ you need to use squares with side 0.0005 —and there are 4,000,000 of these to check.

Much better ways of calculating logarithms are provided by calculus—there's even a separate subject called numerical analysis that specif-

1.4 Logarithms and exponentials

ically studies fast ways to get such approximations. Fortunately, we can get all the values of \ln that we'll ever need from our calculator. You'll notice that it has a key labeled LN, and to get a value like $\ln(2) \approx 0.69314718056$ you just type LN 2 ENTER—on other calculators you may need to reverse the order and type 2 LN ENTER. Try to find $\ln(2)$ on your calculator now so you'll know how to compute values when we need them.

PROBLEM 1.4.39: Using the [LOGARITHM PROPERTY 1.4.14](#), the [OTHER LOGARITHM PROPERTIES 1.4.19](#) and the values

$$\ln(2) = 0.69314718056 \text{ and } \ln(3) = 1.0986122887$$

to predict what each of the values of \ln below should be. Then, use your calculator to check your prediction.

- i) $\ln(6)$
- ii) $\ln(\frac{1}{3})$
- iii) $\ln(\frac{2}{3})$
- iv) $\ln(9)$ Hint: $4 = 2^2$.
- v) $\ln(0.125)$ Hint: $0.125 = \frac{1}{8}$.
- vi) $\ln(\sqrt{2})$

PROBLEM 1.4.40: Find $\ln(2.718281828459)$.

Now that we've seen how to use the area definition of \ln to compute approximate values, I have a small confession to make. Right after I have the [AREA DEFINITION OF \$\ln\$ 1.4.21](#), I said it was “certainly a perfectly good way to define *the* number $\ln(a)$ ”. Don't look down, but let me now ask, “Is that really so certain?” Once again, we have an intuitive feel for areas, based on working with simple regions like rectangles, triangles and circles, and we expect them to have certain properties. For example: “The sum is the whole of the parts”—if we cut up a region into pieces the area of the region as a whole is the sum of the areas of the pieces we cut it into; or, “Stretching horizontally by a factor of 3 and squishing vertically by a factor 3 leaves the

area of any region unchanged”. I mention these properties because we used both of them to see that \ln had the logarithm property.

How do we know that areas inside “curvy” regions, like the region under the graph of $y = \frac{1}{x}$, have these same properties? When we come right down, to it, “How do we know exactly *what* number we mean when we speak about the area of a curvy region ?”—like the ones associated to $\ln(2)$ or $\ln(3)$ in [FIGURE 1.4.23](#) and [FIGURE 1.4.24](#). Once again, the fact that we can compute approximate areas like $\ln(2) \simeq 0.69314718056$ does *not* mean we can stick a fork into the exact area $\ln(2)$. So, how do we know *exactly* what number an area like $\ln(2)$ stands for? Ahem. Er, well, um. OK dammit, we don’t.

But I have some good news for you and some better news. The good news is that it’s not hard to answer this question. The better news is that I’m not going to make you learn the details of the answer. Not, I emphasize, because the answer is hard, or because its not important and interesting. In fact, in every one of the physical and social sciences, quantities that can be viewed as areas are all the time, and understanding how to recognize and compute such areas is a basic skill. But because we will not need to work further with areas in the course and I’d rather spend the time on our central topics, I’ll just sketch the basic idea.

Once again, the answer is something that is part of any freshman calculus course, and comes directly from the idea of approximating areas by “counting the squares”. Let’s take $\ln(2)$ as an example. Whatever we mean by this number, it had better be bigger than any of the “areas of the squares completely under” and smaller than any of the “area of the squares completely above” the graph of $y = \frac{1}{x}$. What needs to be checked is that, if we take small enough squares, then we can make the *difference* between the areas of the squares “completely under” and “completely above” the curve *as small as we like*. Because “there are no holes in the real number line”, this tells us that there *is* a number—and only *one*—that is both bigger than all the

1.4 Logarithms and exponentials

“areas completely under” and smaller than all the “areas completely above”. That’s the number we mean by $\ln(2)$.

There’s even a simple picture that makes it clear we can make that difference as small as we like. All we need to do is use rectangles instead of squares. In the picture below, I’ve shown the area defining $\ln(2)$ divided into 10 small strips, and in each strip, I’ve shown the largest rectangle “completely under” and the smallest rectangle “completely above” the graph of $y = \frac{1}{x}$.




FIGURE 1.4.41: Stacking up the differences for $\ln(2)$ using strips

The differences between these areas are shaded. Notice how the bottom of one difference is the top of the next. Because of this, I can stack up all the differences as shown in the “extra” strip on the right. This means that the differences fit into a rectangle of height 1 and width $\frac{1}{10}$ so they total at most $\frac{1}{10}$. To see that I can make the differences less than $\frac{1}{100}$, I’d just need to use strips of width $\frac{1}{100}$; to see they can be made smaller than $\frac{1}{1000000}$ is just choose strips of width ... Right, $\frac{1}{1000000}$! And so on. The general case is not really any harder but, once again, since we’re just going to agree that we “understand” areas and promise that we won’t look down, I’ll just move on.

Next, we state a formula that plays a central role in the mathematics



1.4 Logarithms and exponentials

of finance and that we'll use a lot in [CHAPTER 5](#).

BERNOULLI'S LIMIT FOR \ln 1.4.42: *As n gets larger and larger—we say that n goes to infinity and write $n \rightarrow \infty$ —the product $n \cdot \ln(1 + \frac{1}{n})$ approaches 1. In fancier language, the **limit**, as n goes to infinity, of $n \cdot \ln(1 + \frac{1}{n})$ is 1.*

By [OTHER LOGARITHM PROPERTIES 1.4.19.iv](#)), a consequence is that as $n \rightarrow \infty$, $\ln((1 + \frac{1}{n})^n) \rightarrow 1$. The picture that explains these limits also shows that, in both cases, the numbers we get are always slightly smaller than 1.

Before we explain this limit, let's note that it provides another typical example of how it's easiest to work with the natural logarithm \ln . If we had chosen to work with $\log_{10}(x) \simeq 0.4342944819 \ln(x)$, we'd get a result like [BERNOULLI'S LIMIT FOR \$\ln\$ 1.4.42](#)—but a lot messier: the limit would be 0.4342944819 not 1.




FIGURE 1.4.43: Estimating $\ln(1 + \frac{1}{n})$

Once again, [BERNOULLI'S LIMIT FOR \$\ln\$ 1.4.42](#) becomes clear when we draw the right picture, in this case, [FIGURE 1.4.43](#). What are the coordinates of the right endpoint on the graph of $y = \frac{1}{x}$ marked with a green dot? Since the x -coordinate is $1 + \frac{1}{n}$, the y -coordinate is

$$\frac{1}{1 + \frac{1}{n}} = \frac{n \cdot 1}{n \cdot (1 + \frac{1}{n})} = \frac{n}{n + 1}.$$

1.4 Logarithms and exponentials

The key observation is that this y -coordinate is—as the figure indicates—greater than $1 - \frac{1}{n} = \frac{n}{n} - \frac{1}{n} = \frac{n-1}{n}$. To see this, we just compute the difference:

$$\frac{n}{n+1} - \frac{n-1}{n} = \frac{n \cdot n - (n+1)(n-1)}{(n+1)n} = \frac{n^2 - (n^2 - 1)}{(n+1)n} = \frac{1}{(n+1)n} > 0.$$

What this observation means is that the area defining $\ln\left(1 + \frac{1}{n}\right)$ lies between the area of the rectangles extending from the x -axis to the upper and lower blue horizontals in [BERNOULLI'S LIMIT FOR \$\ln\$ 1.4.42](#).

Both rectangles have width $\frac{1}{n}$. The height of the upper rectangle is 1 since the upper left corner is at the point $(1, 1)$. The height of the lower rectangle is $1 - \frac{1}{n}$. So we conclude by taking areas, that

$$\frac{1}{n} \left(1 - \frac{1}{n}\right) < \ln\left(1 + \frac{1}{n}\right) < \frac{1}{n} \cdot 1.$$

Multiplying these inequalities by n they become,

$$1 - \frac{1}{n} < n \cdot \ln\left(1 + \frac{1}{n}\right) < 1.$$

But now, as we take n larger and larger the left term approaches 1 since $\frac{1}{n}$ goes to 0. Thus, [BERNOULLI'S LIMIT FOR \$\ln\$ 1.4.42](#) gets squeezed between a number rising towards 1 on the left and the number 1 on the right, hence it must equal 1. Note also that, for any n , the value $n \cdot \ln\left(1 + \frac{1}{n}\right)$ is always just a bit *smaller* than 1 as claimed in the last statement of [BERNOULLI'S LIMIT FOR \$\ln\$ 1.4.42](#).

Before we close this subsection, we note two facts that will be important in the next one.

\ln IS INCREASING 1.4.44:

- i) *The function \ln is increasing. That is, if $a < b$, then $\ln(a) < \ln(b)$.*
- ii) *No horizontal line meets the graph of \ln in more than 1 point.*
- iii) *If a and b are positive and $\ln(a) = \ln(b)$, then $a = b$*

Part ii) follows from part i) because, if a horizontal line $y = c$ did meet the graph twice—say at (a, c) and (b, c) with $a < b$, then we'd

1.4 Logarithms and exponentials




FIGURE 1.4.45: Comparing $\ln(a)$ and $\ln(b)$

have numbers $a < b$ with $\ln(a) = c = \ln(b)$. And [iii](#)) is just another way of putting [ii](#)) .

[FIGURE 1.4.45](#) demonstrates why [i](#)) holds when a and b are greater than 1: $\ln(a)$ equals the area of region on the left and $\ln(b)$ is the entire area—both left and right pieces, which is clearly larger. The case when both are less than 1 then follows from the rule $\ln(\frac{1}{b}) = -\ln(b)$ of [OTHER LOGARITHM PROPERTIES 1.4.19.ii](#)). This also shows that if $a < 1 < b$, then $\ln(a) < 0 < \ln(b)$.




FIGURE 1.4.46: The graph of $\ln(x)$

A less formal way to convince yourself that [ln IS INCREASING 1.4.44](#) is just to look at it's graph shown in [FIGURE 1.4.46](#), with a few key values marked. Because [ln IS INCREASING 1.4.44](#), there is a unique number a for which $\ln(a) = 1$ which is shown on the graph. The number turns out to be so important in virtually every area of mathematics that it has a name.

THE NUMBER e 1.4.47: *We use the letter e to denote the number*

such that $\ln(e) = 1$.

Why use a name and not just write down the number e ? Because we can't write down e ! In [PROBLEM 1.4.40](#), we saw that e is pretty close to 2.718281828459. Pretty close but *not* not equal, as we can see by using a calculator that gives a few more digits:

$$\ln(2.718281828459) \approx 0.99999999999998335884.$$

It turns out, that like $\sqrt{2}$, e is irrational, although it's much harder to see this. But, because it is, as I explained when we were discussing [IS IRRATIONAL 1.4.8](#), we can't *ever* write it down exactly as a decimal. Since it's so important, the only solution is to name it.

One question you might have after looking at [FIGURE 1.4.46](#) is whether $\ln(x)$ keeps getting bigger and bigger as x does, or whether there is some ceiling it never gets beyond. The answer is the former, as you can see from [OTHER LOGARITHM PROPERTIES 1.4.19.iv](#)). This shows that if I want a number whose natural log is 1000000, I can just take $e^{1000000}$ because $\ln(e^{1000000}) = 1000000 \ln(e) = 1000000$. So \ln does eventually become as large as you please—we write $\ln(x) \rightarrow \infty$ as $x \rightarrow \infty$. Likewise, we can get any negative number: to get -1000000 take $\ln(\frac{1}{e^{1000000}})$.

One disclaimer is in order: $\ln(x)$ may get as big as you like, but it gets big *ve-e-e-ry* slowly. For example, $e^{1000000}$ has over 400000 *digits*! The number with natural logarithm 100 is about 5184705528587072464087.4533229.

ln TAKES ON ALL REAL VALUES 1.4.48: *For any real number c , there is a positive real a for which $\ln(a) = c$.*

This is just another way to say that we can make the value of $\ln(a)$ as positive or as negative as we like by choosing the right a .

UNIQUENESS OF VALUES OF \ln 1.4.49: *Every horizontal line meets the graph of \ln in a unique point.*

1.4 Logarithms and exponentials

The line $y = c$ *does* meet the graph at the point (a, c) for which $\ln(a) = c$, and cannot meet it anywhere else because [ln IS INCREASING 1.4.44.ii](#)).

exp and exponentials

Since we digressed for a long time to understand the natural logarithm and its properties, permit me to remind you that this whole section is an attempt to find an answer to the [SIMPLE QUESTION 1.4.1](#), “What number is $10^{\sqrt{2}}$?” We haven’t even mentioned an exponential base 10 for quite a few pages. Yet, it turns out that we’re actually almost ready to answer that question. In fact, we’ll be able to explain what we mean by b^a for any positive real base b and any real exponent a . There’s just a bit more work to do, but, as we undertake it, I’d like you to keep this goal in mind.

The last big step is to use natural logarithm function—in particular, the fact that [ln IS INCREASING 1.4.44](#)—to define a new function \exp . We’ll do this by reflecting the graph $y = \ln(x)$ in the line $y = x$.




FIGURE 1.4.50: Reflection in $y = x$

[FIGURE 1.4.50](#) shows what this reflection does both geometrically and in coordinates. The reflection of a horizontal line like $x = a$ (shown in blue) is the vertical line $y = a$ (of the same color), and vice-versa;

1.4 Logarithms and exponentials

the point (a, a) where these lines meet stays fixed and both lines make angles of 45° with $y = x$. The reflection this takes the point (a, c) where $x = a$ meets $y = c$ to the point (c, a) where $x = c$ meets $y = a$; the line $y = x$ bisects the line segment joining the points (a, c) and (c, a) and is perpendicular to it. Simply put, the reflection just swaps x - and y -coordinates.

REFLECTION DEFINITION OF exp 1.4.51: *The function exp is the function whose graph is the reflection in the line $y = x$ of the graph of the natural logarithm \ln . In coordinates,*

$$a = \exp(c) \iff c = \ln(a)$$




FIGURE 1.4.52: The graphs of the exp and \ln functions

There's one implicit assumption in [REFLECTION DEFINITION OF exp 1.4.51](#) that needs checking. We claim to be defining a *function* exp. But we'll only get a function if our recipe produces a unique value $a = \exp(c)$ for every input real c . In other words, we need to check any vertical line $x = c$ meets the graph of exp in just 1 point. But, by reflecting in the line $y = x$ again, the points in which $x = c$ meets the graph of exp are exactly the points in which the horizontal line $y = c$ meets the graph of \ln . There's exactly 1 such point because [UNIQUENESS OF VALUES OF \$\ln\$ 1.4.49](#).

FIGURE 1.4.52 shows both graphs with a few of the most important pairs of corresponding points on each. In particular, the points $(0, 1)$ and $(e, 1)$ are on the graph of \exp because the points $(1, 0)$ and $(1, e)$ are on the graph of \ln . In other words, $\exp(0) = 1$ because $\ln(1) = 0$ and $\exp(1) = e$ because $\ln(e) = 1$. Similarly, $\exp(\ln(2)) = 2$ and since $\ln(2) \approx 0.693$, we find that $\exp(0.693) \approx 2$.

Once again, in the fact that $(\ln(2), 2)$ is on the graph of \exp —in other words, that $\exp(\ln(2)) = 2$ —there's nothing special about 2. It's a general fact that taking the \exp of *any* \ln takes you back where you started. The fancy way to express this is to say that \exp and \ln are **inverse** functions.

exp AND ln ARE INVERSES 1.4.53: *For any $a > 0$, $\exp(\ln(a)) = a$ and for any c , $\ln(\exp(c)) = c$. That is, each of the functions \exp and \ln undoes the transformation wrought by the other.*

Two comments are in order. First, the use of the term “inverse” has nothing to do with the 1-over— $\frac{1}{\cdot}$ —kind of inverse by division. When we speak of two functions being inverses, we mean that applying first one and then the other to any starting value returns us to that starting value. In this sense, the inverse of putting on your pants is taking off your pants. More mathematical examples are that the inverse of “adding 5” is “subtracting 5” or that the inverse of “squaring” a positive number is taking its “square root”.

PROBLEM 1.4.54: Here's a bit of easy practice with inverse functions.

- i) Suppose that f and g are two operations that have inverse operations F and G , respectively. Show that the operation “first f , then g ” has an inverse operation and express this inverse in terms of F and G . Hint: What if f is “putting on your sock” and g is “putting on your shoes”?
- ii) What is the inverse of “first add 5, then double”?
- iii) Is there any way to fill in the blank in “first halve and then _____” which yields the inverse of “first add 5, then double”?

1.4 Logarithms and exponentials

Second, note that the starting values a we can *input* into \ln must be positive, because we only defined $\ln(a)$ when $a > 0$ but that any real b occurs as an *output* value of \ln by **ln TAKES ON ALL REAL VALUES 1.4.48**. Thus, we can feed any starting value c into \exp because but no matter what c we choose the output value a will be positive. We can do a bit better: not only is any exponential value positive, but $\exp(c)$ gets larger as c does—that is, \exp is also an increasing function.

PROBLEM 1.4.55:

- Use **ln IS INCREASING 1.4.44** to show that, if $c = \ln(a)$, $d = \ln(b)$ and $c < d$, then $a < b$. Hint: Draw this on the graph of $y = \ln(x)$.
- Show that the function \exp is *increasing*. That is, if $c < d$, then $\exp(c) < \exp(d)$. Hint: How is the a above related to $\exp(c)$?

In a similar vein, we saw, in checking that **ln TAKES ON ALL REAL VALUES 1.4.48**, that while \ln gets as big as we like, it does so very slowly. This means that for large a the graph of \ln near a is almost horizontal. Reflecting, the graph of $\exp(b)$ near a large b must be almost *vertical*. In other words, not only does $\exp(b)$ get as big as we like— $\exp(x) \rightarrow \infty$ as $x \rightarrow \infty$ —but it does so *ve-e-e-ry* fast.

We won't need to use these properties but you'll see lots of examples of them when we study interest. Both the \ln and \exp functions are radically different from power functions and their growth for large values of x is the most obvious evidence. For very large x , the function $\ln(x)$ grows more slowly, and the function $\exp(x)$ more rapidly, than any positive rational power of x . Eventually, for example, $\ln(x) < x^{\frac{1}{1000}}$ and $e^x > x^{1000}$ but we could use any positive number in place of 1000.

As a first application of the fact that **exp AND ln ARE INVERSES 1.4.53**, let's see what **BERNOULLI'S LIMIT FOR ln 1.4.42** says about the exponential function. When we apply the \exp function to the last version— $\ln\left(1 + \frac{1}{n}\right)^n \rightarrow 1$ —of **BERNOULLI'S LIMIT FOR ln 1.4.42**, we get $\exp\left(\ln\left(1 + \frac{1}{n}\right)^n\right) \rightarrow \exp(1)$. Since the \exp on the left undoes

1.4 Logarithms and exponentials

the \ln leaving just the exponential inside and $\exp(1) = e$ by definition, we obtain:

BERNOULLI'S LIMIT FOR \exp 1.4.56: *As n goes to infinity, the exponential $(1 + \frac{1}{n})^n$ approaches the limiting value e . In fancier notation, as $n \rightarrow \infty$, of $(1 + \frac{1}{n})^n \rightarrow e$. Moreover, for any given n , because the logarithm is slightly smaller than 1, the corresponding exponential is always a bit smaller than e .*

I'll make a few comments on this limit at the end of the section when I finish reviewing the history of \ln and \exp . It's a key tool in [CHAPTER 5](#), allowing us to use logarithms and exponentials to study the time value of money and many other phenomena.

We are now going to squeeze a series of amazing properties out of the fact that [exp AND ln ARE INVERSES 1.4.53](#). Broadly speaking, this lets us turn any property of \ln into an inverse property of \exp . The model uses the fundamental [LOGARITHM PROPERTY 1.4.14](#). This says that \ln converts products of inputs into sums of outputs, so it translates to say that \exp converts sums of inputs into products of outputs. In formulae, because $\ln(a \cdot b) = \ln(a) + \ln(b)$, we expect that $\exp(c + d) = \exp(c) \cdot \exp(d)$.

Our method for checking such inverse property equations is stupidly simple. Suppose that a and b are two positive numbers and that we know that $\ln(a) = \ln(b)$. Then $a = b$ because [FIGURE 1.4.45.iii](#)). We take a and b to be the left and right hand sides of the equation we want to check. Why can we assume that these are positive? First, all values (outputs) of \exp are positive because they are inputs to \ln and we only allow these to be positive. This means that as long as we only perform multiplications, divisions and powers of \exp values, we can only produce positive values.

Now let's try our method out above. We have $a = \exp(c + d)$ and $b = \exp(c) \cdot \exp(d)$. Then $\ln(a) = \ln(\exp(c + d)) = c + d$ because [exp AND ln ARE INVERSES 1.4.53](#). But $\ln(b) = \ln(\exp(c) \cdot \exp(d)) = \ln(\exp(c)) +$

1.4 Logarithms and exponentials

$\ln(\exp(d))$ by the [LOGARITHM PROPERTY 1.4.14](#). And $\ln(\exp(c)) + \ln(\exp(d)) = c + d$ by using that [exp AND ln ARE INVERSES 1.4.53](#) on each term separately.

PROBLEM 1.4.57:

i) Show that $\exp(-d) = \frac{1}{\exp(d)}$. Hint: Use [OTHER LOGARITHM PROPERTIES 1.4.19.ii](#)) to see that the left and right hand sides have equal logarithms.

ii) Show that $\exp(c - d) = \frac{\exp(c)}{\exp(d)}$. Hint: You can copy the argument in the preceding paragraph almost verbatim. All you need to do is to replace the use of the [LOGARITHM PROPERTY 1.4.14](#) by [OTHER LOGARITHM PROPERTIES 1.4.19.iii](#)).

[OTHER LOGARITHM PROPERTIES 1.4.19.iv](#)) saying that, for rational c , $\ln(b^c) = c \cdot \ln(b)$ will be the key that will unlock general exponentials for us. We can think of it as saying that \ln converts (rational) exponentiation of the input to multiplication of the output. Hence, we expect that:

KEY EXPONENTIAL PROPERTY 1.4.58: *The function \exp converts rational multiplication of its input into exponentiation of its output. That is:*

$$\text{If } c \text{ is rational then, } \exp(c \cdot d) = \exp(d)^c.$$

We've already seen the footwork needed twice, in checking [OTHER LOGARITHM PROPERTIES 1.4.19.iv](#)) and in answering [EVEN SIMPLER QUESTION 1.4.2](#). First handle whole numbers m , then Egyptian fractions—those that are of the form $\frac{1}{m}$ —and finally general rational fractions $c = \frac{m}{n}$.

The first step is easy. First, $\exp(2 \cdot d) = \exp(d + d) = \exp(d) \cdot \exp(d) = (\exp(d))^2$. Then, $\exp(3 \cdot d) = \exp(2 \cdot d + d) = \exp(2 \cdot d) \cdot \exp(d) = \exp(d)^2 \cdot \exp(d) = \exp(d)^3$. Continuing in this way we see that $\exp(m \cdot d) = \exp(d)^m$ for any positive whole number m .

PROBLEM 1.4.59: Show that $\exp(-m \cdot d) = \exp(d)^{-m}$. Hint: $\exp(-(m \cdot d)) = \frac{1}{\exp(md)} = \frac{1}{\exp(d)^m}$ by first using [PROBLEM 1.4.57.i](#)) and then the positive m case just proved.



1.4 Logarithms and exponentials

We handle Egyptian fractions by cancelling them: $d = m(\frac{1}{m}d)$ so $\exp(d) = \exp(m(\frac{1}{m}d)) = \exp(\frac{1}{m}d)^m$. Then, taking m^{th} roots $\exp(d)^{\frac{1}{m}} = (\exp(\frac{1}{m}d)^m)^{\frac{1}{m}} = \exp(\frac{1}{m}d)$.

Similarly, $\frac{m}{n}d = m(\frac{1}{n}d)$ so

$$\exp(\frac{m}{n}d) = \exp(m(\frac{1}{n}d)) = \exp(\frac{1}{n}d)^m = \exp((\frac{1}{n}d)^m)^{\frac{1}{m}}$$

by using the whole number and Egyptian fraction cases. Finally, we use $(b^x)^z = b^{x \cdot z}$ —[RULES OF EXPONENTS 1.4.10.ii](#)—, setting $b = \exp(d)$, $x = m$ and $y = \frac{1}{n}$ to get $(\exp(d)^m)^{\frac{1}{n}} = \exp(d)^{\frac{m}{n}}$ as desired.

An immediate consequence of the [KEY EXPONENTIAL PROPERTY 1.4.58](#) is a way of computing $\exp(c)$ for any rational c : $\exp(c) = \exp(c \cdot 1) = \exp(1)^c = e^c$.

RATIONAL EXPONENTIALS ARE POWERS OF e 1.4.60: [THE NUMBER \$e\$ 1.4.47](#) has the property that, for any rational c , $\exp(c) = e^c$.

We're now at the punch line. We now *define* e^x to be the number $\exp(x)$ for *any* real number x . This is potentially ambiguous when x is rational number like 4. We already have a number in mind when we write e^4 , namely $e \cdot e \cdot e \cdot e$ but because [RATIONAL EXPONENTIALS ARE POWERS OF \$e\$ 1.4.60](#), this number and our new number $\exp(4)$ are the same.

What I want to emphasize, however, is that, when x is an irrational number like $\sqrt{2}$, the *only* way we have of identifying the number $e^{\sqrt{2}}$ is as the function value $\exp(\sqrt{2})$. The exponential notation is a convenient one because it reminds of properties of this number that we may find useful. One way to think of equations like $\exp(c + d) = \exp(c) \cdot \exp(d)$ and the [KEY EXPONENTIAL PROPERTY 1.4.58](#) as saying that the function values $\exp(x)$ walk like exponentials and talk like exponentials. By writing e^x instead of $\exp(x)$, we make these numbers dress like exponentials too. Moreover, the self-deception we practice when we write e^x for $\exp(x)$ will never lead us into any

1.4 Logarithms and exponentials

computational errors: any rule that we “expect” to hold for e^x *does* because all the rules of exponentials hold for $\exp(x)$.

But please remember that except when x is rational, there’s no way to write down e^x in terms of powers of e . Even though we write e^x , the number we are identifying in this way IS the value y output by the \exp function from the input x . Further, the only way we have of identifying that output value y is to say that it is the input value on which the function \ln outputs x .

One further advantage of writing $\exp(x) = e^x$ is that it links up our [AREA DEFINITION OF \$\ln\$ 1.4.21](#) with the more familiar interpretation of logarithms in terms of exponents. Remember that, by the [REFLECTION DEFINITION OF \$\exp\$ 1.4.51](#), saying that $a = \exp(c)$ is the same as saying that $c = \ln(a)$. These are now the same as saying that $a = e^c$. In other words, the natural logarithm c of a positive number a is just the *exponent* to which we have to raise e to obtain a .

DEFINING e^x TO BE $\exp(x)$ 1.4.61: *For any real exponent x , the exponential e^x is defined to be the number $\exp(x)$, although e^x only makes sense as a “power of e exponential” when x is rational.*

EXPONENT INTERPRETATION OF $\ln(a)$ 1.4.62: *For any positive a , we interpret $\ln(a)$ as the exponent c for which $e^c = a$, with the proviso that unless this number c is rational, what we really mean is the number c for which $\exp(c) = a$.*

OK, you can look down now. We’re back on solid ground and unlike in the Road Runner™, there’ll be no drawn out whistling sound as we fall. It’s been a long journey but now we are in a position to understand what number is referred to by any exponential b^x for any real exponent x and any positive real base b . So we are at last in a position to answer the [SIMPLE QUESTION 1.4.1](#) and say what we mean by $10^{\sqrt{2}}$. Isn’t there an easier way? Unfortunately, not. There *are* other ways to make sense out of exponentials like $10^{\sqrt{2}}$ but they are even more circuitous and require an understanding of a lot of delicate

1.4 Logarithms and exponentials

tools from calculus. Our apparently roundabout path through the study of the natural logarithm function \ln and its inverse, the exponential function \exp is, in fact, the *shortest* one to a firm footing with exponentials.

So what is the answer to the [SIMPLE QUESTION 1.4.1](#)? Well, $10^{\sqrt{2}} = \exp(\ln(10) \cdot \sqrt{2})$; or, using the exponential notation, we all prefer $10^{\sqrt{2}} = e^{\ln(10) \cdot \sqrt{2}}$. This answer is *forced* on us if we want the rules of exponents to continue to hold for irrational exponents. In the rule $(b^x)^z = b^{(x \cdot z)}$ —[RULES OF EXPONENTS 1.4.10.ii](#)—set $b = e$, $x = \ln(10)$ and $z = \sqrt{2}$. Because $10 = \exp(\ln(10)) = b^x$, the left side $(b^x)^z$ equals $10^{\sqrt{2}}$. But the right side is $b^{(x \cdot z)} = e^{(\ln(10) \cdot \sqrt{2})}$. So the rules of exponents tell us that $10^{\sqrt{2}}$ can *only* be defined by $10^{\sqrt{2}} = e^{\ln(10) \cdot \sqrt{2}}$. More generally,

A SIMPLE ANSWER 1.4.63: *For any positive real base b and any real exponent x , the exponential b^x is defined to be the number*

$$b^x := \exp(\ln(b) \cdot x) = e^{\ln(b) \cdot x}.$$

Once again we have no choice. When x is rational, this choice agrees—as it had better—with the “powers of b ” definition we gave at the start of the section: $b = \exp(\ln(b))$ so $b^x = (\exp(\ln(b)))^x$. But by the [KEY EXPONENTIAL PROPERTY 1.4.58](#), $(\exp(\ln(b)))^x = \exp(x \cdot \ln(b)) = \exp(\ln(b) \cdot x)$. And when x is not rational, this definition is the only one that ensures that b^x continues to obey the same rules of exponents that hold when x is rational.

Perhaps this is a good time for me to ‘fess up. As long as you didn’t look down, you wouldn’t have fallen anyway. You can ask your calculator to find b^x for any reasonable b and x and it’ll happily oblige you with more decimals than you’ll ever need. What’s more those decimals will be exactly what you would have got by rounding the number $\exp(\ln(b) \cdot x)$ to the same number of places. That’s because what your calculator actually *does* when you ask it to find b^x is to compute $\exp(\ln(b) \cdot x)$.

1.4 Logarithms and exponentials

It's easy to see why this is smart. It's not really practical to find exponentials by taking powers. Even if we are willing to pretend that $\sqrt{2} = \frac{1414213562373}{1000000000000}$ (and not just close to this decimal), it's just not very easy to compute $10^{\frac{1414213562373}{1000000000000}}$. But calculus gives us very fast techniques to compute the two functions \exp and \ln and once we have these functions under control we have everything we need to find b^x —or rather $\exp(\ln(b) \cdot x)$ —for any b and any x .

How *do* we compute values of \exp on a calculator? Most use the 2nd option for the key labeled LN, and to get a value like $\exp(2.31) \simeq 10.074424655$ you just type 2nd LN 2.31 ENTER—on other calculators you may need to reverse the order and type 2.31 LN ENTER. Try to find $\exp(2.31)$ on your calculator now so you'll know how to compute values when we need them. Here is a bit more practice.

PROBLEM 1.4.64: Using the [KEY EXPONENTIAL PROPERTY 1.4.58](#) and the properties in [PROBLEM 1.4.57](#) and the values $\exp(2) \simeq 7.38905609893$ and $\exp(3) \simeq 20.0855369232$ and the [OTHER LOG-ARITHM PROPERTIES 1.4.19](#) to predict what each of the values of \ln below should be. Then, use your calculator to check you prediction.

i) $\exp(5)$.

ii) $\exp(-2)$.

iii) $\exp(-1)$ Hint: $-1 = 2 - 3$.

iv) $\exp(\frac{2}{3})$. Hint: $\frac{2}{3} = \frac{1}{3} \cdot 2$ so you can apply [KEY EXPONENTIAL PROPERTY 1.4.58](#).

PROBLEM 1.4.65: We can write 6 as either $2 \cdot 3$ or $3 \cdot 2$. Use this observation, the [KEY EXPONENTIAL PROPERTY 1.4.58](#), and the values $\exp(2) \simeq 7.38905609893$ and $\exp(3) \simeq 20.0855369232$ to compute $\exp(6)$ in two ways. Then check that both give the right value by computing $\exp(6)$ directly.

OK, so hasn't this whole section just been a gigantic waste of time? Why couldn't we have just left it to our calculator to worry about exponentials? The answer is that we could have if we *only* wanted to



1.4 Logarithms and exponentials

compute exponential values. But later in the course, when we come to statistics, and especially when we study the mathematics of finance, we'll need to know many basic properties of the \ln and \exp functions.

All the properties we will need have been revealed in the course of our attempt to understand what numbers we mean when we write general exponentials. The [SIMPLE QUESTION 1.4.1](#) and the many other questions about exponentials that we have answered in this section have guided us to discover a beautiful pair of functions—probably the two most important functions in all mathematics—and I hope you'll view this as our main achievement. Figuring out what we mean when we write exponentials with irrational exponents turns out to be a bonus tossed in with our wider knowledge about these fundamental mathematical objects.

Before we close this section, here's the end of the story of these functions from the point at which I broke off when introducing logs. At that point, mathematicians and astronomers knew how to calculate approximate values of logarithms and how to use them to simplify calculations but they were still in the Wile E. Coyote position—suspended in midair—because they had no way to really say what number a value like $\ln(2)$ represented. In fact, it was more common to work with logarithms base 10 (usually denoted today as \log_{10} —we won't need these logarithms at all), although, in the earliest days, there wasn't even a completely clear notion of base.

In fact, it was quarter of a century before the [AREA DEFINITION OF \$\ln\$ 1.4.21](#) was discovered and the existence of logarithms was on a sound footing. The idea that areas underneath the hyperbola $y = \frac{1}{x}$ have properties like those of a logarithm was probably first discovered by a Jesuit Mathematician named Grégoire de Saint-Vincent. In particular, he showed the key equality, that of the two areas in [FIGURE 1.4.29](#) in a book on areas and conic sections that he published in 1647.



1.4 Logarithms and exponentials

His ideas were refined by his student, Alphonse Antonio de Sarasa, another Jesuit mathematician. Christian Huygens, a great Dutch mathematician, astronomer, physicist and—clockmaker, clarified the connection to logarithms and exponentials and exploited it to produce some very accurate values. For example, he calculated that $\frac{1}{\ln(10)} \simeq 0.43429448190325180$ which is incorrect only in the last place!

The name *natural logarithm* for our function \ln was first used by the Danish mathematician Nicolaus Mercator (born Kauffmann) in a book published in 1668 in which he showed clearly how to use areas underneath the curve $y = \frac{1}{x}$ to compute logarithms. Curiously, he was best known as a—yes!—clockmaker and was made a member of the Royal Society of England, not for his mathematical work, but for devising a pendulum clock that would keep accurate time at sea and thus enable ships' captains to determine their longitude.

Rather amazingly, given the enormous amount of work that has been done on logarithms, no one had ever put a fork into the fundamental constant e . An approximate value of e appears in tables of logarithms of as early as 1619, but it was not identified and christened for another 70 years. [BERNOULLI'S LIMIT FOR exp 1.4.56](#) calculates e exactly. The Bernoulli in the name is the Swiss mathematician Jakob (not to be confused with Johann, his brother or Daniel, his nephew, both also important mathematicians). Jakob was led to consider the limit by questions about compound interest and it will be to answer these very questions that we'll need his formula in [Section 5.4](#). But he did not give the limiting value any name, only showing that the powers *have* a limit and that this limit lies between 2 and 3.

The number e was baptized, but under the name b , in 1690, in letters written to Huygens by Gottfried Leibniz, a great German mathematician (and polymath who contributed to too many other fields to even list), most famous today for having developed the basic ideas of calculus independently of Isaac Newton. For a while both b and c were

commonly used to denote e .

The man who first identified the number e as we do here (calling it “that number whose hyperbolic [i.e. natural] logarithm = 1”) and used the letter e to denote it was the incomparable Leonhard Euler, a Swiss mathematician and physicist. Euler was the greatest mathematician of the 18th century and one of the greatest of all time. Publication of a complete edition of his collected works began in 1911 and although 76 volumes have appeared is still not complete in 2009! Just the 29 volumes of his mathematical works fetch about \$4,500 but most of the papers are available at [The Euler Archive](#). Believe it or not about half of these papers were written after he went totally blind—dictated from memory to scribes.

Corollarium II.

171. Quamquam autem in ista aequatione ipsa potentia p non inest, tamen eius directio a qua relatio elementorum dx et dy pender, adhuc super est. Data igitur directione potentiae punctum in quouis loco sollicitantis, et ipsa curua in qua punctum mouetur, poterit ex his solis datis determinari puncti celeritas in quouis loco. Erit enim $\frac{de}{e} = \frac{dy}{r dx}$ seu $e = e^{\int \frac{dy}{r dx}}$ vbi e denotat numerum, cuius logarithmus hyperbolicus est 1.

FIGURE 1.4.66: Euler’s first published use of e

Euler first used the letter e in a manuscript from 1727 but this paper was not published until 1862. [FIGURE 1.4.66](#) shows, in the last 2 lines, the first published use of the notation from Chapter 2 of his *Mechanica*, published in 1736. Yup, that’s Latin which was the standard language used for scientific publication into the 19th century.

Chapter 2

Dangerous misunderstandings

The first aim of this chapter is to convince you that intuitive thinking about probabilities is very error-prone. Very often, just when we're surest we understand the answer to a question involving uncertainty or randomness, our understanding is just plain wrong. I'll make this case by asking you a bunch of questions to which the answer is obvious and then explaining how the obvious answer is mistaken. The moral is that to get reliable answers to such questions we need some more formal techniques as a check on and corrective to our hunches.

Developing these ideas the goals of the pair of chapters—[CHAPTER 5](#) and [CHAPTER 4](#)—that follow and will call for a fair bit of work. The second goal of this introductory chapter is to convince you that this effort is worthwhile because many questions to which these tools apply come up in everyday life. You'll often be faced with such questions, and like it or not, the answers will be important to you. So understanding how to think about them is a skill you'll be able to use throughout your adult life. I've chosen here to focus on one question, "How dangerous is it to eat beef?" I'm not talking elevated cholesterol or a 24 hour case of food poisoning. I'm talking about acquiring a slow degenerative disease for which there is no cure or

2.1 A stroll through the minefield

effective treatment, that is invariably fatal but that that will turn you into a human vegetable long before it actually kills you. Is that dangerous enough for you?

When we're done, you'll understand much better how to analyze such questions. Next semester, when you've forgotten all these tools, I hope you'll still be able to recognize and avoid some of the most common pitfalls in thinking about probabilities. And, a decade from now, when you can no longer state the odds that a tossed coin will come up heads, I hope you'll still remember to distrust your instincts and hunches when trying to reach conclusions in the face of randomness.

2.1 A stroll through the minefield

Never forget that probability is a minefield. You always need to tread carefully when walking in it. The purpose of this short chapter is to convince you that such distrust of your intuition is the right attitude. The whole field of probability is mined and, when you venture out into it, you need to be on the alert at every step or you're likely to wind up with your best guesses blown to smithereens.

I also hope that the examples we'll look at will convince you that we need an up-armored vehicle to safely handle reasoning with uncertainty. The whole of [CHAPTER 5](#) is devoted to bolting on all the necessary reinforcing and we'll only come back to deal with probability in relative safety in [CHAPTER 4](#). So another goal of our stroll through the minefield is to motivate the effort it will take before we're ready to study probability in earnest.

We don't yet *know* any of the math of probabilities. But in the problems that follow, you won't need to know any. Just rely on your native intelligence to find the answers to any questions I ask. We'll then

2.1 A stroll through the minefield

test these answers, either by carrying out simple experiments with objects like coins and cards, or by analyzing some data from real world examples.

“\$7’ll get you \$12”

The game of “\$7’ll get you \$12” is a version of the classic short con, three card Monty, played by a grifter, who handles three cards, and a mark, who places bets. Traditionally, the 3 cards are two Kings and a Queen and the goal of the game is for the mark to “find the lady”—the Queen. The mark bets \$7 on a face-down card, and if the card she bets on is the Queen, she receives \$12 (\$5 plus her bet of \$7), if not she loses her \$7.

As it stands, this is too much of a sucker bet to attract even an AIG swaps trader, so the grifter shortens the odds as follows. To begin, the grifter shuffles the three cards face down and the mark puts her finger on one of the three cards. The grifter then looks at the two *other* cards and turns up, or *exposes* a King, reducing the number of face down cards to 2. Note that whether or not the mark has her finger on a King, the grifter can always find a King to expose. The mark then places her bet on either of the two unexposed cards.

PROBLEM 2.1.1:

- i) Suppose the mark places her bets at random. That is, half the time, she bets on the card she placed her finger on, and the other half of the time, she bets on the other unexposed card. Show that, on the average, the mark will lose \$1 each time she plays the game.
- ii) Does the mark’s betting strategy matter or will she always lose an average of \$1 each time she plays no matter how she places her bet?
- iii) If you were going to play “\$7’ll get you \$12” for money, would you choose to be the grifter or the mark?



2.1 A stroll through the minefield

Before we go on, I assume that we're all in agreement that in [ii](#)), the mark's strategy does not matter. The outcome is the same as in [i](#)). The mark loses \$1 each time she plays. Now we're going to check our predictions by actually playing some rounds of "\$7'll get you \$12".

EXPERIMENT 2.1.2:

i) Find a partner in your class. One of you will need to be the grifter in this experiment and the other will need to be the mark. If the roles you picked in [iii](#)) of [PROBLEM 2.1.1](#) are different, you can each take the role you prefer. If you both prefer the same role, just toss a coin to decide who takes what role.

You can use any pair plus a kicker of a different value (instead of a pair of Kings and a Queen)—the grifter exposes one card of the pair and the mark wins if she bets on the kicker. Play the game 30 times, with the mark deciding which unexposed card to bet on each time by flipping a coin: heads the mark bets on the card she fingered, tails, she bets on the other unexposed card. Record the amounts won and lost. How close are you to the prediction in [i](#)) of [PROBLEM 2.1.1](#)? Compare your results with those of other pairs in your class. Explain why the observed losses vary, staying close to but usually not exactly equaling the predicted value of \$30.

ii) Play the game 30 times, but this time with the mark betting on the card she did *not* finger all 30 times? How close are you to the prediction that the mark will lose about \$30 this time? Compare your results with those of other pairs in your class. Do you want to change your answer to [ii](#)) of [PROBLEM 2.1.1](#)?

iii) Predict what will happen if you play the game 30 times, but this time with the mark betting on the card she *did* finger all 30 times. Then play the game 30 times and compare your results with the rest of your class to check your prediction.

In "[\\$7'LL GET YOU \\$12](#)" [REVISITED](#), we explain the results of all the experiments you've run above. For now, in keeping with the spirit of this section, let's just agree that they're pretty convincing evidence

2.1 A stroll through the minefield

that the obvious answer to ii) is obviously wrong. One further point. This example shows that just having been shown some probability trap in the past is seldom solid protection against falling into it in the future. Many of you have already studied the game of **three card Monty** in high school and been bored at length with the explanation of the experiments you just ran. Hint: Monty's last name is Hall.

Trolling for terrorists

Britain's **National Security Strategy**, as laid out in a white paper for a blue-ribbon Commission on National Security contemplates “examination of the innocent as well as the suspect” by the “application of modern data mining and processing techniques”. While recognizing that “privacy issues” will arise, it concludes that: “Finding out other people's secrets is going to involve breaking everyday moral rules.”

Sound a bit sinister? Well, cheer up. You've been subject to this kind of invasion of privacy since 2006. In the US, it goes by the name of the **Automated Targeting System**. If you have travelled abroad, you've got a “risk assessment score”. This score, and the information used to derive it, can be shared with any government (federal, state, local, even foreign), used if you apply for a government job, license, or other benefit and even shared with private organizations and individuals doing business with the Federal government.

But don't worry, it's not available to just anybody. *You* are not allowed to see your score. In fact, you're not allowed to know what information was used to compile it, how your score is derived from that information, of when and how it will be used. For your further protection, there's no way you can challenge your score if you've been wrongly classified. Feeling safer now?

We know a bit more about a similar list, the Transportation Security Administration's No-Fly list. A copy which became public in

2.1 A stroll through the minefield

2006 contained 44,000 names and was the subject of a [60 Minutes segment](#). Many of these names posed little risk—because they were dead (among these were 14 of the 9/11 hijackers). Senator Ted Kennedy was on the list, as was Catherine Stevens, wife of Senator Ted Stevens, in both cases because someone with a similar name had been flagged. In Kennedy’s case it was a “T Kennedy” (Kennedy’s given name is Edward) and in Stevens’ it was the male pop singer Cat Stevens. Nelson Mandela was on the list. Children, some less than a year old, were on the list.

The good news is that the US Government has now consolidated all its watch lists in a single Terrorist Watch List managed by the [Terrorist Screening Center](#). The bad news is that *this* list now has [over a million entries](#) (probably a couple of hundred thousand more).

OK, so it’s pretty clear that the rules being used to add names to the Terrorist Watch List aren’t quite good enough. What I’d like to look into here is how good the rules would have to be for a watch list program like the British National Security Strategy or the US Terrorist Watch List to actually be an effective way of identifying potential terrorists. We’ll work with numbers for the British plan because it’s easier to find the necessary data.

To make such analysis, we need 2 key numbers. First, how large is the population being screened? The plan being proposed is to develop tools to screen the general population. In Britain, this comes to about 60,000,000 people. Second, how large is the number of terrorist suspects for which we are screening? The head of MI5, the agency responsible for tracking terrorists in Britain, [estimated this number](#), in the fall of 2007 as 2000, up by 400 in the last year. So let’s, to be safe, say we are looking for 5000 suspects, more than twice his estimate. To keep the numbers easy, let’s assume that we only need to screen 50,000,000 innocent people (by eliminating, say, the very young or old).

Now, we can assess what the results of a terrorist screening proce-



2.1 A stroll through the minefield

sure would be in terms of its error rate. Such a process can make two kinds of errors. First, it can fail to flag real terrorists, and second, it can label innocent people as terrorists. Stop for a moment and ask yourself which kind of error we need to be most worried about. It's natural to say the first. We're much more worried about an attack from an unknown terrorist we have failed to identify than we are about inconveniencing a few innocent people like Ted Kennedy.

The general wisdom from advertising and political campaigns that try to carry out similar screenings on a regular basis is that a test that is right on either count 90% of the time is a very good one. You'll flag 4500 actual terrorists (90% of the 5000 in Britain) and miss the other 500 (or 10%). We're not happy about missing these 500, But the thought that we can identify another 2500 suspects that we weren't aware of is what spurs interest in such a screening.

The problem is that we'll actually have almost *no* idea who those 4500 terrorists are! That's because our screening will clear 45,000,000 innocent people (90% of the 50,000,000 in Britain) but will flag the other 10%—that's 5,000,000 terrorist suspects! So even after our screening, our 4500 terrorists will be an invisible drop in a Watch List bucket of 5,004,500. Fewer than 1 in 1000 of the people on our list are actual terrorists. At this proportion, we certainly have no basis to lay any charges. We can't even sensibly monitor the people on our list: if we budgeted a billion dollars to do so, we'd just be totally wasting over \$999,000,000 of it!

What if we can devise a better test, say one that is 99% accurate? Nobody thinks that it's realistic to hope for accuracy this high, but that needn't stop us from asking what it would buy us.

PROBLEM 2.1.3: Show that a test that is 99% accurate will results in a watch list of 504,950 of whom only 4,950 will be actual terrorists.

So even in our wildest dreams, fewer than 1 person in 100 on our watch list would actually be a terrorist and if we spent a billion dollars on monitoring these people we'd still waste more than



2.1 A stroll through the minefield

\$990,000,000 of it. For a more formal discussion, see [PROBLEM 4.5.10](#).

What about the US Terrorist Watch List? It's hard to say anything too definite. I have not been able to find any reliable estimates for the number of terrorists targeting the United States: the numbers I have seen range from about 5,000 to about 20,000. Nor do we have any idea what the accuracy of the procedures used by the Terrorist Screening Center to place people on it are. But I'd be a lot more surprised to learn there were 50,000 terrorists on our Watch List than I would be to find out that it contains 1,000,000 innocent people.

The ideas that we have used to analyze watch lists are of great importance in many fields. We'll study applications to medical diagnosis and other problems when we deal with **false positives** and **false negatives** in [DIAGNOSTIC TESTING 4.5.11](#).

Lightning strikes twice

People have a very hard time distinguishing truly rare and unusual events from coincidences that are likely to occur fairly frequently by chance. Here I want you to think about some events that seem unusual and to estimate just how rare they are. Since we haven't learned how to compute such probabilities, I'll only ask multiple choice, "pick-the-right-answer" questions

EXAMPLE 2.1.4: To play a pick4 lottery—a common type of lottery operated by 32 states on a daily basis—you simply choose a 4 digit number, that is a number from 0000 to 9999. There are 10,000 such numbers and each is equally likely to be drawn on any given day. In other words, the chance that you'll win by playing any of the possible numbers (cost \$1) equals $\frac{1}{10,000}$. In most states, if you do you'll win \$5,000.



2.1 A stroll through the minefield

i) On September 9, 1981, the winning number in both the Massachusetts and New Hampshire pick4 was 8092. What was the chance of this happening?

a. $\frac{1}{5,000} = 2 \frac{1}{10,000}$

b. $\frac{1}{10,000}$

c. $\frac{1}{20,000} = \frac{1}{2} \frac{1}{10,000}$

d. $\frac{1}{100,000,000} = \frac{1}{10,000} \cdot \frac{1}{10,000}$

ii) What is the chance that the winning number in both the Massachusetts and New Hampshire pick4 on September 9, 2081 will be the same?

a. $\frac{1}{5,000} = 2 \frac{1}{10,000}$

b. $\frac{1}{10,000}$

c. $\frac{1}{20,000} = \frac{1}{2} \frac{1}{10,000}$

d. $\frac{1}{100,000,000} = \frac{1}{10,000} \cdot \frac{1}{10,000}$

iii) On September 10, 1981, several newspapers reported that both the Massachusetts and New Hampshire pick4 winning number for the previous day had been 8092 and said that the chance of this happening was $\frac{1}{100,000,000}$. Were they right about this probability?

Let's start with [i\)](#). We saw above that the chance that 8092 is drawn in New Hampshire is $\frac{1}{10,000}$. Likewise, the chance that 8092 is drawn in Massachusetts is $\frac{1}{10,000}$. Right away this tells us that [i\)a](#) is wrong: anytime 8092 comes up in both states, it's sure to come up in New Hampshire so the former chance must be less than the latter.

PROBLEM 2.1.5: Explain why the answer in [i\)c](#) must also be wrong.

Now let's turn to [ii\)](#) for a moment. Here it's easy to see what the right answer is. Let's suppose that on September 9, 2081, the number 8092 is chosen in New Hampshire. Since the chance the number 8092 will be chosen in Massachusetts is then $\frac{1}{10,000}$, we'll have a $\frac{1}{10,000}$ that the numbers match. What is some other comes up in New Hampshire, say 7615. The chance that 7615 will be chosen in Massachusetts is again $\frac{1}{10,000}$, so again there's a $\frac{1}{10,000}$ chance that the numbers match. In fact, no matter *what* number is chosen in New Hampshire, there's



2.1 A stroll through the minefield

a $\frac{1}{10,000}$ chance that the *same* number will come up in Massachusetts. So the right answer to **ii)** is **ii)b**.

Now I ask: are the questions in **i)** and **ii)** the same? No. The difference is that in **i)**, we want, not only that the two winning numbers are the same, but, in addition, that they match the specific number 8092. In the year 2081, having 8092 come up in both states is only 1 of 10,000 way the numbers can match. So the $\frac{1}{10,000}$ chance of they will match then must be 10,000 times bigger than the chance that they will both be 8092. Travelling back in time a century, the right answer to **i)** must be $\frac{1}{100,000,000}$.

How, then, can the answer to **iii)** be anything but “Obviously, yes”? Let me ask, and answer, a couple of even more obvious questions. Would it have merited a story if the Massachusetts pick4 number was 8092 and the New Hampshire number was 7615. Pretty much “Dog bites man”: I don’t think so. What if the numbers had been reversed? Still no story. What if the winning number in both states has been 7615? Now the man is biting the dog again. What was unusual about the draws on September 9, 1981 has nothing to do with the fact that the winning numbers were 8092. It was the fact the the *same* number was drawn in both States that was news; exactly *what* number was drawn is irrelevant. And the chance of seeing the same number in both states is the same on September 9, 1981 and on September 9, 2081: $\frac{1}{10,000}$ not $\frac{1}{100,000,000}$.

PROBLEM 2.1.6: On January 22 and 23, 2009 the same 3-digit number, 196 was picked two days in a row in the Nebraska PICK 3 lottery and an [article on this coincidence](#) circulated widely. Check out the pick3 winning numbers for August 10 and 11, 2006 at <http://www.nelottery.com/numbers.xsp>. How likely do you think it is that a “one in a million” event (as the article calls this coincidence) would happen twice in 1200 or so drawings? How likely *is* it that the same number will be drawn twice in a row in pick3?

2.1 A stroll through the minefield

PROBLEM 2.1.7: New York state holds two drawings daily for WIN 4, its name for a pick4 lottery.

- i) How likely is the same number to come up in both daily drawings?
- ii) How likely is the number 8056 to come up in both daily drawings?
- iii) Comment on the probabilities in [this Daily News article](#) from March 10, 2009.

One skill I fervently hope you'll acquire in **MATH4LIFE** is that of distinguishing between the truly rare and unpredictable, the unusual but predictable, and the merely coincidental. It's truly rare for any given *pre-assigned* number to come up twice in a day in the New York win4. So we'll almost certainly never see 8056 twice in one day *again*—having decided *now* to look for this number to come up twice. It's unusual for the *same* number to come up twice in a day in the New York win4 but we can predict that it will happen about once every 27.4 years on average. Since there are about 25 states with daily pick4 type lotteries, we can expect *some* such lottery to draw the same number twice in a row a bit less often than once a year. This happened in Pennsylvania on June 23rd, 2003. This event was also unusual but predictable. That the number that came up twice was 3199 was merely a coincidence, in hindsight. This would have been a truly rare event had we been watching, *before* June 23rd, 2003 for the number 3199 to come up twice.

AIG gives back: a fairy tale with a moral

Early in 2009, during the furor over the huge amount of Federal money (\$180,000,000,000) used to bail out its financial products division bailouts and bonuses (\$180,000,000) it had paid to employees in that division, the Board of Directors of insurer AIG decided that it would be smart to “give back to the American taxpayer”.



2.1 A stroll through the minefield

They decided to return 10% of the bailout money via a lottery. Every American citizen—the census bureau estimates that there are 227,719,424—would have a chance to win one of 18,000 prizes of \$1,000,000 in cash. (Yes, $18,000 \times \$1,000,000 = \$18,000,000,000$ which is only 10% of \$180,000,000,000). A computer was programmed to select American citizens at random using a combination of birth, tax and Social Security records. We'll assume that this task was performed perfectly—that is, on each draw every citizen had an equal chance of being picked. Because the number of prizes was so much smaller than the number of citizens (less than 1 in 10,000 citizens would be a winner), AIG felt that it was not necessary to take any action to prevent the same citizen winning more than 1 prize.

PROBLEM 2.1.8: Which of the following numbers is closest to the chance that some citizen's name would be drawn twice?

- i) 0
- ii) 0.00000000625 (This is $\frac{18,000}{227,719,424} \frac{17,999}{227,719,424}$.)
- iii) 0.00007904464 (This is $\frac{18,000}{227,719,424}$.)
- iv) 0.00010000000 (This is $\frac{1}{10,000}$.)
- v) 0.50000000000 (This is $\frac{1}{2}$.)

Can we rule out some of the answers above as obviously wrong? Well, we can rule out 0. The chance of drawing the same person twice may be small, but it's definitely a possibility if we don't try to prevent it, so the right answer is positive. In the other direction, $\frac{1}{2}$ seems crazy: how can there be a 50 – 50 chance of having a double winner when fewer than 1 in 10,000 is going to win? Likewise, 0.00010000000 seems pretty fishy because the fraction $\frac{1}{10,000}$ really has nothing to do with the drawing—I just mentioned it to give you a rough feel for the size of $\frac{18,000}{227,719,424}$.

That leaves 0.00000000625 and 0.00007904464. If, as we expect, there is no double winner, then there will be exactly 18,000 winners among the 227,719,424 citizens, so the chance of winning will be $\frac{18,000}{227,719,424}$. But the chance of winning twice must surely be smaller

2.1 A stroll through the minefield

than this. First you need to win one of the 18000 prizes, then you need to *also* win one of the 17999 *other* prizes. That gives us $\frac{18,000}{227,719,424} \frac{17,999}{227,719,424}$, so the answer must be 0.00000000625.

OK, I owe you an apology. The preceding paragraph is total nonsense, even if it sounded pretty reasonable going by. I just wanted to set you up for one of the most graphic demonstrations that intuition is a very unreliable guide to probability. The closest answer is 0.500000000000. The actual chance of having a double winner is 0.50903428737!! Yes, that's right, there's a *better* than 50–50 chance this will happen. Let me just remark that that's 5,000 times as big as 0.00007904464 and over 80,000,000 times as big as 0.00000000625, the answer that seems right. How's that for unreliable intuition?

BE SCEPTICAL ABOUT PROBABILITIES 2.1.9: *Moral: Be skeptical with probabilities. If you're not sure you understand it, think three times before believing it, and even if you are sure you understand it, think twice before believing it.*

That's all I'll say for now about this example, because I'd need to digress too long to explain how to calculate the number 0.50903428737 and why it's the right answer. If you can't wait to find out, you can peek ahead to the discussion to [EXAMPLE 3.8.42](#). Oh yes, and, of course, there was no AIG lottery. I just made this problem up to have an example where the chance of a double winner would come out close to $\frac{1}{2}$.

“We wuz robbed”

The title for this section was first uttered by boxing manager Joe “Yussel the Muscle” Jacobs (also the creator of “I shoulda stood in bed”) on June 21, 1932. That night, Max Schmeling had, in the view of most observers, outfought Jack Sharkey for the heavyweight championship but lost a split decision.



2.1 A stroll through the minefield

Here we're going to look at another robbery. It concerns the selection of playoff teams for the strike-shortened 1981 Major League Baseball season. Teams played differing numbers of games depending on how many of their games had been scheduled during the strike period, from June 12th to August 9th. For example, the Expos played 6 more games than the Cards. So the playoff teams that year were decided, not by the greatest number of wins, but by the highest winning percentage. Moreover, to create some extra playoff games and make back some of the money lost during the strike, two teams from each division would qualify for the playoffs. We're back in the days before wildcard teams when the division winners used to play each other for their league championship. Instead, the team in each division with the best record *before* the strike began would play off against the team with the best record *after* the strike ended.

You may notice a small problem with this format. What if the same team has the best winning percentage both before and after the strike? It was at first proposed to select the team with the second best overall record, but this was changed, on August 20th, to the team with the second best record in the second half. Here's a very easy set of questions about how this worked out.

PROBLEM 2.1.10: Below are the final standings (for the entire 1981 season) in the National League.

- i) Using TABLE 2.1.11, can you name with certainty all 4 of the 4 teams that qualified for the National League Playoffs? If so, what teams were they? If not, explain why it's impossible.
- ii) Using TABLE 2.1.11, can you name with certainty any 2 of the 4 teams that qualified for the National League Playoffs? If so, what teams were they? If not, explain why it's impossible.
- iii) Using TABLE 2.1.11, can you name with certainty any 1 of the 4 teams that qualified for the National League Playoffs? If so, what team was it? If not, explain why it's impossible.



2.1 A stroll through the minefield

iv) Using TABLE 2.1.11, can you name with certainty any 1 of the 8 teams that did *not* qualify for the National League Playoffs? If so, what team was it? If not, explain why it’s impossible.

Rank	Club	Wins	Losses	Percentage	GB
East Division					
1st	St. Louis Cardinals	59	43	.578	-
2nd	Montreal Expos	60	48	.556	2.0
3rd	Philadelphia Phillies	59	48	.551	2.5
4th	Pittsburgh Pirates	46	56	.451	13.0
5th	New York Mets	41	62	.398	18.5
6th	Chicago Cubs	38	65	.369	21.5
West Division					
1st	Cincinnati Reds	66	42	.611	-
2nd	Los Angeles Dodgers	63	47	.573	4.0
3rd	Houston Astros	61	49	.555	6.0
4th	San Francisco Giants	56	55	.505	11.5
5th	Atlanta Braves	50	56	.472	15.0
6th	San Diego Padres	41	69	.373	26.0

TABLE 2.1.11: 1981 NATIONAL LEAGUE FINAL STANDINGS

I hope you answered “No” to part i). For example, in the East Division, it’s pretty clear that one of the teams must be the Cards and the other must be either the Expos or Phillies. But which of the Expos and Phillies made it is too close to call. After all, they’re only separated by a single game, so which of them did better in the second half of the season is impossible to say “with certainty”. And whichever of the two that was would have qualified, either as the outright winner of the second half in the East, or as the second best second half team in the East.

Similar reasoning applies in the West. The Reds look like a lock but picking between “with certainty” between the Dodgers and the As-

2.1 A stroll through the minefield

tros is not possible, although since the Dodgers led overall by 2 games, you have a hunch they made it.

What about part ii)? I'll bet that many of you answered "Yes", the Cards and Reds must have qualified. After all, each has the best record in its division for the entire season. So each must have had the best record either in the first half or in the second half—if not in both. You can't be best overall without being best in at least one half. Any way you cut it, they'd both qualify. Right?

Wrong, as TABLE 2.1.12 shows. The Cards were best overall but only second best in each half, so the Phillies and Expos made the playoffs in the East.

Rank	Club	Wins	Losses	Percentage	GB
First half season					
1st	Philadelphia Phillies	34	21	.618	–
2nd	St. Louis Cardinals	30	20	.600	2.5
3rd	Montreal Expos	30	25	.545	4.0
4th	Pittsburgh Pirates	25	23	.521	5.5
5th	New York Mets	17	34	.333	15.0
6th	Chicago Cubs	15	37	.288	17.5
Second half season					
1st	Montreal Expos	30	23	.566	–
2nd	St. Louis Cardinals	29	23	.558	0.5
3rd	Philadelphia Phillies	25	27	.481	4.5
4th	New York Mets	24	28	.462	5.5
5th	Chicago Cubs	23	28	.451	6.0
6th	Pittsburgh Pirates	21	33	.389	9.5

TABLE 2.1.12: 1981 NL EAST STANDINGS BY HALF

I'll bet pretty much all of you answered "Yes", the Reds *must* have qualified to iii). After all, they don't just have the best record in the West, they have the best record in the entire National League. You probably felt, as before that you can't be best overall without being

2.1 A stroll through the minefield

best in at least one half. Well, we now know that’s not so, but so what? The Reds aren’t just a couple of games ahead like the Cards were, they’re a full 6 games ahead of the Astros. They *must* have made the playoffs. Well, read [TABLE 2.1.13](#) and weep. It was Dodgers versus Astros in the NL West in 1981. They’re still saying “We wuz robbed” in Cincinnati.

Rank	Club	Wins	Losses	Percentage	GB
First half season					
1st	Los Angeles Dodgers	36	21	.632	–
2nd	Cincinnati Reds	35	21	.625	0.5
3rd	Houston Astros	28	29	.491	8.0
4th	Atlanta Braves	25	29	.463	9.5
5th	San Francisco Giants	27	32	.458	10.0
6th	San Diego Padres	23	33	.411	12.5
Second half season					
1st	Houston Astros	33	20	.623	–
2nd	Cincinnati Reds	31	21	.596	1.5
3rd	San Francisco Giants	29	23	.558	3.5
4th	Los Angeles Dodgers	27	26	.509	6.0
5th	Atlanta Braves	25	27	.481	7.5
6th	San Diego Padres	18	36	.333	15.5

TABLE 2.1.13: 1981 NL WEST STANDINGS BY HALF

What we’ve seen in parts [ii\)](#) and [iii\)](#) is usually referred to as **Simpson’s paradox**. If you start with several component sets of averages (like the 2 sets of winning percentages in the first half and second half above), and aggregate them to get a single overall set of averages (the winning percentages for the full season), you can’t tell very much about the component averages you started with from the overall averages. Even an obvious overall best, like the Reds, needn’t be best in any of the components.

OK, so what about the worst teams. Well, Simpson’s paradox also

2.1 A stroll through the minefield

means that we can't be sure that the worst teams overall were worst in at least one half. As it happened they were, as [TABLE 2.1.12](#) and [TABLE 2.1.13](#) show. In fact the hapless Padres were not only worst overall but worst in both halves. But that wasn't what [iii\) of PROBLEM 2.1.10](#) asked. Can't we at least be sure that the worst teams did *not* make the playoffs? After all, the Cubs finished 21.5 games behind the Cards and the Padres 26 behind the Reds.

Rank	Club	Wins	Losses	Percentage	GB
First half season					
1st	Atlanta Braves				
2nd	Cincinnati Reds				
3rd	Los Angeles Dodgers				
4th	Houston Astros				
5th	San Francisco Giants				
6th	San Diego Padres				
Second half season					
1st	San Diego Padres				
2nd	Cincinnati Reds				
3rd	Houston Astros				
4th	San Francisco Giants				
5th	Los Angeles Dodgers				
6th	Atlanta Braves				

TABLE 2.1.14: 1981 NL WEST “STANDINGS” BY HALF

PROBLEM 2.1.15: Complete [TABLE 2.1.14](#) for the two halves of the NL West season so that:

- i) the total wins and losses for each team for the full season match those of [TABLE 2.1.11](#);
- ii) the total number of games played by each team in each half season matches that in [TABLE 2.1.13](#); and,

2.1 A stroll through the minefield

iii) the standings in each half season are as given.

Other than matching these totals, your table does not need to be realistic—and it won't be, because you'll have to make some teams like the Padres impossibly good in one half and impossibly bad in the other.

When you've finished, you'll have shown that the answer to part iv) of [PROBLEM 2.1.10](#) is “No”. At least mathematically, the Braves and Padres *could* have played for the 1981 NL championship. Hey, that's no stranger than putting a team from Atlanta in the NL West!

We'll see other examples of [Simpson's paradox](#), in less frivolous contexts, including a famous case dealing with discrimination in admission to graduate school in [SIMPSON'S PARADOX](#).

He's on Fire!

What is a random pattern? That question turns out to be very difficult to answer authoritatively. To keep things fairly simple, I'm going to restrict attention in this section to runs of heads and tails, which will write as a sequence of letters H and T, and which we can generate easily either out of our heads or by flipping a coin.

We don't want to work with sequences that are too short because no such sequence is really unlikely and there's no way to distinguish between randomness and order. For example, there are 4 sequences of length 2—we'll write such sequences without any angle brackets here, as HH, HT, TH and TT—each of which we expect to see $\frac{1}{4}$ th of the time if they appear randomly. In this section, I'm just going to ask you to take frequencies like this on faith: we'll see soon (in [EXAMPLE 3.3.5](#)) that there are 2^ℓ sequences of Hs and Ts with ℓ letters and that, if we choose at random, we expect each to occur 1 time in 2^ℓ .

On the other hand, we don't need to take really long sequences to study randomness. If, as we will, we work with sequences with 100 Hs and Ts, then each will come up $\frac{1}{2^{100}}$ th of the time. By the way,

2.1 A stroll through the minefield

$$\frac{1}{2^{100}} \simeq 7.8886 \times 10^{-31} \simeq 0.000000000000000000000000000078886.$$

So one answer to the question, “What runs of 100 Hs and Ts are random?” might be: they all are. This is not wrong, in that all such runs are equally likely—or better, unlikely. But it misses the point we’re trying to get at.

Consider the three sequences below:

[illegible]

H	T	H	T	H	T	T	H	H	H	T	H	H	H	H	T	H	T	T	T	H	H	H	H	H	T	H	H	T	H	H	H	H	T	H	H	T
T	H	H	T	H	T	H	H	T	H	T	H	T	T	H	T	H	T	T	H	H	H	H	H	T	H	T	H	T	H	H	H	T	H	T	T	H

H	T	T	T	T	T	T	H	H	T	H	T	T	H	H	H	H	T	T	T	H	H	H	H	T	H	H	H	T	T	T	H	T	H	T	T	H	H	T	H	T	H	H	H	T
T	T	T	H	T	T	T	T	T	T	H	H	H	H	T	H	T	T	T	T	H	T	H	T	T	H	H	T	T	H	H	H	H	T	H	T	H	H	T	H	H	H	T	H	H

Each of these is equally likely. They'll come up 1 time in $\frac{1}{2^{100}}$. But the first is clearly very unrandom. It contains *no* Ts. Since we expect a coin to come up tails about half the time, we'd be surprised, to say the least, to see *no* tails in 100 tosses. In fact, I'd be pretty sure the coin we were tossing had heads on both sides.

The second sequence looks more random, but it still contains 59 Hs and only 41 Ts. Weren't we expecting 50 of each? Not quite. We only expect the coin to come up tails *about* half the time. Going back to runs of length 2, we're not surprised to see no tails (or all tails) because HH (and TT) come up 1 time in 4. So we see exactly one T only half the time. When we have learned the [BINOMIAL DISTRIBUTION FORMULA 4.7.23](#), we'll see exactly 50 tails in 100 tosses a bit less than 8% of the time (but if you are interested you can check out [EXAMPLE 4.7.26](#) now). So *not* seeing 50 tails is the more common outcome. We'll see exactly 41 tails about 1.5% of the time. That's less than $\frac{1}{5}^{\text{th}}$ as often as we see 50 tails, but it's still not all that unusual.

The third sequence gives an impression somewhere in between. This time there are 45 Hs and 55 Ts. Those numbers come up almost 5% of the time, so are nearly as common as a 50-50 split. On the other

2.1 A stroll through the minefield

hand, look at the left side. There are two blocks of 8 Ts in a row, one at the left of each line. Moreover, 12 of the first 13 tosses on the bottom line are Ts. Shouldn't this make us suspicious that this run is not random? But how suspicious should we be? That's not so easy to say.

So no tails in 100 tosses is definitely not random. It's not so obvious, but even 20 tails is definitely not random: it happens less than 1 time in a billion. In fact, *any* number of tails less than or equal to 20 happens less than 1 time in a billion.

Where do we draw the line between a “mildly” unusual (41 tails) and a “wildly” unusual (20 or fewer)? Or between a “mildly” unusual run of 8 Ts in a row and a “wildly” unusual run of 100 Hs? The answer depends on how much confidence we want to have that what we're seeing is *not* random. We can never be *completely* sure: remember even 0 tails will *eventually* appear if we look at enough runs of 100 tosses. But since we'd expect to have to look at about 10^{31} such runs to see 0 tails once, we're *very* well-justified in deciding that something's fishy if we ever do see it. Likewise, something's definitely fishy if we see fewer than 20 tails.

A run like the second one above with only 41 Ts is right at the cusp. As we'll see when we apply the [CENTRAL LIMIT THEOREM 4.9.12](#) in [SECTION 4.9](#), a better question to ask is what's the chance of seeing *no more than* 41 Ts, and this turns out to be about 4.4%. In other words, we expect this to happen less than 5% of the time, or less than 1 time in 20. The figure 5% is very commonly used (for historical reasons) as the cutoff for outcomes that are unusual enough to make us seriously question their randomness. When you hear about experimental results at the 95% significance level, what is meant is that something was observed that we'd expect to see less than 5% of the time by pure chance.

As I said above, distinguishing the random from the meaningful is a very tricky problem, but one that arises in just about every discipline

2.1 A stroll through the minefield

in the modern world. The following experiment is designed to give you a first inkling of just how tricky.

EXPERIMENT 2.1.16: Below are two sets of empty boxes, in each of which you are to record a run of 100 random Hs and Ts.

- i) **Made-up method:** fill in the upper set of boxes by just writing down 100 Hs and Ts “out of your head”, trying to make the sequence look as random as possible.
- ii) **Coin-toss method:** fill in the lower set of boxes by tossing a coin 100 times, putting down an H each time you get a head and a T each time you get a tail.

[illegible]

Before going on, take a look at the run of Hs and Ts that you made up. Does it look random to you? If not, feel free to go back and change some of the letters to make it look more random. Please don't edit the boxes that you filled out by tossing a coin.

I'm now going to "Spot the quarter", that is, tell you which set of boxes you filled out and which were filled out by the coin. Since I won't be able to look at the Hs and Ts in your boxes, I'm going to ask you to help by computing a pattern number for each set of boxes. Then, I'll ask you to concentrate very hard on these two numbers for 30 seconds to communicate their random essence to me.

Here's how to find the pattern number. Find the greatest number of consecutive Ts in the top row. Then find the greatest number of consecutive Ts in the bottom row. Then repeat, but this time find the greatest number of consecutive Hs in each row. Add up the *squares* of these 4 numbers to get the pattern number. For example, for the boxes below, where the runs are shown in red, the pattern number is 66: 4^2 above and 5^2 below for the runs of Ts and 3^2 above plus 4^2 below for the runs of Hs.

2.1 A stroll through the minefield

H	T	H	H	T	H	T	T	T	H	H	T	T	H	T	H	H	H	T	T	T	H	T	H	T	H	T	T	T	T	H	T	H	H	H	T	T	T						
T	T	T	H	H	T	H	H	H	T	T	T	T	H	T	T	T	H	H	T	H	T	T	T	H	H	T	H	T	T	H	H	H	H	T	T	T	H	T	H	H	H	T	T

OK. When you have your two pattern numbers, write them down and concentrate hard on communicating them to me for 30 seconds. Thank you.

I can now reveal that the boxes you filled out of your head are the set with the smaller pattern number. What's going on here? Naturally, I cannot reveal the secrets of my psychic powers. That's mainly because I don't *have* any psychic powers. What I do know is that what people think is random is usually not.

In a case like the sequences of Hs and Ts, people think that long runs of consecutive Hs or consecutive Ts are not random. So sequences of Hs and Ts written down “out-of-your-head” to look random do not contain very long streaks of Hs or Ts. In fact, a random sequence of Hs and Ts, like that produced by tossing a coin, will have surprisingly long runs of consecutive Hs or consecutive Ts. When there are a total of 50 tosses there’s a run of 5 or more heads more than half the time (and likewise a run of 5 or more Ts more than half the time). More than half the time there’ll be at least one run of 6 or more. Without going into all the details, the upshot is that the pattern number of a random sequence of 100 Hs and Ts is almost always 100 or larger; the pattern number of a sequence generated by a person is almost never bigger than 75.

So my pattern number was just a gadget to “peek” at your sequences and see which contained long streaks of Hs and Ts: the one that has these was almost always generated by the coin and the one that doesn’t by you. I’d have been able to tell which was which by taking a real glance at your boxes, just by noting which set had the longer streaks in it.

One way to express the intuition about randomness that this example reveals is that people feel that the more consecutive heads we have tossed, the more we should expect a tail on the next throw. The belief that “runs tend to stop” is what causes us to avoid long streaks



2.1 A stroll through the minefield

when trying to simulate random sequences of Hs and Ts. But “runs tend to stop” is wrong. There’s a huge amount of evidence that “The coin has no memory”: regardless of *what* has happened on the previous tosses, we should expect the next one to come up heads (and tails) half the time. In other words, a run of (say) heads is *exactly equally likely* to continue and to stop on the next toss. In [SECTION 4.7](#), we look extensively at this very important property for which the general technical name is **independence**

What’s really odd about our intuition is that most people simultaneously believe even more strongly in the opposite fallacy: “runs tend to continue”. Listen to any broadcast of a college or professional basketball game and I guarantee that, at some point after a player has made 2 or 3 shots in a row, the announcer will invoke the “hot hand”. A player with a streak of made shots is *more than ordinarily* likely to make the next shot he attempts. So the rest of the team should “get him the ball”. Make another shot or two and you can expect to hear the dreaded cliché, “He’s on *fire!*”.

Wrong. Wrong, wrong, wrong. There’s no such thing as the hot hand. Entire seasons of shot by shot records of both field goals and free throws have been examined—the classic reference is *The Hot Hand in Basketball: On the Misperception of Random Sequences* by Thomas Gilovich, Robert Vallone and Amos Tversky [Cognitive Psychology **17**, 295-315 (1985)]—and the evidence solidly confirms that a shooter’s success on recent attempts has no influence on his chance of making his next shot.

I know you don’t believe me. The common argument is that shooters aren’t coins and that it’s the role of confidence in shooting that affects their success. Wrong. But, for now, let’s agree to disagree. We’ll come back to this topic when we have some more tools in hand and you can make your own experiments.

Chuck-a-luck

Chuck-a-luck or birdcage is the name of a carnival or midway game in which 3 dice are thrown by spinning a cage like that shown in FIGURE 2.1.17. In the most common variant, the rube bets on a number from 1 to 6 and wins the value of his bet for each die that comes up with the number bet showing, losing the bet if the chosen number does not show on any die.



FIGURE 2.1.17: A chuck-a-luck dice cage

In a striking example of convergent evolution¹, functionally equivalent games exist in China (Hoo hey how), England (Crown and anchor), Flanders (Anker en Zon), France (Ancre, pique et soleil) and Vietnam (Bầu cua cá cọp). The genetic trees of these games can be traced by comparing the dice used, as in the anchor common to the European versions. A Chinese set is shown in FIGURE 2.1.18. With the

¹Convergent evolution occurs when two genetically unrelated lines develop the same functional biological trait, as in the development of winged flight in bats, birds and insects. Usually, differences in implementation, as in this example, confirm the independence of the development. The most famous example is the camera eye which is “wired backwards” (nerves enter the front of the retina creating a blind spot) in vertebrates but forwards in cephalopods.

2.1 A stroll through the minefield

addition of many bells and whistles, the game is now found in most Asian and US casinos under the name **Sic bo**.



FIGURE 2.1.18: Hoo hey how dice

Suppose you bet on 4. Since a die has 6 sides, each of which has an equal $\frac{1}{6}$ chance of coming up on any roll, you can expect the first die to show a 4 $\left(\frac{1}{6}\right)^{\text{th}}$ of the time you play. Another way to put this is that you expect the first die to contribute $\left(\frac{1}{6}\right)^{\text{th}}$ of a 4 to your winnings. Likewise, you can expect each of the second and third dice to contribute $\left(\frac{1}{6}\right)^{\text{th}}$ of a 4 to your winnings. All told you expect to see $3 \times \frac{1}{6} = \frac{1}{2}$ of a 4 each time the case is spun.

Of course, there's nothing special about the number 4. Since all 6 numbers are equally likely to come up, you expect to see $\frac{1}{2}$ any number you bet on $\left(\frac{1}{6}\right)^{\text{th}}$ of all dice and hence to see $\frac{1}{2}$ of any of the 6 numbers each time the case is spun. We can check this by noting that 6 numbers times $\frac{1}{2}$ an appearance per spin gives an expected total of 3 appearances per spin for all 6 numbers which matches the 3 dice in the case. If we replaced $\frac{1}{2}$ with a larger or smaller number, we get either fewer than 3 or more than 3 expected appearances.

Another way to express the preceding argument is to say that for every \$1 you bet at chuck-a-luck, you expect to win 50¢, or to keep the numbers round every 2 times you bet \$1, you expect to win \$1. Of course, if you bet \$1 twice and win \$1 once, you'll also lose \$1 once. So, in the long run, rubes should expect to break even at Chuck-a-luck.

PROBLEM 2.1.19: Why do they call them rubes? Why is the carney spending 14 hour days spinning that dice cage? What do you really

2.1 A stroll through the minefield

think happens to the rubes when they play Chuck-a-luck?

That's right. The rubes lose money. On average, they will lose about $\$ \frac{17}{216} \approx 0.07870370370$ or roughly 7.87 cents for every dollar bet. I'll postpone explaining what the catch is until we have learned how to count the dice rolls in the game, but feel free to look ahead to the discussion after [EXAMPLE 3.8.36](#).

Let's note that, once again, the obvious answer about the game is wrong. You're sure the game is **fair**—meaning that you neither expect to win nor to lose—and it's not. That fact explains why the game is popular in so many different cultures. The rube gene seems to be one of the most widespread in the human population and Chuck-a-luck does a brilliant job of expressing it.

Instead, I want to use Chuck-a-luck to look at another question. How random are the losses of the rubes and the winnings of the operator? To make this question a little more concrete, let's consider several versions of it.

CHUCK-A-LUCK SCENARIOS 2.1.20:

- i) You like to gamble and Sic Bo is your game. We want to consider your chances of beating Dump Casino's house edge in a couple of scenarios.
 - a. Suppose you visit the Dump Casino in Atlantic City on August 25, 2008 and bet \$1 at at Sic Bo 100 times. What is the chance that you'll walk away a winner on the day?
 - b. Suppose you visit the Dump Casino in Atlantic City twice a week for a year—let's call this 100 visits so you can take a 2 week vacation—and bet \$1 at Sic Bo 100 times on each visit. What is the chance that you'll walk away a winner on the year?
- ii) You're the owner of a string of Chuck-a-luck booths each of which you book on a state fair circuit for 100 days each year and lease to operators. Each day, each operator reports the net amount of money won on bets placed at his booth and you split this profit



2.1 A stroll through the minefield

50–50. Let's assume, to keep things simple, that each day each booth covers exactly 10,000 \$1 bets.

- Suppose that one of your operators reports that, on August 25, 2008, his booth won \$675. Does this indicate a problem with the operator or his booth?
- Suppose that one of your operators reports that, on August 26, 2008, his booth won \$443. Does this indicate a problem with the operator or his booth?
- Suppose that over the course of a season, one of your operators reports that his booth won \$76,034. Does this indicate a problem with the operator or his booth?
- Suppose that over the course of a season, one of your operators reports that his booth won \$84,741. Does this indicate a problem with the operator or his booth?

Before we start to try to answer these questions, a few comments. First, despite the frivolous context, questions of this type are very important. Other questions of this type—in the sense that giving answers requires applying exactly the same mathematical ideas to different observations—include:

- Does smoking cause lung cancer?
- Is the recent change in average global temperatures due to rises in carbon dioxides levels in the atmosphere or to random climatic variations?
- Is the SAT culturally biased?
- Does adding omega-3 fatty acids to your diet reduce your risk of heart disease?
- Can the mutual fund managers consistently outperform stock market indexes?
- Does pledging sexual abstinence reduce your risk of acquiring STDs?

So, let's try to get a bit more of a bird's eye view of what it is we are asking. In each of the [CHUCK-A-LUCK SCENARIOS 2.1.20](#), we have



2.1 A stroll through the minefield

a probability prediction of what should happen, based on the knowledge that rubes lose 7.870370370 cents for every dollar bet. Let's work out these expectations.

PROBLEM 2.1.21: Calculate the expected loss of the player and the expected winnings of the operator for each of the **CHUCK-A-LUCK SCENARIOS 2.1.20**.

Partial Solution

I'll work out **i)a** and **ii)c** and leave the other parts to you. In all cases, the calculation is very easy: just multiply the total number of \$1 bets placed by the expected win or loss of \$0.0787 per dollar bet and you'll have the expectation.

In **i)a**, there are 100 bets so the expectation is that the rube will lose *about* $\$100 \times 0.0787 = \7.87 .

In **ii)c**, there are 10,000 bets on each of 100 days, making 1,000,000 bets in all, so the expectation is that the operator will win *about* $\$1,000,000 \times 0.07870370370 = \$78,703.70$.

OK, so we have a clear idea of what we expect, but there's kicker. Despite all the decimal places and the apparent accuracy, these expectations are only approximate. The kicker is the italicized word *about* that we need to insert because we are dealing with probabilities. It's the same about as the one in the sentence: "If we toss a coin 100 times, it will come up heads *about* 50 times." When we say that, we do *not* mean that we think exactly 50 heads is all that likely; as I mentioned above in discussing randomness in **HE'S ON Fire!**, that'll happen only about 8% of the time. What we mean is that we expect the number of heads to be *close* to 50 *most* of the time. But sometimes, we'll see a lot more heads—even 100 heads although only 1 time in about 10^{31} . Saying how near to 50 heads is *close* and how often is *most* of the time is not so easy. Are the 45 heads and the 59 heads that we saw in random examples in **HE'S ON Fire!** *close*? How often can we expect numbers this far from our expectation of 50? When is a value far enough from what we expect for us to say "I



2.1 A stroll through the minefield

smell smoke” and when is it far enough away to yell “Fire”.

The questions posed in [CHUCK-A-LUCK SCENARIOS 2.1.20](#) are questions of this type. We are comparing an particular observation (in [He's ON Fire!](#), these were the particular runs of 100 heads and tails) to an expectation (we expect to see about 50 heads), and we ask: “How likely is it that a random observation would deviate from our expectation as much as the particular observation does?” If the likelihood is big enough, we say the deviation is due to randomness. An observation like 45 heads falls squarely in this range. As that likelihood get smaller, we start to smell smoke: “I suspect that my particular observation is *not* random.” In the discussion following, [UNLIKELY SUCCESSES RULE-OF-THUMB 4.9.20](#), we'll see that an observation like 64 heads smells distinctly smoky. When the likelihood gets small enough, we yell fire: “My observation *isn't* random; if it is, I'll eat my hat.” Given an observation like 31 heads, we see flames. At 81 heads, we bring in the arson squad. But things are seldom completely cut and dried. An observation like the 59 heads that appeared in of our random runs is almost, but not quite, in the smoky range.

Let's rephrase the questions in [i](#)) in these terms. In [i\)a](#), our expectation is that you'll lose \$7.87. If you come home a winner on August 25th, then we've observed a loss that deviates by more than \$7.87 from our expectation. How likely is a deviation this big? In [i\)b](#), our expectation is that you'll lose \$787.03 on the year. If you are a net winner over the year, then we've observed a loss that deviates by more than \$787.03 from our expectation. How likely is a deviation this big?

Statistics is the science of finding ways to answer this, and similar, questions. It's a huge subject and we'll just learn a few of its most basic lessons, at the end of our study of probability, in [SECTION 4.9](#). These tell us how to estimate such likelihoods in terms of the number amount of data in our observation (the number of tosses of the coin, or the number of spins of the Chuck-a-luck cage) and a key



2.1 A stroll through the minefield

number, called the variance of our problem. I'm not even going to try to define the variance roughly here, except to say that it tries to capture how far from our expectation a typical observation will be.

Once we have learned a bit of statistics, we'll be able answer questions like those in the [CHUCK-A-LUCK SCENARIOS 2.1.20](#) about deviations from expectations almost as easily as we were able to determine what the expectations were above. Warning: don't take this to mean that statistics is easy. Far from it; it's simply that we're only going to look at a tiny corner of the subject that is easy. The kind of answers we'll be able to give will be deductive, meaning that we'll just have to study problems mathematically to reach our conclusions.

There's also an **inductive**, or observation based, approach to deciding how likely the deviations in the [CHUCK-A-LUCK SCENARIOS 2.1.20](#) are. We simply make the observation in each question *a large number of* times, and then see what fraction of our observations exhibit the deviation from our expectation asked about. For example, consider the observation or **experiment** of playing Sic bo 100 times, as in [i\)a](#), and recording the net win or loss for the rube. Suppose we repeat this observation n times and see p occasions in which the rube went home a winner. Then, the fraction $\frac{p}{n}$ gives us an estimate of how often we can expect the rube to win, and hence an approximate answer to the question in [i\)a](#). If n is reasonably large, we expect this estimate to be reasonably accurate.

In a moment, we going to look at a set of $n = 100$ such observations for each question, which is enough to give a good feel the answers to the [CHUCK-A-LUCK SCENARIOS 2.1.20](#). For this number of observations, we can rephrase our conclusion as: the rube wins about $p\%$ of the time.

There are a couple of draw backs to the inductive approach. First, it's a lot of work. To get each observation we need to make 100 spins to get one observation, and hence $100 = 10,000$ spins to find p . In [ii\)c](#), each observation calls for 1,000,000 spins so finding p would call for

2.1 A stroll through the minefield

100,000,00 spins. A homework assignment like that might make for a couple of pretty late nights. One of the challenges facing statistics is to find ways to reach reliable conclusions *without* having to take a very large number of observations. Being able to do this well is of immense practical significance, because in many problems—think of testing medical procedures—it is very expensive to gather even a single observation.

I have two pieces of good news for you. First, our look at statistics won't reach the difficulties of efficient testing of hypotheses. Second, I've done all the experiments needed. I wrote a small computer program to simulate chuck-a-luck spins; then I used this program to produce 100 observations of each of the [CHUCK-A-LUCK SCENARIOS 2.1.20](#). The 100,000,00 spins needed for parts ii)c and ii)d took 13 hours on my iMac™. All you'll have to do is analyze the resulting observations.

The second drawback to the inductive approach is that the answer we'll get is itself somewhat random. Suppose we carry out a second set of 100 observations of the rube in the casino, and see the rube come home a winner q times. There's no guarantee that p and q will be equal. In fact (as with the 50 heads in 100 tosses) this is usually rather unlikely. What we expect is that p and q will be fairly close to each other. But how close is close? And if they're not close, which one should we trust and which should we reject? If we *don't* repeat our observations, how do we know that the p we observed isn't a wildly unlikely value? These questions are, as I hope you now see, just further statistical questions about our collections of observations. Once again, we could try to answer them either deductively (if we are prepared to learn some more statistics) or inductively (if we have a lot of free time for experimenting).

More good news. I'd like you to remember is that even inductive answer based on a lot of observations may be misleading. This is one of the main reasons that there's no substitute for deductive statisti-



2.1 A stroll through the minefield

cal methods. But you do not need to worry about the variability of the observations that I'll provide you with. Just trust me that they are reasonable and draw the conclusions they suggest. I'll periodically tell you what the deductive methods say about various likelihoods. Again, just trust these figures for now and we'll see how to arrive at them in [PROBLEM 4.9.26](#).

Each cell in [TABLE 2.1.22](#) shows the net winnings (or, if negative losses) of the rube in 100 Sic Bo spins. There are also 100 cells or observations: the cells are in the order that I made the observations in reading order (from left to right and top to bottom). Just by glancing at the table you can see that the rube does not go home a winner too often. On the other hand, there are cells when he does go home a winner by a fair bit—by \$17 on two occasions.

-5	-26	8	-19	-20	-17	-29	4	-7	0	-12	-7	3	-20	-7	-11	-5	-16	-3	-12
-8	-7	-8	-1	-9	-17	-20	-12	1	-20	-17	-18	-7	-4	2	-6	0	1	5	2
-22	-10	-31	9	2	-20	7	-9	-13	17	10	-9	-10	-35	-5	-2	-2	-14	8	-2
-11	-24	-1	-12	6	7	-24	-4	-6	-23	-9	-1	17	-1	14	-2	-11	-6	16	0
-4	-18	-5	-35	-11	11	-4	-9	-18	-2	-14	-19	-8	-17	-15	11	-11	-29	-3	-2

TABLE 2.1.22: 100 RAW OBSERVATIONS FOR PART i)a

[TABLE 2.1.23](#) contains the same data as [TABLE 2.1.22](#), but the numbers have been sorted into increasing order to make it easier to answer the question, “How often does the rube go home a winner after a night of Sic Bo?” The answer is that he wins 21 nights out of 100 and wins or breaks even 24 times.

-35	-35	-31	-29	-29	-26	-24	-24	-23	-22	-20	-20	-20	-20	-20	-19	-19	-18	-18	-18
-17	-17	-17	-17	-16	-15	-14	-14	-13	-12	-12	-12	-12	-11	-11	-11	-11	-11	-10	-10
-9	-9	-9	-9	-9	-8	-8	-8	-7	-7	-7	-7	-7	-6	-6	-6	-5	-5	-5	-5
-4	-4	-4	-4	-3	-3	-2	-2	-2	-2	-2	-2	-1	-1	-1	-1	0	0	0	1
1	2	2	2	3	4	5	6	7	7	8	8	9	10	11	11	14	16	17	17

TABLE 2.1.23: 100 SORTED OBSERVATIONS FOR PART i)a

[TABLE 2.1.24](#) contains a second set of 100 “night in the casino”. Once again the rubes winnings have been sorted into increasing order.

2.1 A stroll through the minefield

Now, he only wins 20 nights out of 100 but he wins or breaks even 26 times. Both sets of observations suggest that the rube wins on about 20% of the nights he plays and breaks even on about 25%.

-35	-33	-32	-30	-29	-28	-26	-26	-25	-25	-22	-21	-21	-21	-20	-20	-20	-20	-20	-20
-20	-19	-19	-19	-19	-19	-19	-17	-17	-17	-16	-14	-14	-14	-13	-13	-13	-13	-13	-12
-12	-12	-12	-11	-11	-11	-10	-9	-9	-9	-7	-6	-6	-5	-5	-5	-4	-4	-4	-4
-4	-3	-3	-3	-3	-3	-2	-2	-2	-2	-2	-2	-1	-1	0	0	0	0	0	0
1	1	1	2	3	3	3	3	4	4	4	5	5	5	7	8	8	11	13	14

TABLE 2.1.24: 100 MORE SORTED OBSERVATIONS FOR PART i)a

Looking at the 2 sets of observations, it’s pretty clear that the rube didn’t do as well in the second set as in the first. The numbers in corresponding cells are generally a little more negative in the second table than in the first. You can confirm this by summing the numbers. In the first table, they sum to -752 and in the second to -898 . That’s a difference of $\$146$. Another way to put these totals is to saw that in the first table the average loss per night is $\$7.52$ —close to, but a bit less than our expectation of $\$7.87$ —while in the second the average loss is $\$8.98$ over a dollar more than we expected.

This differences illustrate what I meant when I said that inductive answer like those in the table are themselves somewhat random. The statistic we were trying to focus on—how many winning nights were there—was quite close in the two tables. But, the two average losses differ substantially.

Before I present 100 observations for i)b of the CHUCK-A-LUCK SCENARIOS 2.1.20, note that you already have 2 in hand. We’ve been thinking of TABLE 2.1.23 and TABLE 2.1.24 as 100 observations for i)a—the rube plays for one night—but each can be thought of as *one* observation of the rube playing for a whole year. Notice also that these 2 observations don’t really offer much guidance in answering i)b. Yes, the player loses in both observations, $\$752$ and $\$898$. But if I picked 2 cells at random in one of these tables, I’d probably get 2 negative numbers ($\frac{3}{4}$ of the cells are negative). Yet the player seems

2.1 A stroll through the minefield

to win about 20% of the time. So we really need lots of tables like the two above to answer i)b.

TABLE 2.1.25 gives a sorted summary of the data in 100 tables like the 2 above. Each cell of this table contains the total won or lost by the rube in 100 nights of Sic Bo play, so each *cell* summarizes an amount of data equal to *all* of one of the tables above. For example, the -898 in the top row stands for as much data as all of TABLE 2.1.24.

-1114	-1103	-1077	-1031	-1031	-982	-963	-963	-958	-946	-942	-922	-905	-902	-902	-901	-899	-898	-895	-889
-887	-882	-878	-873	-872	-869	-868	-859	-858	-855	-850	-848	-845	-843	-841	-841	-839	-833	-826	-824
-824	-824	-819	-817	-816	-816	-816	-815	-812	-806	-805	-805	-800	-789	-788	-787	-786	-784	-782	-779
-776	-765	-762	-758	-755	-752	-751	-746	-746	-737	-736	-727	-726	-723	-722	-722	-721	-720	-719	-710
-709	-703	-701	-697	-691	-687	-678	-668	-666	-664	-658	-654	-645	-578	-574	-570	-521	-517	-516	-493

TABLE 2.1.25: 100 SORTED OBSERVATIONS FOR PART i)b

What strikes you about this table? All the entries are negative; quite negative. The rube never ends the year a winner. He doesn't even come close. If he goes to Dump 100 times a year for a century, he won't break even in a single year.

The total of the 100 cells in the table is $-\$79,748$. and so the rube's average loss over 100 nights—the average value of a call—is $\$797.48$. We expected the rube to lose an average of $\$787.30$ and now we're getting pretty close. If we divide the *entries* in the table by 100 to get the corresponding average loss on the 100 nights at Dump that each cell summarizes, we see that these average nightly losses range from a low of $\$4.93$ to a high of $\$11.14$. That range is much smaller than the ranges of nightly losses we saw in the previous tables (which were between $\$35$ and $-\$17$ (where we're now viewing a win of $\$17$ as a loss of $-\$17$).

That last observation is the most important point about this example. We expect each the losses on an average night to be about $\$7.87$. But lots of individual nights are far from average: we see nights that $\$25$ above and night that are $\$25$ below this figure. When we take

2.1 A stroll through the minefield

average of groups of individual nights, we still see some variability. But this variability is much smaller. Our best and worst averages are now less than \$4 away from our expectation.

We can also use TABLE 2.1.25 to start to answer ii)a and ii)b. Each cell in TABLE 2.1.25 summarizes 100 nights of 100 spins of Sic Bo. But we can also think of each cell as summarizing 10,000 spins of chuck-a-luck, or one day’s action at a state fair.

PROBLEM 2.1.26:

- i) How many cells in TABLE 2.1.25 are smaller than 675? Estimate the chance of seeing a days winnings equal to or smaller than \$675?
- ii) How many cells in TABLE 2.1.25 are smaller than 443? Estimate the chance of seeing a days winnings equal to or smaller than \$443? Would you be suspect your operator of falsifying his report in either case?

Let me ask that question again, showing you a different set of observations. TABLE 2.1.28 was generated by the same random process as TABLE 2.1.25. Each cell summarizes 10,000 simulated spins of chuck-a-luck. The total of the cells in this table is −80152 meaning that the operator averaged winnings of \$802.52 a day, a few dollar more than in the previous table.

PROBLEM 2.1.27:

- i) How many cells in TABLE 2.1.28 are smaller than 675? Estimate the chance of seeing a day’s winnings equal to or smaller than \$675?
- ii) How many cells in TABLE 2.1.28 are smaller than 443? Estimate the chance of seeing a day’s winnings equal to or smaller than \$443? Would you be suspect your operator of falsifying his report in either case?

-1128	-1076	-1029	-1023	-1012	-987	-972	-972	-970	-963	-953	-950	-949	-945	-937	-933	-921	-911	-897	-894
-893	-888	-887	-887	-885	-871	-868	-862	-861	-861	-858	-856	-849	-848	-844	-842	-838	-838	-833	-832
-830	-827	-818	-809	-808	-804	-800	-798	-798	-796	-796	-794	-781	-780	-776	-776	-776	-773	-772	-772
-772	-770	-768	-761	-760	-757	-755	-750	-744	-740	-735	-734	-731	-727	-726	-726	-722	-721	-719	-716
-713	-709	-709	-706	-705	-697	-697	-694	-687	-687	-685	-643	-641	-623	-620	-585	-575	-562	-512	-361

TABLE 2.1.28: 100 SORTED OBSERVATIONS FOR PARTS ii)a AND ii)b

2.1 A stroll through the minefield

This example illustrates what makes statistics hard and why we need to supplement inductive methods with deductive ones. Looking at [TABLE 2.1.25](#), I presume you said that winnings as low as \$443 must be very rare. We only saw 1 day in 100 below \$500.

But then we get [TABLE 2.1.28](#). One of the 100 numbers in that table is -361 which is not only a lot smaller than -500 , but a lot smaller than -443 . An observation this far from its fellow is usually called an **outlier**. What should we conclude? Do numbers as small as -361 come up close to once in every 200 observations? Or did we just happen to observe something unusual? There's no easy way to tell. Either we need a lot more observations, which means a lot more work, or we need some theory. I hope this example will convince you that that latter is the better answer.

It turns out that the deductive theory shows that winnings less than \$500 can be expected less than 1 time in 100. Winnings as low as \$443 occurs less than 1 time in 700. So we might be very suspicious of the operator who reported this figure but realize that his report is not wildly unlikely. When you look at 100 observations, seeing a 1 in 700 value is surprising, but only mildly so. The number -361 really is unusual. We expect to see an observation that small only about 1 time in 10,000. Put more graphically, suppose we looked at 100 tables like the 2 we have: how often would we expect to see an entry this small? Most of the time we'd see either exactly 1 or 0, and only rather rarely fairly often would we see more than one. So it's surprising that we did observe the number 361 in looking at just 2 such tables.

But note also that we can't just throw an outlier away because it messes up our tables and annoys us. We'll also see how important respecting the data is to drawing sound statistical conclusions. I was shocked when my program spit out that -361 and I went back and did a few more checks of my simulator. But I found no problems. So I swallowed hard and passed the number on to you. If I saw an

2.1 A stroll through the minefield

outlier than far off *again*, I would much more seriously question my program. Right now, I'm more worried that some reviewer will think I deliberately inserted the number -361 to set off this discussion of outliers.

PROBLEM 2.1.29: We can also view [TABLE 2.1.28](#) as a second set of 100 yearly observations of the rube in [i\)b?](#) Should the outlier -361 make us revise our conclusion that the rube will never come home a winner at the end of the year?

With all these cautions, we're ready to see the striking effects of accumulating large numbers of observations. Each cell in [TABLE 2.1.31](#) below summarizes 1,000,000 spins of chuck-a-luck: that's as much data as in all of [TABLE 2.1.28](#). Notice, once again, how much more "predictable" all the numbers in this table are. We expected the total loss in 100,000,000 spins to equal $-7,870,370.37$ and we observe a total loss of $-7,887,034$ (the total of all the cells). Observation and expectation now match to 2 places (almost to 3). If we divide each cell by the 1,000,000 spins we get numbers ranging from $-.081928$ to $-.075908$, all of which are clustered near the expected value of $-.078703$.

PROBLEM 2.1.30: Use the data in [TABLE 2.1.31](#) to answer parts [ii\)c](#) and [ii\)d](#).

-81928	-80899	-80820	-80568	-80552	-80513	-80431	-80429	-80286	-80274
-80217	-80198	-80129	-80108	-80038	-79995	-79965	-79954	-79936	-79933
-79929	-79861	-79850	-79847	-79807	-79755	-79737	-79611	-79570	-79557
-79551	-79495	-79465	-79456	-79425	-79381	-79358	-79326	-79237	-79236
-79234	-79176	-79167	-79160	-79102	-79068	-79044	-79042	-79004	-78995
-78994	-78913	-78867	-78831	-78828	-78797	-78776	-78756	-78705	-78700
-78657	-78519	-78490	-78442	-78409	-78395	-78387	-78383	-78126	-78040
-78008	-77981	-77937	-77931	-77928	-77902	-77898	-77832	-77824	-77771
-77759	-77755	-77736	-77728	-77702	-77694	-77672	-77667	-77630	-77448
-77445	-77410	-77361	-77307	-77196	-77079	-77032	-76704	-76155	-75908

TABLE 2.1.31: 100 SORTED OBSERVATIONS FOR PARTS [ii\)c](#) AND [ii\)d](#)

Let me close by quickly discussing [PROBLEM 2.1.30](#). Winnings of

2.1 A stroll through the minefield

\$76,234 are observed twice in our table to this figure is a bit unusual but not wildly so. In the long run, you'd see winnings this low about 1 year in 60. The closest we come to seeing winnings of \$84,741 in our observations is the \$81,928. Is that extra \$1,800 or so significant? Almost certainly.

There are 61 observations within 1000 away from the expectation of -78703.7 and 39 further away, but only 5 more than 2000 away and only 1 just over 3000 away. Winnings of \$84,741 would be about 6000 away from the expectation—twice as far as any of our 100 observations, and that's far enough to yell "Fire". A value this far away can be expected in fewer than 1 in every 10,000,000 observations.

If your operator reported winnings that were low by this much, you'd suspect him of holding out on you. Perhaps you're not worried about making a few thousand more than you expected. But you should be. That \$84,741 tells you that something is not right—and it's not the laws of probability—so you'd be smart to investigate that booth.

The very important moral here is called the:

LAW OF LARGE NUMBERS 2.1.32: *Even outcomes that are individually random become predictable when you look at large collections of them.*

In an evening of Sic Bo, the rube can often go home a winner (about 20% of the time) even though we expect him to lose. In a year of playing the game, he'll always end up a loser, although his total losses can vary substantially (by a factor of 2 or more). Similarly, a chuck-a-luck booth operator can always expect to be up at the end of a day, but his net winnings can vary a lot. But we can predict, with very high certainty that over a season of running a booth, his winnings will not vary more than 5% percent (i.e. \$4,000) from the expected \$78,703.70. We'll see many more examples later in the course.

2.2 You want prions in that burger

One of the roles of the US Department of Agriculture, or **U.S.D.A.**, is to ensure the safety of meat consumed by Americans. It does this through an Animal and Plant Health Inspection Service or APHIS program whose mission is “To protect the health and value of U.S. agriculture, natural and other resources”.

In recent years, the disease of greatest concern has been **BSE** or **bovine spongiform encephalopathy** common know as **mad cow disease**. This is a disease of cattle, humans and other mammals generally felt to be caused by misfolded proteins or prions. Such prions, although much simpler even than viruses, have the ability to reproduce in mammalian brains causing the gradual degeneration of the brain (encephalopathy) and spinal cord into a spongy (spongiform) mass—hence the name. In humans, the disease is more commonly called **Creutzfeldt-Jakob disease** or **CJD** but here I’ll call all forms BSE for simplicity.

The progress of the disease is if often very slow, especially in humans. It often takes a decade or more to declare itself by progressive dementia (memory loss and hallucinations) and physical impairment (jerky movements, rigidity, ataxia and seizures) and another before death ensues. There is currently no cure or effective treatment for BSE/CJD; the disease is invariably fatal.

The most common route of infection is the ingestion of prions contained the flesh of other infected mammals, although a variant called scrapie that infects sheep is conjectured to be transmitted through both urine and maternal milk. Cattle infections occur mainly when feed contains by-products from carcasses of sheep, goats and other cattle. Although such by-products are usually cooked as part of the feed manufacturing process, ordinary cooking neither destroys nor

2.2 You want prions in that burger

disactivates the offending prions. Most countries now ban the use of such “ruminant-to-ruminant” feed.

Infection due to cannibalism was common until recently in areas such as Papua-New Guinea where the disease is known as kuru. But in the last 30 years, most infections in humans have been caused by consumption of infected beef. As with cattle feed, even very well-cooked beef that contains prions remains infectious. It is estimated that several hundred thousand infected cattle have entered the human food chain since the 1980s, but little good data exists and the figure could easily be an order of magnitude (10 times) lower or higher. An outbreak of BSE/CJD in Britain in the 1980s led to the slaughter of hundreds of thousands of animals and the banning of importation of British beef. This infection has since led to the death of at least 150 people but, given the long incubation period, the ultimate toll is difficult to estimate.

So, it seems urgent to ask, the

BSE QUESTION 2.2.1: *How many US cattle are infected with BSE?*

I’d sure like to know and I bet you would too. It would be easy to answer this question. Although it is currently difficult to test living animals for BSE, examination of the brain and spinal cord tissues after slaughter is a reliable and easy method of identifying infected meat. Some countries, notably Japan, perform this test on every carcass intended for human consumption. However, as we’ll see at the end of [BSE TESTING](#), it’s possible to screen only a small fraction of slaughtered animals and still say that with a high degree of certainty that rates of infection are lower than pretty much any pre-determined threshold we’d like to set. Testing programs in Europe, Canada and most of Asia are of this type with the threshold typically set at about 1 in 1,000,000.

Unfortunately, it’s also possible to design a program that tests tens of thousands of carcasses and is likely to find *no* infections even if infection rate is as high as 1 in 100,000. It turns out that



2.2 You want prions in that burger

In this section, I want to convince you that the APHIS screening program in the US is deliberately designed **not** to answer the [BSE QUESTION 2.2.1](#).

That's right. The U.S.D.A. does *not* want you to know how many cattle in the US have BSE and has designed a screening program that is pretty much guaranteed not to discover this information. Why? Because, if *you* find out, so will the authorities in the many countries that import large quantities of US Beef. Even if the answer turns out to be “very, very few”, recent history shows that what happened to exports of British beef in the 1980s will happen to ours: they'll be prohibited. So to protect US beef exports, you must remain in the dark.

Before we look at the program, let me say that my personal hunch is that rates of BSE infection in US cattle are probably pretty low and that your risk of acquiring CJD is very low indeed. But I don't really know. I want to try to convince you that you don't either and to convince you that the reason we don't know is because APHIS has not done the tests that would answer the [BSE QUESTION 2.2.1](#). To see this, we need to consider two further questions.

Our first question is, “What is an acceptable rate of *exposure* of Americans to BSE infected cattle?” The first reaction to this question is to say, “What are you, a moron? There's *no* acceptable rate of exposure to BSE. Only *zero* exposure is acceptable.” The U.S.D.A. would agree with you. When a *single* case of a BSE infected cow was discovered in Canada in mid-2003, it banned all imports of Canadian beef into the US. Deboned beef from younger cows was allowed back in in 2004. (Cattle less than 30 months old do not show BSE pathologies, although, given the often slow progress of the disease, it's not completely clear that they pose no risk.) But imports of live cattle were only renormalized in November 2007. By that time, several more cases of BSE had been found in Canada, but these all affected cattle born before Canada banned “ruminant-to-ruminant” feed and

2.2 You want prions in that burger

such cattle were still excluded.

Speaking of “ruminant-to-ruminant” feed bans, the [FDA enacted one](#) in 1997. But in 2005, [a General Accounting Office report](#) found that many feed firms had never been inspected and many were still processing ruminant protein. According to another [FDA report](#) from later in 2005, showed that 4,553 firms—27% of that *had* been inspected—were handling prohibited materials. Such materials are allowed in feed intended for export (apparently it’s perfectly acceptable to infect our foreign customers with BSE via their cattle) though to avoid scaring off export customers, such infected feed need not be *labeled* as containing ruminant protein. So some—no one knows just how much—is almost certainly making its way onto US feed lots.

Our second question is, “What is an acceptable rate of *discovery* of BSE infected cattle in the US?”. Here again we know that the answer is, “Only *zero* discovery is acceptable.” On December 9th 2003, Dave Louthan shot a dairy cow that he feared would trample “downers” (cows too weak or ill to stand) in its trailer at Vern’s Moses Lake Meats in Washington. The cow’s brain was sent for testing; because APHIS was temporarily paying \$10 an animal for test samples, Vern’s had begun sending in about 80 samples a month in October, although in the two years ending in mid-2003 not a *single* cow was tested in Washington state. On December 11th, the animal’s split carcass was sent to Midway Meats for deboning, and on the 12th the meat was shipped on to two licensed meatpackers. By the time the cow’s BSE test came back positive on December 23rd, the meat had long since been distributed to wholesalers, and from them to retailers in 8 states. In a press conference called to try to quench a firestorm of outcry the next day, the Secretary of Agriculture revealed that barely 20,000 cattle had been tested nationally in 2002. A [Christmas Day editorial](#) in the New York Times commented that “This single case will expose the holes in the American system of meat production and disease testing.”

2.2 You want prions in that burger

The discovery of that single case of BSE eventually led 50 countries to ban imports of US beef. It was not until mid-2008 that normal exports were resumed in some of the biggest Asian markets, like Japan and Korea. Japan, which accounted historically for about half of all beef exports from the US with an annual value close to \$2 billion, had allowed in some beef earlier but closed its markets again in 2006 when another case of BSE was discovered in the US. Those \$2 billion a year are the biggest reason why “Only *zero* discovery is acceptable.”

So an ideal U.S.D.A. testing program for BSE would meet 2 goals: zero discovery and zero exposure. That’s easy to do, if there are no BSE infected cows in US herds. You could test every animal slaughtered, as Japan does: you wouldn’t find any BSE and everyone would know they were not being exposed to BSE. In Europe, where there is some BSE present, all animals older than 30 months at slaughter are tested; because of its slow development tests generally do not reveal BSE in younger animals even if they are infected. These countries also enforce farm-to-fork traceability of all food products, so they can easily locate infected meat at later processing stages and have it destroyed. This has benefits in controlling other types of contamination, like *escheria coli*. Currently the U.S.D.A. is planning to track animals *to* the slaughterhouse, and to track meat *from* the slaughterhouse, but has no plans for tracking meat *in* the slaughterhouse, where similar cuts from many different animals are routinely mixed.

But a dilemma arises if some US cattle actually do have BSE. Then the only way to ensure zero exposure is to identify those cattle with BSE so they can be kept out of the human food chain. But that, of course, violates zero discovery. The only way to ensure zero discovery is not to test any cattle for BSE. And that, of course, violates zero exposure. One of the two goals has to give.

APHIS’s solution has been to abandon *both* as absolute goals. Public concern over BSE is very high—over 80% of American consumers

2.2 You want prions in that burger

would like to see every animal tested—so *some* BSE testing must be done. This means there’s *some* risk of BSE discovery. But the testing program is designed to ensure that zero discovery is *very likely* in most years even if the rates of infection or **prevalence** is of the order of 1 animal in 100,000, meaning that about 300 infected animals are consumed every year. That’s an amount of meat weighing as much as about 600,000 quarter-pounders. APHIS claims that its program shows that the prevalence of BSE in American herds are actually less than 1 in 1,000,000(or even in 10,000,000) but this claim is based on very questionable assumptions. Other U.S.D.A. policies that we’ll review late convince me that the goal is *not* to answer the **BSE QUESTION 2.2.1**.

Let’s look at the math, as summarized in **TABLE 2.2.2**, and then I’ll close by comparing what it tells us about APHIS’s testing program to APHIS’s claims about its program. You’ll have to take my word for the numbers in the tables that follow, but you’ll soon be able to derive them yourself using the **BINOMIAL DISTRIBUTION FORMULA 4.7.23**: see . The numbers assume that cattle to be tested are chosen at random. In particular, whether or not one animal being tested has BSE has no influence on the chance that any other does. This is another example of the **independence** property we met in the preceding section in the notion that “The coin has no memory”. We’ll also assume that the APHIS tests successfully identify as infected all cattle with BSE. This assumption is open to debate: see the discussion of the middle portion of of the table below.

The first set of numbers shows what we’d expect when testing 20,000 cattle a year; this was the number of tests conducted in the years before the BSE infected cow was detected in 2003. If 1 in 100,000 cattle were infected, we’d expect to uncover at least 1 case of BSE (i.e. more than 0) in about 18% of years. In other words, we’d see a case about every 5 years, and more than one case only once every 50 years.

2.2 You want prions in that burger

Prevalence	1 in 100,000		1 in 1,000,000	
20,000 cattle tested	Probability that number of infected cattle is			
k	exactly k	more than k	exactly k	more than k
0	0.818730	0.181270	0.980199	0.019801
1	0.163748	0.017522	0.019603	0.000197
2	0.016374	0.001148	0.000196	0.000001
3	0.001092	0.000057	0.000001	0.000000006
360,000 cattle tested	Probability that number of infected cattle is			
k	exactly k	more than k	exactly k	more than k
0	0.027323	0.972677	0.697676	0.302324
1	0.098365	0.874312	0.251164	0.051160
2	0.177058	0.697255	0.045209	0.005950
3	0.212470	0.484784	0.005425	0.000526
40,000 cattle tested	Probability that number of infected cattle is			
k	exactly k	more than k	exactly k	more than k
0	0.670319	0.329681	0.960789	0.039210
1	0.268130	0.061551	0.038431	0.000778
2	0.053625	0.007926	0.000769	0.000010
3	0.007150	0.000776	0.000010	0.0000001

TABLE 2.2.2: NUMBER OF CATTLE WITH BSE DETECTED PER YEAR

At a prevalence of 1 in 1,000,000, we only expect to find any infected animal once every 50 tears of so and we'd essentially *never* see more than 1—well, once every 5,000 years or so. In summary, none of this data provides much basis for claiming that infection rates are less than 1 in 100,000 and we could expect such results even if prevalence was considerably higher.

For just over two years, between June 2004 and July 2006, testing rates were increased to about 1,000 animals a day—say 360,000 a

2.2 You want prions in that burger

year. These tests uncovered 2 infected cows, one in Texas in November, 2004 or June, 2005 and one in Alabama in February, 2006. The ambiguity about the date for the Texas animal arises because the November, 2004 tests did not disclose BSE (but for other reasons the animal was kept out of the human food supply). However, the U.S.D.A.'s Inspector General, against the wishes of agency's head who was "irked" by the move, requested that a more sensitive Western blot test be run. This was done in June, 2005 and the test came back positive for BSE. Later, that month the BSE testing protocol was altered to include the Western blot test. So it is hard to say how many infected animals were missed in the first year of the program. We'll assume as mentioned before the table that this number is 0.

So we observe a single infected animal in each year of the program. If prevalence was as high as 1 in 100,000, these outcomes *would* be surprising. We'd expect to see no more than 1 infected animal only about 1 year in 8. I won't explain the reasoning needed here, or in what follows, but the arithmetic is involved easy and I'll show it: $0.027323 + 0.098365 = 0.125688 \approx \frac{1}{8}$. So we'd expect this two years running only about 1 time in 64 or 1.5% of the time. At least, that's well below the standard minimal cutoff of 5%.

What about a **prevalence** of 1 in 1,000,000? Then the jury is out. At this prevalence, only expect to see any infected animal in 0.302324 of all years. To see at least 1 for 2 years running will happen only about 1 time in 11 ($0.302324 \cdot 0.302324 \approx \frac{1}{11}$). In contrast, we'd expect to see *no* infected animal in either year almost half the time ($0.697676 \cdot 0.697676 \approx \frac{1}{2}$). We can't draw a very strong conclusion, but the observed results *suggest* that the prevalence may well be quite a bit higher than 1 in a 1,000,000.

PROBLEM 2.2.3: After June 2006, APHIS reduced the number of tests to 40,000 a year. In the first 2 years of this testing program, no infected animal has been detected. By imitating the calculations I have given above, try to use the numbers in the third section of the table



2.2 You want prions in that burger

to obtain the following conclusions:

- i) At a prevalence of 1 in 100,000, we expect to see no infected animal in about 67% of years, or in about 2 out of 3 years.
- ii) At a prevalence of 1 in 100,000, we expect to see no infected animal in 2 consecutive years about 45% of the time.
- iii) At a prevalence of 1 in 1,000,000, we expect to see no infected animal in about 96% of years, or in about 19 out of 20 years.
- iv) At a prevalence of 1 in 1,000,000, we expect to see no infected animal in 2 consecutive years about 92% of the time.

What is no infected animal is found in the third year of this program? By cubing numbers in the table, argue that:

- i) At a prevalence of 1 in 100,000, we expect to see no infected animal in 3 consecutive years about 30% of the time.
- ii) At a prevalence of 1 in 1,000,000, we expect to see no infected animal in 3 consecutive years about 87% of the time.

How many years will the APHIS testing program have to continue without finding a single infected animal before the chance of seeing such an outcome drops below the standard 5% level?

- i) Show that, at a prevalence of 1 in 100,000, we must see no infected animal in 8 consecutive years before the chance of this happening is less than 5%. Hint: Compute powers of 0.670319.
- ii) Show that, at a prevalence of 1 in 1,000,000, we must see no infected animal in 75 consecutive years before the chance of this happening is less than 5%.

The numbers above make it pretty clear what the basic problem with APHIS' testing program is: they just aren't testing enough cattle. If we tested 3,000,000 cattle—that just 10% of the animals slaughtered each year—and found no BSE we'd be 95% certain that prevalence was below 1 in 1,000,000. If we found no BSE for 2 years running, we be over 99.7% certain. If we tested 10,000,000 cattle for a single year without finding any BSE, we could essentially rule out the possibility of a prevalence of 1 in 5,000,000: we'd have observed something that

2.2 You want prions in that burger

would happen only 1 year in 300 at that prevalence. So it would be very easy for APHIS to answer the [BSE QUESTION 2.2.1](#).

In fairness, I should explain that APHIS claims that it has answered this question. More precisely, it claims that the testing program described above gives us 99% confidence that BSE prevalence in US cattle is below 1 in 10,000,000 (yes, that's *ten* million). Do they just not understand the math? No. A Background Report prepared by APHIS in 2004 (cited in the [Federal Register](#), but no longer on the APHIS website) makes this clear. It states that, at a prevalence of 1 in 1,000,000, "to achieve a 95 percent confidence level in detecting at least one case from a random sample of adult cattle, we would have to randomly sample and test approximately 3 million animals". That's exactly right.

Instead, APHIS assumes that infected animals will "all be found in the high-risk cattle population" (again, this and further quotes are from the APHIS Background Report). Specifically, they assumed that in the roughly 100,000,000 US cattle, only the roughly 45,000,000 adult cattle might be infected. There's no good evidence for this, but, hey!, we can't test for BSE in younger cattle anyway. Next, they decided to try to rule out a prevalence of 1 in a million. Thus they were looking for 45 infected animals. Now comes the punch line. They assume that: *the majority of these potential infected cattle will be in a sample of 195,000 "high-risk" cattle that they selected for testing.*²

That is, they assume that they are dealing with a population of cattle where the prevalence of BSE is roughly 0.00023. That's 23 times a prevalence of 1 in 100,000 and 230 times a prevalence of 1 in 1,000,000. Not finding an infected animal in 12,500 tests gives roughly 95% confidence that prevalence is not this high; not finding one in 20,000 tests gives 99% confidence.

²They also assume, as we did, that they can test accurately for BSE. We have no way to evaluate this assumption, so we'll pass over it, but the missed animal in 2005 makes you wonder.

2.2 You want prions in that burger

The italicized assumption makes life so easy that one can just as well rule out much lower **prevalences**. So the U.S.D.A. does. In July 2006 (after finding the third infected animal), it claimed that 2 positives in 759000 tests (over all years) gives a best estimate of 4 to 7 infected animals amongst 42,000,000 adult cattle. This is still in the **current document on their site** as of April, 2009, but prevalences less than 1 in 18,000,000 are claimed online.

What about that italicized assumption? Well, there's some evidence from much larger European testing programs that prevalence rates *are* much higher in certain groups of like **downers** (cattle slaughtered when unable to walk). But there's been no well designed experiment to show how much higher such prevalence rates might be. Indeed, the same APHIS survey goes on to note that "It is important to note that *no estimations of prevalence are done when designing these surveillance plans.*" (italics mine).

Remember, I said that my belief was that the APHIS program was designed *not* to answer the **BSE QUESTION 2.2.1**. Well, at least we can be sure that it was *not* designed to answer it. The Background Report states: "*The objective of the surveillance plan is not to estimate prevalence of BSE in the U.S. cattle population*" and goes on to make this very clear: "We would like to clarify that, at this time, there have been no sampling-based, quantitative estimates of the **prevalence** of BSE in the United States. Certain assumptions have been made to assist in developing the surveillance plan, but *this is very different from calculating or estimating the **prevalence** of BSE in the United States.*" Again, my italics.

OK, but not having designed a plan that answers the question is not the same as designing a plan not to answer it. Why do I claim that APHIS' goal was the latter (as well as the former)? Well, imagine you are a beef exporter whose main business is in the lucrative Japanese market (where prices per pound are much higher than in the US or most other countries). Japan has just closed its border to all US beef

2.2 You want prions in that burger

and said that it won't open them unless *every* animal whose meat is shipped to Japan is tested for BSE. That kind of testing is clearly not going to be done for you by the U.S.D.A. What should you do?

The answer is pretty obvious. Test every animal *your company* slaughters for BSE. Then you can give the Japanese the assurances they're demanding and start exporting *your* beef to Japan again. Remember we're talking a billion dollar market here. There's just one catch. Federal Law gives APHIS legal responsibility for animal health so you can't just go off and start your own testing program. Your program must be authorized by APHIS.

No sweat. You're going to pay for testing your animals yourself—and you can well afford to since the beef you sell in Japan goes for much more per pound than beef sold domestically. You're happy to have APHIS monitor the actual testing and to compensate it for the time of its staff. So you apply to APHIS to permission to start testing your animals for BSE on these terms. That's just what [Creekstone Farms](#), a Kentucky and Kansas based processor of black Angus beef that sold a third of its beef in Japan and Korea, did in 2006. Specifically, Creekstone asked for permission to test for BSE in a testing site built to U.S.D.A. specifications, the roughly 300,000 cattle it processes annually.

The U.S.D.A. *denied* this request, citing the 1913 Virus-Serum-Toxin Act, intended to ensure the safety of animal vaccines, as justification. This act gives the U.S.D.A. authority to regulate “any worthless, contaminated, dangerous, or harmful virus, serum, toxin, or analogous product for use in the treatment of domestic animals” and the agency held that BSE testing kits were an “analogous product” and that they were “worthless, contaminated, dangerous, or harmful”.

Wait a minute, back up the bus here: we're talking about *exactly the same* BSE testing kits that the U.S.D.A. uses in its *own* BSE testing program! It gets better. Creekstone took the U.S.D.A. to court. The case made its way through the courts: Creekstone was upheld in US

2.2 You want prions in that burger

District Court in March, 2007 but that decision was overturned in the US Court of Appeals in September, 2008 and the case sent back to District Court.

In a dissent to the Appeal Court decision, Chief Judge Sentelle noted, “It seems that the Department’s fear is that Creekstone’s use of the test kits would enable it to provide buyers with a *false assurance that the cattle from which its beef is obtained are free of Bovine Spongiform Encephalopathy*.” But APHIS claims that its own testing program shows that *all* American cattle are free from BSE. Creekstone could only provide *false* assurance of safety if its cattle tested negative but were actually BSE infected. In fact, Creekstone’s cattle are too young for the tests to reveal BSE and Creekstone agrees that the tests do not demonstrate that its cattle are BSE free. They only want to test because that’s what Japan and Korea demand and they want to sell beef in those markets.

My guess (and purely a guess) is that the U.S.D.A. fears that if come producers began universal testing, the US public would begin to demand such testing on the beef *it* eats. Of course, that’s only a worry if you’re not really *sure* that such testing wouldn’t uncover an unacceptably high prevalence of BSE. I’m not.

So that’s my case for claiming that APHIS has designed a testing program not to answer the [BSE QUESTION 2.2.1](#). Why else should it act so aggressively to prevent any private testing that might answer that question?

So, is it safe to eat beef? Let’s first discuss long term risks. We don’t really know what long term mortality from CJD may turn out to be. Many Western countries are currently experiencing an epidemic of dementia but the causes are unclear. At least a few of these cases are almost certainly due to CJD caused by eating BSE tainted beef. How many? I can find no convincing estimate. The incubation period for CJD is very long so the majority of infections may still be undeclared for a long time. Moreover, reliable diagnosis of CJD is only possible

2.2 You want prions in that burger

by autopsy. I think that it will be a long time, if ever, before we know real rates of BSE infection over the past 20 years. It may be very low—so far only about 200 deaths are attributed to it—but it could also be very much higher.

Next, what about your risk from BSE today? Well, only certain tissues in cattle (principally the brain, spinal cord and perhaps marrow) seem to carry prions. Unfortunately, a lot of beef is probably contaminated with such tissues. The first step in preparing a carcass for deboning is splitting: this involves passing a large saw through the carcass longitudinally, bisecting the brain and spinal cord and typically spreading tissues from these areas widely over the carcass.

Meat is often removed by Advanced Meat Recovery or AMR systems which essentially squeegee it off the bone under high pressure and produce a sort of "blood, marrow, and muscle gumbo". A recent study by the Food Safety and Inspection Service, another arm of the U.S.D.A., found that 35% of AMR meat sampled contained marrow and spinal cord. But the real piece of data we'd need to assess our BSE risk is an answer to the [BSE QUESTION 2.2.1](#) and we're not going to have that anytime soon.

Ok, so is it or isn't it safe to eat beef? My personal feeling is that the primary risk in eating beef is not from BSE but from saturated fat. A [major recent study](#) suggests that reducing red meat consumption could lower the mortality of the 20% of the population who are the biggest consumers by 11% in men and 16% in women. Nationally, that translates to tens of thousands of deaths per year attributable to the dietary role of beef and other red meats. That's almost certainly much higher than the number of CJD deaths we'll see even if infection rates are much higher than we currently think. Despite this knowledge, I eat beef. So what the heck? Give me a side of prions with that burger.

Chapter 3

When it really counts

This chapter introduces the objects we use to describe collections in mathematics—sets to give them their mathematical name. Set are very much the air we breathe in mathematics, fundamental, simple objects that underlie all the more ornate structures in the subject.

Our study of sets has two main themes. The first is to review and relate ways of creating new sets from old, using a variety of set operations: products, subsets, power sets, intersections, unions and complements. The goal here is to learn how to use these operations to simplify large, complicated sets by disassembling them into small, basic components.

The second theme is learning ways to count sets: that is, calculate their orders, the number of objects or elements they contain. This is the problem to which we'll apply our study of set operations, using them to find the orders of sets that are much too big to list or count directly. Our approach will be to develop shorthands for the orders of common basic component sets, and then to understand how to reassemble such shorthand counts into an overall count for a large set built from such components.



3.1 Sets: the air we breathe

The language of sets and our ability to count them will then be key tools in our study of probability and statistics in the next chapter.

3.1 Sets: the air we breathe

The purpose of this section is mainly linguistic. We're going to learn some terminology for dealing with sets. Things will be pretty informal for two good reasons.

First, although a lot of the mathematics we'll be doing later on uses sets, they're very much in the background. They're a bit like air. We're surrounded by it, we breathe it all the time, but we never need to think about it. Like air, sets will surround us and we'll use them all the time, but, if we have a feel for what sets are and *solid command* of a few basic set relationships, we seldom need to think consciously about them.

But be warned! This is one of those times when you either really master the necessary ideas and then everything else is easy as breathing, or you're just a bit vague on the basics and then you'll have asthma for the next two chapters. Do yourself a favor and force yourself to really memorize the (very few) definitions and formulae in this section.

The second reason we're going to be informal about sets is that defining them carefully is quite tricky. It took decades for mathematicians to see what some of the subtle points involved were and decades more for them to find the right ways of dealing with them. Fortunately, for you, only professional mathematicians really need to come to grips with these niceties.

Objects and oracles

OK, let's start by recalling the definition of an **atom**. To the Greeks, an atom was something that could not be divided into parts—the Greek word *atomos* means uncuttable. Democritus' philosophical school held that all matter was composed of eternal atoms though he had no idea what these atoms might be. The modern use of the term in chemistry is due to Dalton around 1800 who claimed that atoms were of many types, each type giving rise to one of the chemical elements discovered by Lavoisier. The viewpoint of the time was that these smallest elemental components—our modern atoms—in *chemical* reactions must be the smallest components of all matter. By the late 19th century, physics had begun to reveal that this use of the term atom was a misnomer because these atoms were composed of even smaller “sub-atomic” particles like electrons, protons and neutrons. Today we know that these particles are themselves made up of even smaller constituents like quarks. In hindsight, being an atom—indivisible—is relative, not absolute.

I recall all this to explain what we will mean by an *object*. Well, an object is just any thing we want to distinguish. What's a thing? An object, of course! I hope you start to see the difficulty of giving careful definitions. On the other hand, as I've already said we don't really need to be more precise. For us, an object will be like one of Democritus' atoms. It's something we choose to view as indivisible but we don't worry about specifying its properties much more carefully. Some objects are mathematical. Each of the natural numbers—the numbers 0, 1, 2, 3, ...—is an object. So is each of the letters *a*, *b*, *c*, ..., *z* in the alphabet. But many objects are completely non-mathematical. You are an object. More generally, each student in your section of **MATH4LIFE** is an object. Your left sock is an object. So is your right sock. The primary colors *red*, *green* and *blue* are each objects. So far (I hope) so good.

Now comes the only mildly tricky point. Being an object is, like be-



3.1 Sets: the air we breathe

ing an atom, relative. The difference is that, unlike the chemists and physicists who found out that their atoms could be decomposed, we'll more often want to consider of group of objects as forming a single object. When we do this we have to *choose* to view the group as atomic or indivisible. The advantage of doing this is that it makes the concept of being a object very general. For example, the *pair* of socks you have on is an object. I know that it *could* be divided into a left sock and a right sock but I choose to view it as one object (the pair) rather than two (the socks). Likewise, you can choose to view the 25 students in your section of **MATH4LIFE** is an single object, by simply deciding that the section itself is the object and ignoring the possibility of dividing it into 25 individual students. Even the natural numbers can be viewed as a single object by simply choosing to view the whole collection as an atom and forgetting the possibility of decomposing it into (infinitely many!) individual numbers. The natural number object is important enough in mathematics to have its own notation \mathbb{N} . The basic idea is the same in all three cases. We get to decide what's an object to suit our own convenience, just as, even today, it's useful to view atoms as "atomic" when studying chemical reactions but as collections of particles when studying nuclear physics.

Here's the good news. You can pretty much forget everything I have just said. Just remember that an object can be anything that we find it useful to label as an object as long as we agree to view the object as atomic (even if we know better).

The advantage of this long discussion of objects is that it makes saying what a **set** is pretty easy. Informally, a set is just an object that we have *decided* to view as a **container** that holds other objects. We'll denote sets by upper case letters like A, B, C or U, V, W to make it more visible that we view them as containers holding other objects rather than as atomic objects.

EXAMPLE 3.1.1: The natural numbers \mathbb{N} is a set that contains the



3.1 Sets: the air we breathe

numbers 0, 1, 2 ...and so on. The Latin alphabet is a set L that contains the 26 letters from ‘a’ to ‘z’, your pair of socks is a set P containing just two objects, your left sock and your right sock, the primary colors form a set C with contents are the colors red, green and blue, and the students in your section of **MATH⁴LIFE** is a set S whose contents you can describe better than I.

So, if sets are just special container-objects, how do we tell the players without a program? In other words, what makes an object a set other than saying so? Formally, a set is an object that comes equipped with an **oracle**. I’ll call the oracle attached to the set A the A -oracle.

What do I mean by an oracle? Historically, oracles have existed in many cultures and in many forms. What they have in common is that they provide authoritative answers to questions. You’ve probably heard of the oracle of Delphi who gave answers to the ancient Greeks through a priestess to the god Apollo, many of which are the subject of many classic poems and dramas. But the Chinese used turtle bones as oracles far earlier, they appear as voices from the sky in the great Indian epics the Mahabharata and Ramayana, and are found in African, Central American and many other cultures. Set oracles are so named because they also provide answers to questions.

What kind of questions must the A -oracle be able to answer to qualify an object A as a set? Only very special ones, that let us decide what objects make up the set A . *For any object x* , the A -oracle must be able to answer yes or no to the question “Is the object x in the container A ?”. If the answer is “Yes”, we say that x is an **element** of the set A and write

$$x \in A.$$

This last is read in several ways. In addition to saying “ x is an element of A ”, we say, more simply, “ x belongs to the set A ”, or simplest of all “ x is in A ”. If the answer is “No”, we say that “ x is not an element



3.1 Sets: the air we breathe

of A ” or “ x is not in A ” and write

$$x \notin A.$$

Another way to think of the A -oracle is as an admissions test for A . The elements x of A are the ones that pass the test and are admitted to A (by getting a “Yes” from the A -oracle). All other objects get a “No” and fail.

If you’re wondering where I plan to erect the A -temple of Delphi or look for A -turtle shells, I’m afraid that set oracles are much more down-to-earth. What they do have in common with classical oracles is that come in a variety of forms. We’ll often be able to give multiple oracles for a single set A . This does *not* mean that we’re giving multiple sets. All that counts about an oracle is the answers it gives—that is, which objects it says belong to A . The embodiment of the oracle—be it a list, a formula, a description, a priestess or a shell—doesn’t matter. A set A is simply its members or contents. Put differently, two sets A and B are the same if exactly the same objects x belong to both, even if the A -oracle and the B -oracle at first seem very different.

The simplest way to give an A -oracle is simply to write down a list of all the elements of A . To indicate that this list is a set we surround it in what are called **set braces**: $\{ \}$. You’ve probably seen this way of writing the set \mathbb{N} of natural numbers:

$$\mathbb{N} = \{0, 1, 2, 3, \dots\}.$$

But we can also write the letters in the alphabet as

$$L = \{a, b, c, d, e, f, g, h, i, j, k, \ell, m, n, o, p, q, r, s, t, u, v, w, x, y, z\}$$

or the primary colors C and your socks P as

$$C = \{\text{red, green, blue}\} \quad \text{and} \quad P = \{\text{left sock, right sock}\}.$$

One point about such lists that is a bit misleading is that they have a left-right order. Changing this order *does not* change the set listed.

3.1 Sets: the air we breathe

Why? Because, as I've stressed above, changing the order does not change which objects x belong—that is, appear in the list, it just changes *where* they appear in the list. We don't care where, just whether. So we could just as well have written $C = \{\text{green, red, blue}\}$ and $P = \{\text{right sock, left sock}\}$.

SETS VERSUS LISTS 3.1.2: A brief digression. In [SECTION 3.5](#), we'll want to consider object rather like sets except that the elements *do* come with an order. We'll call these **lists** and use brackets, rather than set braces to emphasize the difference.

So $[1, 2, 3]$ is a list and $\{1, 2, 3\}$ is a set. The difference is that $[1, 2, 3]$ and $[3, 2, 1]$ are *different* lists (because the change in order matters for lists) but $\{1, 2, 3\}$ and $\{3, 2, 1\}$ are the *same* set (because the change in order does not matter for sets). But I'll wait until we come to them in earnest to say more about lists.

Back to sets and a very important point. Listing elements is usually *not* how we want to give an A oracle. The problem is that it's simply too cumbersome when the set has lots of elements. It was annoying to have to write down the letters of the alphabet to describe L . You certainly wouldn't want to have to do too many homework problems with your section S of **MATH⁴LIFE**, if you have to write out a class list every time you mentioned S . Soon we'll be working with sets with so many elements that it would be impossible to list them all in a life-time, much less at the end of your English class when you're rushing to finish your math homework. Even worse, it was *impossible* to list all the elements of \mathbb{N} because there are infinitely many, and I sleazed out above with those ellipses (...) indicating that I was omitting all the rest.

We need set oracles that are more concise than those that simply list their elements. What we're almost always going to use is what experts call an informal “natural language” description. In plain english, plain english! In fact, we already did this in [EXAMPLE 3.1.1](#) before I had even said carefully what a set was and I'm sure you had



3.1 Sets: the air we breathe

no problem understanding what each of the example sets there was. When I said the letters in the alphabet formed a set L , we both knew that the elements of L were the 26 letters I listed above. Likewise there was no problem in speaking about the set S of students in your section of **MATH4LIFE**. We both know that you're in S and Britney Spears isn't. Even the "list" of natural numbers I gave above is really an informal description. I only listed the first 4 elements of \mathbb{N} but we both interpreted those ellipses as indicating that \mathbb{N} contains the objects 5, 6, 7, ... (and we both know what objects comes after these, and so on).

In all these examples, there are some implicit assumptions. If you were Russian, you might have said "Don't you mean the 33 letters in the alphabet?" and we'd have had to settle which alphabet we meant. Likewise, I'm comfortable thinking about the set S of students in your section, but if you asked me how many elements S has—that is, how many students are in your section—I'd have to don my priestess' robes and email your Registrar to find out. And we both share the intuitive notion of the natural numbers \mathbb{N} as all those numbers you'd get if you could "count forever".

I'm not going to worry about such potential confusions and you don't need to either because, in practice, they just don't come up for the finite sets we'll be working with. The informal description of any set A that we need will be sufficiently clear that we'll both understand what set is meant and be able to serve as our own oracles about what objects are elements of A . And, if we need more information, like how *many* elements we'll be able to derive an unambiguous answer from the description.

You may be wondering why it took all those mathematicians so long to settle on the right definition of sets. The answer is connected with problems that arise with certain very big infinite sets and with sets whose elements are other sets. I won't go into any further details but if you are interested you can google "Russell's paradox"—Bertrand



3.1 Sets: the air we breathe

Russell was one of the mathematicians who helped clarify and resolve these issues.

Let's sum up. An object is something we choose to view as atomic. A set is a special kind of object that we view as both an atomic entity and as a collection of objects called its elements. What identifies an object as a set is saying so. We need to specify how to tell what the elements of the set are—give its oracle—but we're free to provide this oracular specification in whatever form is most convenient for us.

The order of a set

Before we end this subsection, there's one more basic term to introduce.

ORDER OF A SET 3.1.3: *If a set A is finite, we define the order $\#A$ of A to be the number of distinct elements in A . If A is infinite, we simply say that A has infinite order and write $\#A = \infty$.*

EXAMPLE 3.1.4: Just so there's no confusion we give the order of some of the sets above. The order of the Latin alphabet L is 26. The order of the set P of socks is 2. The set \mathbb{N} of natural numbers has infinite order.

EXAMPLE 3.1.5: How many elements does the set $A = \{a, b, b, c\}$ have? The answer is 3 not 4! The point is that, although we have listed b twice, it only gives us 1 element of A .

This is the meaning of the emphasized adjective distinct in the definition of the order of a set. An element is either in a set A or it's not. You can't hold multiple memberships in the A club. It's not totally wrong to list an element more than once, as b was above, but it's pointless and confusing to do so. So we don't.

Warning: These examples are misleading. It's stupidly easy to find the order of these sets because we just count. The most difficult

3.2 Subsets

skill in this course for most students is learning to “count” without counting. That is, to find the number of elements in a set A which is informally described—rather than listed—without trying to list the elements of A . This skill will be our basic tool for working with probabilities and statistics. And, there’s no way of avoiding it because, I have already indicated, we’ll very soon be working with sets with far too many elements to list.

3.2 Subsets

Here is where we start to lay out the basic set relationships and operations I mentioned at the beginning of this section. Unlike the preceding subsection, where all you needed was to get a feel for what sets were, in this subsection and the next, it’s critical to really master the ideas. From here, on you need to pay careful attention.

In this subsection, we discuss the fundamental relationship of inclusion or containment between sets, then see how it can be used to construct new sets from old ones.

SUBSET 3.2.1: *We say that B is a subset of A and write $B \subset A$ if, informally, every element of B is also an element of A .*

There are lots of other ways to say or write this. Other ways to say that B is a subset of A are to say that B is contained in A or that A contains B , and, instead of $B \subset A$ we often write $A \supset B$.

There are also lots of ways to restate the condition that every element of B is also an element of A , all with the identical meaning. Using element notation, B is a subset of A , if $x \in A$ whenever $x \in B$. Or, B is a subset of A , if $x \in B$ implies that $x \in A$. In terms of oracles or admissions tests, this simply means any object x that passes the admissions test for B automatically passes the test for A . The “club” or collection B is more exclusive than A is.



3.2 Subsets

EXAMPLE 3.2.2: Consider the sets $D = \{0, 1, 2, 3, 4, 5, 6, 7, 8, 9\}$, $E = \{0, 2, 4, 6, 8\}$, $O = \{1, 3, 5, 7, 9\}$, $L = \{5, 6, 7, 8, 9\}$, and $S = \{0, 1, 2, 3, 4\}$.

Is $E \subset D$? Yes because, each of the 5 elements 0, 2, 4, 6, and 8 of E is also an element of D . In fact, O , L and S are also subsets of D . Is D a subset of D ? Yes. Why? See [EXAMPLE 3.2.6](#) for a general solution, but try to work this out for yourself first.

Is E a subset of S ? No, because $6 \in E$ but $6 \notin S$. Equally, no because $8 \in E$ but $8 \notin S$. It doesn't matter that *most* of the elements of E are in S : E is only a subset of S if *every* element of E is in S . Show that E is not a subset of O or L either.

EXAMPLE 3.2.3: In the previous example, we listed the elements of each set. Remember we don't want to make this a habit because it just doesn't work when the sets get bigger. So let's try a descriptive example. Consider the set T of two digit numbers (from 10 and 99), the set H of whole numbers from 1 to 100, and the set E of even whole numbers.

Checking that a set B is not a subset of a set A is easy: just exhibit any element of B that is not an element of A . Here are some examples: Is $E \subset T$? No, because $100 \in E$ but $100 \notin T$. Equally, no, because $6 \in E$ but $6 \notin T$.

Is $E \subset H$? This time $100 \in H$ and $6 \in H$. Nonetheless, $102 \in E$ but $102 \notin H$ and since there's an element of E not in H , E is *not* a subset of H .

Is $T \subset E$? No, because $99 \in T$ but $99 \notin E$. Equally, no because $47 \in T$ but $47 \notin E$.

Is $H \subset T$? No, because $100 \in H$ but $100 \notin T$. Equally, no because $8 \in H$ but $8 \notin T$.

Checking that a set B is a subset of a set A is a bit harder: we have to check that every element of B is an element of A . Here's an example: Is $T \subset H$? Yes. If x is any two-digit number, that is, any element of T

3.2 Subsets

then $10 \leq x \leq 99$. If so, then it's also true that $1 \leq x \leq 100$ and these inequalities are just a different way of stating the test for admission to H . Thus we've checked that every element of T is in H so $T \subset H$.

In the last example, notice how useful it was to be able to replace the natural language “numbers from 1 to 100” oracle for the set H with the more formal “ x satisfying $1 \leq x \leq 100$ ” oracle. This freedom is why it's important that we identify a set by its members (what objects it contains), not by its oracle (how these objects are described).

I want to define two more concepts so that I can point them out when they arise in the problems that follow.

EMPTY SET 3.2.4: *The empty set \emptyset is the set with no elements. That is, no object x is a member of \emptyset . This is one set it's easy, if pointless, to list: $\emptyset = \{ \}$.*

DISJOINT SETS 3.2.5: *We say sets A and B are disjoint if no x is a member of both A and B . Informally, B disjoint from A is as far as B can get from being a subset of A .*

EXAMPLE 3.2.6:

- i) For which A is the empty set \emptyset a subset of A ?
The empty set is a subset of every set A . If $\emptyset \not\subset A$, there would have to be an element $x \in \emptyset$ such that $x \notin A$. No such x can exist for the stupid, but very adequate, reason that no x is a member of \emptyset .
- ii) For which sets A is A itself a subset of A ?
Every set A is a subset of itself. What does $A \subset A$ mean? It means that if $x \in A$ (the subset), then $x \in A$ (the containing set). That's tautologically true—a tautology is an implication in which the hypothesis (the if-part) is the same as the conclusion (the then-part)
- iii) Give an example of a set A whose only subsets are the empty set \emptyset and itself.
The empty set itself is such a A . So also is any set A with just a single element x . A subset B of such an A can contain no element except x .

3.2 Subsets

If B contains x , then $B = A$; if not, then $B = \emptyset$. Are there any other examples?

PROBLEM 3.2.7: Consider the set S of states in the U.S., the set C of states in the continental U.S., the set T of the 13 original colonies, and the set B of blue (i.e. Democratic) states in the 2008 Presidential election. Which of these sets are subsets of which others?

PROBLEM 3.2.8: Consider the set D^2 of pairs of numbers from 1 to 6 like $(2, 3)$ and $(6, 4)$. We consider the order of the numbers to matter, so $(2, 3)$ and $(3, 2)$ are different elements of D^2 . This means that D^2 has $6 \times 6 = 36$ elements running from $(1, 1)$ to $(6, 6)$, although this count will not be needed in this problem. We will work a lot with this set later on, viewing the numbers from 1 to 6 as the 6 faces of a standard die, and the pairs of numbers in D^2 as the numbers on a pair of dice. Since the order of the pair of numbers matters, we need to distinguish the two dice which I'll usually do by imagining that they have different colors, say **blue** and **red**. I hinted at this by the coloring a few of the pairs above, but usually I'll rely on you to make the distinction in your own mind.

Here are several sets that are subsets of D^2 *by construction* or *by definition*—that is, because, we *only* admit elements to them that already lie in D^2 :

ES is the set of pairs with an “Even Sum”.

$S7$ is the set of pairs with “Sum 7”.

$S4$ is the set of pairs with “Sum 4”.

BE is the set of pairs which are “Both Even”.

BO is the set of pairs which are “Both Odd”.

OE is the set of pairs with “one Odd, one Even”.

FO is the set of pairs with “First number Odd”.

Which of these sets are subsets of which others? Try *not* to list the pairs in each of these sets—remember we'll soon be working with sets too big to list. Instead, try to give informal arguments about why elements in one set are or are not in the other. I'll get you started.



3.3 Product sets

Partial Solution: Let's see which are subsets of ES .

$S7$ is not a subset of ES because if a pair has sum 7 it's sum is not even. In this case, not as single element of OE is in ES . In fancier terms, the sets OE and ES are disjoint. If we like, we can give a specific example: $(4, 3)$ has sum 7 so $(4, 3) \in S7$ but $(4, 3) \notin ES$, hence $S \not\subset ES$

$S4$ is a subset of ES because if a pair has sum 4 (i.e., $(a, b) \in S4$), then their sum is even ($(a, b) \in ES$).

BE is a subset of ES because if both entries in a pair (a, b) are even (i.e., $(a, b) \in BE$), then their sum is $a + b$ certainly even (i.e. $(a, b) \in ES$).

$B0$ is a subset of ES . This is a bit trickier. If both entries in a pair (a, b) are odd (i.e., $(a, b) \in BE$), then their sum is even ($(a, b) \in ES$) because "odd plus odd is even".

OE is not a subset of ES . If one entries in a pair (a, b) is odd and the other is even (i.e., $(a, b) \in BE$), then their sum is odd ($(a, b) \notin ES$). As an more specific example, $(4, 3) \in OE$ but $(4, 3) \notin ES$. Once again OE and ES are disjoint.

FO is not a subset of ES . The pair $(3, 4)$ is in FO but not in ES . This time FO and ES are not disjoint: for example, $(3, 3)$ is an element of both FO and ES .

3.3 Product sets

Sequences

One aspect of working with sets that is often a bit confusing for the beginner is that many fields have their own special terms that mean

3.3 Product sets

nothing more or less than “set”, “element” or “subset”. The most important example is the field of probability which we’ll come to in the next chapter, but there are lots of others. Sequences provide another example and I’m going to introduce them now, both as a warmup for probability and because they provide an unbeatable source of exercises in the coming subsections.

A **sequence** s in the Latin alphabet is just a sequence of ordinary a–z letters. Examples of such s are the sequences (t,o,p) , (p,o,t) , (p,o,p,p,y) or (x,z,p,p,q) . The order of the letters matters (so (t,o,p) and (p,o,t) are different sequences) and a letter may be repeated (as in (p,o,p,p,y) or (x,z,p,p,q)). Of course, there’s more familiar term for a sequence of letters in the Latin alphabet: a word. That’s how we’d prefer to view sequences most of the time.

Here’s how we’ll do so. We’ll use those parentheses $()$ when the context does not make it clear that we’re dealing with sequences, but we’ll also drop them when they’re not needed. Likewise, we’ll often omit the commas that separate the successive letters in a sequence when no confusion is likely, and rely on the typewriter font to remind us that we’re viewing the word that’s left as a sequence of letters.

Eventually, in [SECTION 3.5](#), we’ll want to view **lists** with **sequences** and we’ll even drop the typewriter font to make it easier to pass between the two. And one of the key skills we’ll learn in [THE FOUR QUADRANTS AND THE THREE HARD WORDS](#) is how to tell, from an informal description, whether objects are subsets, lists or sequences. At that point, we’ll drop the typewriter font since its presence or absence would giveaway the answer. But in this section, we’ll usually write the example sequences above as `top`, `pot`, `poppy` and `xzppq`. Warning: as the last example shows, the letters in a sequence need not form a dictionary word (even if you might love to be able to put down `xzppq` the next time you play Scrabble™).

The **length** of a sequence is just the number of letters in it—my four

3.3 Product sets

example sequences have lengths 3, 3, 5 and 5. We usually denote this length by ℓ and also often refer to a ℓ -letter sequence instead of a sequence of length ℓ .

You may sometimes see sequences referred to as lists but, in this course, we will use the term lists to refer to a special kind of sequence—one in which no letter is repeated: see [LISTS FROM A AS SEQUENCES 3.5.11](#). Such lists are the subject of the whole of section [SECTION 3.5](#).

It will often be convenient to work with sequences that use other alphabets. Don't worry, you won't need to know Greek or Cyrillic.

ALPHABET AND LETTER 3.3.1: An alphabet is simply a set; *any* set. The elements of the alphabet are called letters. In other words, alphabet is a synonym for set and letter is a synonym for element.

Why introduce such synonyms when we just learned the perfectly good primary terms? There are several reasons. The best is that they indicate how we're going to use a set. Set is a very general term and saying that A is a set gives me no indication of why I am interested in it. If we call A an alphabet, we're declaring that we're going to be working with sequences.

SEQUENCE 3.3.2: A sequence in the alphabet A is a just a sequence of A -letters—that is, elements of A . As for ordinary sequences the order of the letters matters and repetition of letters in a sequence is permitted. The length ℓ of a sequence s is the number of letters in the sequence, exactly as for ordinary Latin alphabet sequences (or words).

You *will* have to get used to using the term letter in a broader than everyday sense. The next example introduces the most common case, where we view digits as letters,

EXAMPLE 3.3.3: One alphabet we'll work with quiet a bit is the set $10 = \{0, 1, 2, 3, 4, 5, 6, 7, 8, 9\}$. In this alphabet, as I warned you above, the letters are simply the 10 decimal digits—hence the name. We are

3.3 Product sets

going to work with the set $\mathbf{10}^2$ of sequences of length 2 in this alphabet: the exponential notation I am using here and in the following examples of sets of sequences will be explained in [PRODUCTS OF A SET WITH ITSELF 3.3.20](#). Just bear with it for now.

The set $\mathbf{10}^2$ is almost, but not quite, the set T of two-digit numbers that appears in [EXAMPLE 3.2.3](#). We can think of each number from 10 to 99 as a 2-digit sequence: that is, every element of T can be viewed as an element of $\mathbf{10}^2$. But $\mathbf{10}^2$ also contains the ten sequences 00, 01, 02, 03, 04, 05, 06, 07, 08, 09 starting with 0 and these do correspond one-digit numbers.

Is $\mathbf{10}^2$ the same as the set H of numbers from 1 to 100? No, for several reasons. Again, the leading 0s cause a problem: the sequence 03 is like the number 3 but it is not the same object. We could agree to ignore leading 0s to get around this, and in some later problems we will. More seriously, however, the element 100 of H is not in $\mathbf{10}^2$ and the element 00 of $\mathbf{10}^2$ is not in H (even if we ignore leading 0s). In other words, H is not a subset of $\mathbf{10}^2$ and $\mathbf{10}^2$ is not a subset of H .

Show that the set $\mathbf{10}^1$ of sequences of length 1 in the alphabet $\mathbf{10}$ is just $\mathbf{10}$ itself. In fact, you should be able to convince yourself that this is true for any alphabet A . Then try to find a formula for the number of elements of the set $\mathbf{10}^\ell$ of sequences of length ℓ in the alphabet $\mathbf{10}$. Hint: $\mathbf{10}^1$ had 10 elements, $\mathbf{10}^2$ has 100; first predict how many elements $\mathbf{10}^3$ has then take a leap to any length ℓ .

PROBLEM 3.3.4: Consider the set L^6 of 6-letter sequences in the Latin alphabet L . The set L^6 already demonstrates that we can't hope to work with sets and subsets by listing elements. It has 308915776 elements, so if you could write out one sequence a second, it would still take you almost 10 years to list its elements! Some subsets of L^6 are:

- NR , the set of sequences with “No Repeated letters”.
- NV , the set of sequences with “No Vowels” (we'll call ‘y’ a vowel for this problem).



3.3 Product sets

NC , the set of sequences with “No Consonants”.

BVC , the set of sequences that contain “Both Vowels and Consonants”.

OE , the set of sequences that contain “at least 1 letter e”.

TE , the set of sequences that contain “at least 2 e’s”.

NE , the set of sequences that contain “No letter e”.

EF , the set of sequences that contain “at least 1 letter e and at least 1 letter f”.

DW , the set of sequences that are “Dictionary Words” (in say, Merriam-Webster’s Online Dictionary).

i) Which of these sets are subsets of which others?

ii) Which pairs of subsets in the list above are disjoint?

Partial Solution: Let’s see which subsets are disjoint from NV .

NR and NV are not disjoint because, for example, “bcdfgh” is in both NV (it has no vowel) and NR (it has no repeated letter). There are lots of other correct solutions because there are lots of other sequences in both sets and it would have been just as good to exhibit any other one.

NC and NV are disjoint because if a sequence has no vowels and no consonants it has no letters so must have length 0 (the sequence with 0 letters is called the null or empty sequence in analogy with the empty set with 0 elements).

BVC and NV are disjoint because if a sequence has no vowels it can’t contain both vowels and consonants.

OE and NV are disjoint because if a sequence contains an ‘e’ it contains a vowel. Likewise TE and NV are disjoint and EE and NV are disjoint.

NE and NV are not disjoint: again “bcdfgh” is in both.

Believe it or not, DW and NV are *not* disjoint. Check out [crwth](#)!

iii) Try to guess a formula for the number of elements of the set L^ℓ of sequences of length ℓ in the Latin alphabet. Hint: L^1 has 26 elements. You can check your prediction by seeing whether it gives 308915776 for the number of elements of L^6 .



3.3 Product sets

EXAMPLE 3.3.5: A **binary sequence** is a sequence in the binary alphabet $\mathbf{2}$ with just two letters 0 and 1. Such sequences may seem pretty silly but they are actually among the most widely studied and used. The main reason is that computers ultimately represent all the data they handle—text, photos, music, whatever—as binary sequences. Lo-o-o-ng binary sequences are we’ll see in this example.

My goal here is to look ahead to the counting we’ll be doing soon and find a formula for the number of binary sequences of length ℓ . It’s not too hard to discover such a formula. If we just list the binary sequences of small length ℓ , we’ll quickly see a pattern. To make sure we don’t miss any, we’ll list sequences in increasing order. Here are the sequences of lengths 1 to 4:

ℓ	list of binary sequences of length ℓ	number of binary sequences of length ℓ
1	0 1	$2 = 2^1$
2	00 01 10 11	$4 = 2^2$
3	000 001 010 011 100 101 110 111	$8 = 2^3$
4	0000 0001 0010 0011 0100 0101 0110 0111 1000 1001 1010 1011 1100 1101 1110 1111	$16 = 2^4$

Looking at this table, it’s not hard to predict that the number of binary sequences of length ℓ —the order of the set $\mathbf{2}^\ell$ —is 2^ℓ . If we look a bit more carefully, we can even see why this is. Why did I divide each list of sequences into two rows? To make it easy to compare the rows. Notice that for each length ℓ , the entries in the top row and the bottom are identical except for the first **bit**—a bit is just a shorthand way of saying a **B**inary **d**igIT, that is, a single 0 or 1. The first bit in each top row is a 0 and in each bottom row is a 1.

This pattern continues to hold for any ℓ . Every sequence of length

$\ell + 1$ consists of a leading bit that is either a 0 or a 1 followed by a sequence of length ℓ . Conversely, from any sequence of length ℓ we get exactly 2 of length $\ell + 1$ by adding either a 0 or a 1 at the start. In other words, the number of binary sequences of length $\ell + 1$ is always twice the number of binary sequences of length ℓ . If the number of length ℓ is 2^ℓ as we predict, then the number of length $\ell + 1$ will be $2 \cdot 2^\ell = 2^{\ell+1}$, again, just as we'd predict.

There's nothing special about binary sequences. If our alphabet has m letters instead of 2, then there will be $m = m^1$ sequences of length 1. Each of these will yield m sequences of length 2 by adding one of the m letters in the alphabet to the start of the sequence so there will be $m \cdot m = m^2$ sequences of length 2. And there'll be m times as many— $m \cdot m^2 = m^3$ —of length 3 and so on. We sum this up for future reference:

SEQUENCE COUNTING FORMULA 3.3.6: *For any $\ell \geq 0$, the number of sequences of length ℓ in an alphabet A with m letters is m^ℓ .*

Notice that we allow $\ell = 0$ in this formula. What's a sequence of length 0? Just what it says, an empty sequence with 0 letters. You can see other examples of the [SEQUENCE COUNTING FORMULA 3.3.6](#) in [EXAMPLE 3.3.3](#) for decimal sequences ($m = 10$) and in [iii\)](#) of [PROBLEM 3.3.4](#) for Latin alphabet sequences ($m = 26$).

Before we close this subsection, I'd like to emphasize how fast the number, 2^ℓ , of binary sequences of length ℓ gets big. Let's compute some examples to get a feel—try calculating these in your calculator to check me. First $2^{10} = 1024$. Before your time, computers had disks whose storage was measured in kilobytes or kb—units of 2^{10} or 1024 bytes—although kilo is the Greek work for thousand. Next $2^{20} = 1048576$, roughly a million: in computerese, this is a megabyte or mb. The next two are $2^{30} = 1073741824$ (a gigabyte or gb) and $2^{40} = 1099511627776$ (a terabyte or tb). By this point your calculator has probably already given up on an exact answer and started to use scientific notation.



3.3 Product sets

A 1tb drive is a pretty large hard drive even today. It's enough space to record address book entries (with a cell number and an email) for everyone on earth! But you'd need 5 such hard drives just to list the binary sequences of length 40! Each such sequence is 40 bits long (there are 40 0s or 1s) which is 5 bytes (a byte is 8 bits). And since there are 2^{40} such sequences they'd take up $5 \cdot 2^{40}$ bytes or 5tb). This kind of very rapid growth is why we don't want to describe sets by listing elements. In fact, even binary sequences of length 20, that is, the set 2^{20} , would make the point. This set has over a million elements so if I locked you in a room and refused to feed you until you'd listed them all, you *probably* wouldn't starve to death. But you'd need a whole lot of ruled pads and, since a million seconds is over 11 days, you'd be right peckish when you got out.

CHALLENGE 3.3.7: This is just to illustrate the limitations of even a calculator with such large numbers. Show that there are

1267650600228229401496703205376

binary sequences of length 100.

Hint: Your calculator can tell you what 2^{25} is. Also, $2^{50} = 2^{25} \cdot 2^{25}$ and $2^{100} = 2^{50} \cdot 2^{50}$. This is a challenge because your calculator cannot compute either of these products,

Products of sets

A recurring theme in mathematics is that of reducing complex problems to simpler ones by a strategy that goes by the name of “**Divide and Conquer**”. The term comes from a Roman military and political strategy, used against the Jewish Confederation and the Greek Achaean League, of dividing potentially difficult opponents into small groups that are individually easy to conquer. The mathematical paradigm involves an extra step. In addition to dividing—disassembling a big problem into smaller component problems—and conquering—solving the easier component problems individu-



ally, we need to have a way of reassembling the answer to the component problems into an answer to the original big problem. The key tools in both the dis- and re-assembly stages are usually operations that relate the components to the whole.

We're going to ease into our study of set operations by looking the product of sets. It's a good place to start because most of you are probably already familiar with it and it's the easiest to work with. In fact, we've been working with products for some time now without explicitly saying so. But we'll soon see that sets of sequences are just special kinds of product sets. Let's start with the easiest case.

PRODUCT OF TWO SETS 3.3.8: *The product $A \times B$ of two sets A and B is the set whose elements are all ordered pairs (a, b) in which a is an element of A and b is an element of B .*

As the term suggests, we consider that the order in which the elements a and b in an ordered pair are written to matter. This means that $(a, b) = (a', b')$ only if both the first and second components of the pair are equal: $a = a'$ and $b = b'$. In particular, (a, b) will almost never equal (b, a) : this happens if and only if $a = b$.

The product $A \times B$ is also often called the **Cartesian product** or the **direct product** of A and B . Since it's the only kind of product set we'll need to deal with, we'll just call it the product of A and B .

A very easy, but also very useful, formula tells us how many elements a product set has.

PRODUCT SET COUNTING FORMULA 3.3.9: *If A is a set with m elements and B is a set with n elements then the number of elements of the product set $A \times B$ is just the product $m \cdot n$ of the order of A times the order of B .*

This is almost too obvious to prove, but let's not let that stop us. For any fixed element a in A , there are exactly n pairs (a, b) in $A \times B$ —one for each of the n elements b in B . To count all the elements of

3.3 Product sets

$A \times B$, we need to add up one term n for each of the m elements a in A . Thus we find that $A \times B$ has

$$\underbrace{n + n + \cdots + n}_{m \text{ terms}}$$

elements and just need to recall that the product $m \cdot n$ is simply a shorthand for this sum.

EXAMPLE 3.3.10: Here’s an example we’ll work with later quite a bit, the standard deck D of playing cards shown below.




FIGURE 3.3.11: A standard deck of cards

Let $S = \{\spadesuit, \heartsuit, \diamondsuit, \clubsuit\}$ be the set consisting of the 4 **suits** in a standard deck of cards. In the order, I have written them—pretty standard, because its the ranking they have in the game of bridge—they are called spades, hearts, diamonds and clubs.

Let $V = \{A, 2, 3, 4, 5, 6, 7, 8, 9, 10, J, Q, K\}$ be the set consisting of the 13 **values** in a standard deck. The letters A, J, Q, and K stand for ace, jack, queen and king. Jacks, queens and kings are called **face cards** because carry portraits of persons. The other cards are called **spot** or numbered cards because they carry a corresponding number of “spots”: the Ace has just a single spot. But, in many games, like bridge, the Ace is treated as a sort of honorary face card that out-ranks all the other values while in others, especially poker, it has a

3.3 Product sets

sort of dual citizenship and can be treated as either the highest or lowest value.

The deck of cards D shown above is just the product $S \times V$. Note that, as predicted by [PRODUCT SET COUNTING FORMULA 3.3.9](#), the 4 suits (rows) of 13 values (columns) have a deck (table) with $4 \cdot 13 = 52$ cards.

There's no reason why we can't form a product of a set with itself. We just have to be extra careful to pay attention to the order in ordered pairs. There was not much chance of getting the pairs $(\heartsuit, 9)$ and $(9, \heartsuit)$ in the deck product above confused since the two components are so different. But, if we are considering the product of the $L \times L$ of the Latin alphabet with itself, then there's some risk of mixing up (e, i) and (i, e) if we don't pay attention. We'll usually abbreviate $A \times A$ as A^2 just as we would if A were a number instead of a set. We've already seen an example of such a self-product in [PROBLEM 3.2.8](#) where we formed the product D^2 of the set of numbers from 1 to 6 with itself, viewed as the pair of numbers on a blue die and a red die. One point that this example makes and that I want to underline here is that we are allowed to repeat a component in a product. In particular, D^2 contains ordered pairs like $(2, 2)$ and $(5, 5)$.

Going further, there's no reason why we can't take the product of more than 2 sets. It's easy to define the triple product $A \times B \times C$ to be the set of ordered triples (a, b, c) with $a \in A$, $b \in B$ and $c \in C$.

PROBLEM 3.3.12:

i) Show that if A has order m , B has order n and C has order p , then $A \times B \times C$ has order $m \cdot n \cdot p$.

Solution

The idea is to use [PRODUCT SET COUNTING FORMULA 3.3.9](#) twice. It first says that the set $D = B \times C$ has order $n \cdot p$. Then it says that the set $E = A \times D$ has order $m \cdot (n \cdot p)$. But an element of E is a pair (a, d) with $a \in A$ and $d \in D$ and since $D = B \times C$ this is the same as a pair $(a, (b, c))$ whose second element is another

3.3 Product sets

pair. By removing the interior parentheses, this is the same as a triple $(a, b, c) \in A \times B \times C$ so the order of this product must equal the order of E : $m \cdot (n \cdot p) = m \cdot n \cdot p$, once again, by simply removing the parentheses.

ii) What shorthand would you suggest for the triple product $A \times A \times A$?

PROBLEM 3.3.13: caterer offers a set A of 5 appetizers, a set E of 8 entrees and a set D of 6 desserts. To order your wedding banquet you have to pick one of each. What set describes the possible menus for your banquet and how many such menus are possible?

There's an easy way to summarize the formulas for the order of product sets in both [PRODUCT SET COUNTING FORMULA 3.3.9](#) and i) of [PROBLEM 3.3.12](#).

GENERAL PRODUCT SET COUNTING RULE 3.3.14: *The order of a product of sets is the product of the orders of the sets.*

PROBLEM 3.3.15: Joan has 23 pairs of shoes, 15 pairs of stockings, 18 skirts, 14 blouses and 9 hats. Each day she chooses her outfit by picking one item from each of these categories. Use a product set to describe Joan's outfits and determine how many different outfits she can select from.

Rules like [GENERAL PRODUCT SET COUNTING RULE 3.3.14](#) come up very often in mathematics—we'll see several more in this course—so I am going to digress briefly to make a few points about them. I'm going to name them.

BANG ZOOM RULES 3.3.16: *A bang zoom rule is any rule obtained from the statement, "The bang of the zoom(s) is the zoom of the bang(s)" by replacing the terms "bang" and "zoom".*

Mathematicians say this in a fancier way: banging and zooming *commute*. This term may be familiar to you in the context of a *single* operation: we say, for example, that addition or multiplication are

3.3 Product sets

commutative to mean that, if you perform several additions (or multiplications), you get the same answer regardless of what order you do them in. Here we say two *different* operations commute if they give the same answer in either orders. From [SECTION 1.1](#), you should realize that this is usually *not* the case. But when it is true, it makes life easier so we want to take note of the fact.

PROBLEM 3.3.17: [BANG ZOOM RULES 3.3.16](#) include the familiar distributive law $a(b+c) = ab+ac$. Show this by completing the sentence “The *multiple* of the *sum* is the ...”.

Why do I want to bother identifying [BANG ZOOM RULES 3.3.16](#)? There are two reasons. First, rules stated in “Bang Zoom” are much easier to state and remember: just compare the simple but general [GENERAL PRODUCT SET COUNTING RULE 3.3.14](#) with the complicated statements in [PRODUCT SET COUNTING FORMULA 3.3.9](#) and [PROBLEM 3.3.12](#) which only cover double and triple products. Second, when [BANG ZOOM RULES 3.3.16](#) are true when there are 2 zooms, they’re almost always true when there are 3, or 4, or any number of zooms. The reason is that the trick we used in [PROBLEM 3.3.12](#) of applying the 2-zoom rule twice almost always works in some form or other. In the future, I’m not going to bore you with checking such statements. Instead, once we understand the 2-zoom case, I’ll take it for granted that we know the rule in all cases and (if I feel any justification at all is needed) just cite the:

BANG ZOOM PRINCIPLE 3.3.18: *If a bang zoom rule is true when there are 2 zooms it is true when there are any number of zooms.*

PROBLEM 3.3.19: Check that the [BANG ZOOM PRINCIPLE 3.3.18](#) holds for the distributive law when there are 3 or 4 zooms—that is, 3 or 4 terms in the sum being multiplied.

Back to product sets. The notation needed to define a product of an arbitrary number ℓ of sets gets a bit cumbersome. Fortunately, we’ll only need such general product sets in the easy special case where

3.4 Power sets

they are all the same. What's more, we have already dealt with this case above!

What is an element of the triple product A^3 ? It's an ordered triple (a, a', a'') with each of a , a' and a'' in A . But observe that there is really no need for all that punctuation. If we just write $aa'a''$ we know everything we need to about the triple, including the order of the 3 components. Doesn't that $aa'a''$ look familiar? I hope so, because it's nothing more than a 3-letter *sequence* in the alphabet A . In both the ordered triples and the sequence, order matters and repeating a letter is allowed— $aa'a$ is another legal triple corresponding to (a, a', a) . So every ordered triple determines a unique sequence and vice versa; there's no need for duplicate of definitions and notation.

PRODUCTS OF A SET WITH ITSELF 3.3.20: *The ℓ -fold product A^ℓ of the sets A is the set of sequences of length ℓ in the alphabet A . It's then immediate from the [SEQUENCE COUNTING FORMULA 3.3.6](#) that, if A has m elements, then A^ℓ has m^ℓ elements.*

Why don't we think of all products as sequences? Because the components of an element of a general product come from different sets; for example, in [EXAMPLE 3.3.10](#), from the suits S and the values V . All the components—letters—in a sequence have to come from the *same* set, so we can only think of products of a set with *itself* as sequences. The good news is that almost all the examples we'll need to work are nonetheless simple, either because they involve just 2 (or occasionally 3) factor sets and then we just need to use ordered pairs or triples, or because they are products of a set with itself so that we can use sequences.

3.4 Power sets

The next operation for producing new sets from old comes from the subset relation.



3.4 Power sets

POWER SETS 3.4.1: *If A is any set, the power set $\mathcal{P}(A)$ is the set of all subsets of A . That is, the objects B that are elements of $\mathcal{P}(A)$ are exactly those that are subsets B of A .*

This sounds tricky but, perhaps for that very reason, it isn't. A few examples will make it easy to see what's going on.

EXAMPLE 3.4.2: What are the power sets of the sets $A_1 = \{1\}$, $A_2 = \{1, 2\}$ and $A_3 = \{1, 2, 3\}$?

Recall from [EXAMPLE 3.2.6 i\)](#) and [ii\)](#) that every set contains two trivial subsets, the empty set \emptyset and itself and from [iii\)](#) that these are the only subsets when the set, like A_1 has only a single element. So $\mathcal{P}(A_1) = \{\emptyset, \{1\}\}$. Could I have written, more simply $\{\emptyset, 1\}$? No! Why not? Because, 1 is an element of A_1 and elements of A_1 are *not* elements of $\mathcal{P}(A_1)$. Elements of $\mathcal{P}(A_1)$ are *subsets* of A_1 —so the subset $A = \{1\}$ itself *is* an element of $\mathcal{P}(A)$ but the element 1 of A is *not* an element of $\mathcal{P}(A)$.

PROBLEM 3.4.3: Let $A = \{\emptyset\}$ and let $A' = \{\emptyset, \{\emptyset\}\}$. How many elements does A have? (The answer is *not* 0!) Does $A = \emptyset$? (Do these two sets have the same number of elements?) How many elements does A' have?

Don't worry. Once you grasp that the power set $\mathcal{P}(A)$ is a collection of sets—the subsets of A —and can contain no “naked” elements, everything's easy. The power set of A_2 is $\mathcal{P}(A_2) = \{\emptyset, \{1\}, \{2\}, \{1, 2\}\}$. Notice that the first two subsets in this list are the subsets of A_1 , and the next two are just these subsets with the new element 2 added. The subset of A_3 will likewise consists of the subsets of A_2 , and these subsets with the new element 3 added. So

$$\mathcal{P}(A_3) = \{\emptyset, \{1\}, \{2\}, \{1, 2\}, \{3\}, \{1, 3\}, \{2, 3\}, \{1, 2, 3\}\}.$$

PROBLEM 3.4.4: What is the power set of the set $A_4 = \{1, 2, 3, 4\}$?
Hint: Add one element to each set in $\mathcal{P}(A_3)$.

3.4 Power sets

You may have noticed that the sets powers sets of A_1 , A_2 , and A_3 have 2, 4 and 8 elements and, if you computed it correctly, your power set of A_4 has 16. Of course, these numbers are 2^1 , 2^2 , 2^3 and 2^4 . I hope that this is reminding you of [EXAMPLE 3.3.5](#) above. It doesn't take much nerve to predict the:

POWER SET COUNTING FORMULA 3.4.5: *If A is any set with m elements, then A has exactly 2^m subsets. That is the power set $\mathcal{P}(A)$ has exactly 2^m elements.*

To check this prediction, we just need to list the subsets of A “one element at a time” just as we did in [EXAMPLE 3.4.2](#) and [PROBLEM 3.4.4](#). Each time we add an element we get two new subsets out of each old one. So when we have added all m elements the number of subsets will have doubled ℓ times to 2^m in all.

Is is a coincidence that the number of binary sequences of length ℓ is the same as the number of subsets of a set A with m elements? Not at all. We can build a binary sequence s from a subset B as follows. The first bit of the sequence s is 1 if the first element of A is in the subset B , and 0 if this element is not in B . Likewise, The second bit of the sequence s is 1 if the second element of A is in the subset B , and 0 if this element is not in B . And so on: I'm cheating a bit since I'm talking about the first element of A and order does not matter for sets. This is harmless; whatever order the elements of A come in will do. For example, if $A_5 = \{d, a, c, e, b\}$ then the subset $B = \{a, c, b\}$ gives the sequence $s = 01101$ and the sequence $s' = 11000$ gives the subset $B' = \{d, a\}$. Since we can pair up subsets B and sequences s in this way—with every subset paired with exactly 1 binary sequence—the number of subsets of A and the number of binary sequences of length m must be the same.

We can go even further. What we have really done by pairing up subsets and sequences is to show that the power set $\mathcal{P}(A)$ of a set A with m elements and the sequence set 2^m are pretty much the same; it's just that the objects have different names, subsets in one and

3.4 Power sets

sequences in the other. For this reason, mathematicians often like to write $\mathcal{P}(A)$ as 2^A . This is the reason that the set $\mathcal{P}(A)$ is called a *power set*.

We'll continue to write $\mathcal{P}(A)$, so we can save exponential notation for self-products or sets of sequences and avoid any possible confusion. Note the key difference. If A is a set with m elements, then the ℓ -fold *product* of A with itself is the set A^ℓ which has m^ℓ elements, but the *power set* $\mathcal{P}(A)$ can be viewed as the set 2^m which has 2^m elements. In a self-product the order m of A is the base of the exponential. In a power set, the the order m of A is the exponent.

OK. Time for a slightly nasty problem.

PROBLEM 3.4.6:

i) What is the power set $\mathcal{P}(\emptyset)$ of the empty set?

Hint: the empty set has $\ell = 0$ elements so it must have $2^\ell = 2^0 = 1$ subset. What is this set?

ii) What is the power set $\mathcal{P}(\mathcal{P}(\emptyset))$? Hint: $\mathcal{P}(\emptyset)$ has 1 element so it has $2^1 = 2$ subsets.

Hint: The answer is in an earlier problem in this section.

iii) What is the power set $\mathcal{P}(\mathcal{P}(\mathcal{P}(\emptyset)))$?

Infinites and an argument from *The Book*

Let's next ask a question that illustrates the power of the power set operation (if you'll pardon the pun). We've seen that for any finite set A , the power set of A has "a lot more" elements than A . In other words, for any positive ℓ , 2^ℓ is bigger than ℓ , usually a whole lot bigger. For example, we saw in the preceding subsection that when $\ell = 40$, $2^\ell = 1099511627776$.

What happens if A is infinite? Is $\mathcal{P}(A)$ still bigger than A in some sense? More generally, is all we can say about two infinite sets that they're both infinite or is there some way to compare them? If there



3.4 Power sets

is a way to compare, then there must be different *sizes* of infinity. Such a notion at first seems counter to our intuition.

What if I try to compare the infinity of the sets \mathbb{N} and $\mathbb{N} \times \mathbb{N}$?

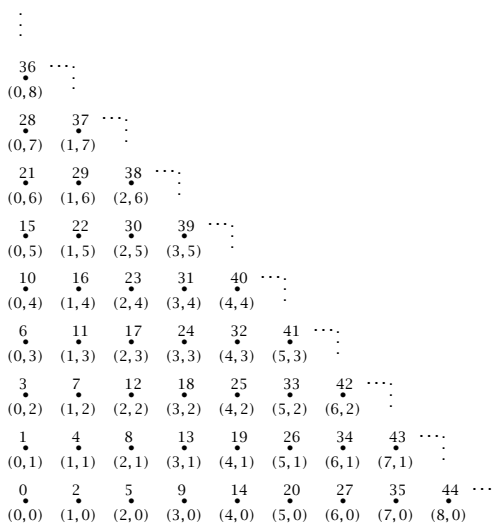


FIGURE 3.4.7: Comparing the infinities of \mathbb{N} and $\mathbb{N} \times \mathbb{N}$

FIGURE 3.4.7 shows a picture of both sets in the xy -plane. The elements of \mathbb{N} are the dots in the bottom row. Every dot is an element of $\mathbb{N} \times \mathbb{N}$. So there are *infinitely* many elements of $\mathbb{N} \times \mathbb{N}$ above every element of \mathbb{N} . Surely this means that the order of $\mathbb{N} \times \mathbb{N}$ is a bigger infinity than that of \mathbb{N} .

Not so, as FIGURE 3.4.7 also shows. By following the diagonals down-and-right from the y -axis to the x -axis, each time returning to the next highest point on the y -axis, we will pass by all the points of $\mathbb{N} \times \mathbb{N}$. Suppose, as we go by, we drop successive whole numbers from \mathbb{N} at each point: these numbers are shown above the point. Then, we eventually pair up every pair in $\mathbb{N} \times \mathbb{N}$ with a different number in \mathbb{N} and vice versa. So these sets have the *same* order—that is, the same infinity! You can do basically the same thing with $\mathbb{N} \times \mathbb{N} \times \mathbb{N}$ and with

3.4 Power sets

$\mathbb{N} \times \mathbb{N} \times \mathbb{N} \times \mathbb{N}$ and so on. You can even pair off \mathbb{N} and \mathbb{Q} : the set \mathbb{Q} of rational numbers is somehow much “thicker” than any of the other examples. For example, you can’t draw even the part of \mathbb{Q} that lies between 0 and 1: you can start with the $\frac{1}{2}$, then $\frac{1}{3}$, $\frac{2}{3}$, then $\frac{1}{4}$ and $\frac{3}{4}$, then $\frac{1}{5}$, $\frac{2}{5}$, $\frac{3}{5}$ and $\frac{4}{5}$ and so on, but, however long you continue you’ve always left out fractions with infinitely many denominators. So now it certainly seems like we’re back to square one: all infinities are the same.

Yet, it turns out that there *are* different infinities. Even infinitely many of them! Working out this kind of question is one of the things that mathematicians spent all those years on when they were try to understand the fine points of working with sets. Even today, many basic questions remain unsolved.

However, there’s one argument that I cannot resist including it here¹ It an argument so simple and elegant that it’s certainly in The Book, in which the great Hungarian mathematician **Paul Erdős** imagined that God collected the most beautiful arguments in mathematics. It’s Cantor’s Diagonalization Argument, named after **Georg Cantor**, the German mathematician who came up with it in 1891. Don’t be put off by the fancy sounding name, because the idea is very simple. It’s really just the argument we used to prove the **POWER SET COUNTING FORMULA 3.4.5** above. The argument can be used to show that for any infinite set A , the power set $\mathcal{P}(A)$ has a *bigger* infinity of elements.

I’m just going to give Cantor’s argument for the set \mathbb{N} of natural numbers which I’ll think of in its usual order as $\mathbb{N} = \{0, 1, 2, 3, \dots\}$. I’ll encode a subset B of \mathbb{N} as a binary sequence as we did for finite sets above. The only difference is that because \mathbb{N} is infinite, I’ll need to use a set S of infinite sequences $s = b_0b_1b_2b_3 \dots b_{n-1}b_nb_{n+1} \dots$. We can convert such a sequence s to a subset B exactly as before bit

¹Let me answer the question I’m sure is on many lips. This will *not* be on the exam. So, if you’re a member of one of the invertebrate orders, feel free to go watch that episode of *American Idol* that cued up on your DVR and come back at the end of this section.

3.4 Power sets

b_n is 1 if n is an element of the subset B and 0 if n is not an element of B . For example, the empty set is given by $s = 0000 \dots 000 \dots$, the subset \mathbb{N} itself by $s = 1111 \dots 111 \dots$ and the subset E of even numbers by $s = 10101010 \dots$

I claim that there's no way to pair up the natural numbers \mathbb{N} its power set $\mathcal{P}(\mathbb{N})$ —or, what's the same to pair \mathbb{N} and the set S of infinite sequences. We can try but there'll always be unpaired or “left over” sequences, just as there would be if we tried to pair up a finite set with a *larger* finite set. For example, if we say, tried to pair off a set with 14 against one with 15 elements, we'd always have 1 element left over. Intuitively obvious as this is, we need to know something about how the pairing is carried out to have any hope of putting a fork into one the leftover element. What makes Cantor's argument so amazing is that he manages to identify a leftover sequence without knowing *anything* about how we've tried to pair up \mathbb{N} and S . The reason he's able to do this, is that, the infinity of S is so much bigger than that of \mathbb{N} that just about *every* sequence in S is leftover. Precisely, the infinity of the leftovers in S is a big as the infinity of all of S

Suppose we try to pair up the natural numbers \mathbb{N} and the set S of infinite sequences by associating 0 and s^0 , 1 and s^1 and so on. Then what we'd have would be a list of *all* the elements of S . That is, we'd have an infinite list $L = [s^0, s^1, s^2, s^3, \dots, s_{n-1} s_n s_{n+1}]$ that included *every* sequence $s \in S$.

Let's lay out this list L more carefully so we can see the bits that make up each sequence:

3.4 Power sets

s^0	=	b_0^0	b_1^0	b_2^0	b_3^0	...	b_{n-1}^0	b_n^0	b_{n+1}^0	...
s^1	=	b_0^1	b_1^1	b_2^1	b_3^1	...	b_{n-1}^1	b_n^1	b_{n+1}^1	...
s^2	=	b_0^2	b_1^2	b_2^2	b_3^2	...	b_{n-1}^2	b_n^2	b_{n+1}^2	...
s^3	=	b_0^3	b_1^3	b_2^3	b_3^3	...	b_{n-1}^3	b_n^3	b_{n+1}^3	...
\vdots	\vdots	\vdots	\vdots	\vdots	\vdots	\vdots	\vdots	\vdots	\vdots	\vdots
s^{n-1}	=	b_0^{n-1}	b_1^{n-1}	b_2^{n-1}	b_3^{n-1}	...	b_{n-1}^{n-1}	b_n^{n-1}	b_{n+1}^{n-1}	...
s^n	=	b_0^n	b_1^n	b_2^n	b_3^n	...	b_{n-1}^n	b_n^n	b_{n+1}^n	...
s^{n+1}	=	b_0^{n+1}	b_1^{n+1}	b_2^{n+1}	b_3^{n+1}	...	b_{n-1}^{n+1}	b_n^{n+1}	b_{n+1}^{n+1}	...

Cantor now, by a brilliantly simple idea, produces a sequence t in S that *cannot* be on the list L above. That’s just what we’re trying to see: that *no* list L can contain *all* the sequences in S .

The trick is to define a bit c_n to be the *opposite* of the red bit b_n^n above: that is, if $b_n^n = 1$, then $c_n = 0$ and if $b_n^n = 0$, then $c_n = 1$. Then, the sequence $t = c_0c_1c_2c_3 \dots c_{n-1}c_nc_{n+1} \dots$ can’t be *anywhere* in the list L . We can’t have $t = s^0$ because Cantor has set things up so that the 0th bits of t and s^0 —that is, c_0 and b_0^0 are opposite. We can’t have $t = s^1$ because Cantor has set things up so that the 1st bits of t and s^1 —that is, c_1 and b_1^1 are opposite. In general, we can’t have $t = s^n$ because Cantor has set things up so that the 0th bits of t and s^0 —that is, c_n and b_n^n are opposite. So t can’t be anywhere in the list.

OK, the good news is that it won’t matter in the rest of the course if you forget everything I’ve just said about infinite sets. But I hope you won’t, as Cantor’s argument is as elegant and beautiful an achievement of the mind as any Shakespeare sonnet or Giotto fresco. You may have no use for mathematics after you leave this course, but if you didn’t didn’t get at least a tiny tingle from seeing his ideas, I’m afraid your cord just doesn’t reach all the way to the outlet.

Cantor’s proof leaves open one question. Can we find a set A with the larger *uncountable* infinity of $\mathcal{P}(\mathbb{N})$ somewhere in nature? Not only is the answer yes, but you already know A ! If we simply put

3.4 Power sets

a decimal point to the left of each infinite binary sequence in the proof, what we'll get are all the real numbers in the interval between 0 and 1. For example, the sequence .0000... encoding the empty set gives the real number 0, the sequence .1111... encoding all of \mathbb{N} gives the number 1 and the sequence .1010101... encoding the set E of even numbers gives $\frac{2}{3}$. We can check these last values using the [GEOMETRIC SERIES FORMULA 1.3.6](#).

PROBLEM 3.4.8: Recall that a binary decimal $.b_1b_2b_3b_4\dots$ is a shorthand for the series $\frac{b_1}{2^1} + \frac{b_2}{2^2} + \frac{b_3}{2^3} + \frac{b_4}{2^4} + \dots$. Most such series are not geometric but a few are. Use the [GEOMETRIC SERIES FORMULA 1.3.6](#) to show that:

- i) $.1111\dots = \frac{1}{2^1} + \frac{1}{2^2} + \frac{1}{2^3} + \frac{1}{2^4} + \dots = 1$; and,
- ii) $.10101\dots = \frac{1}{2^1} + \frac{0}{2^2} + \frac{1}{2^3} + \frac{0}{2^4} + \frac{1}{2^5} + \dots = \frac{1}{2} + \frac{1}{2 \cdot 4} + \frac{1}{2 \cdot 4^2} + \dots = \frac{2}{3}$.

In fact, the points of any line segment form an uncountably infinite set. How does the infinity of a line segment compare to the infinity of the entire real line, or the infinity of the whole $(x - y)$ -plane of analytic geometry. The very counter-intuitive answer is that the entire plane is no bigger than the shortest line segment!

All this talk of different sizes of infinity may seem like the sort of mathematics that could never interest anyone but the most committed mathematician. Not so. Today, it's probably computer scientists who worry most about these issues. The reason is that computers are inherently finite devices. With "enough" resources, they can compute infinite many answers—even we, if we could live "long enough", would be able to count to any of the infinitely many natural numbers. But arguments like Cantor's show that the infinities that we, or computers, can explore are limited to the small, countable \mathbb{N} type. For example, most real numbers can't be computed—there are too many. That's forced computer scientists to think hard about exactly what they can compute. But understanding this would take us too far from our path so I won't try to explain it further here.

Combinations

Before we close this section, I want to do a bit more counting with power sets. We already have lots of evidence that the most important property of a set is its order or number of elements. So it's natural to feel that two subsets of a set A that have the same number of elements are somehow similar. Our next goal is to use this idea to group the elements of $\mathcal{P}(A)$ —the subsets of A —by the number of elements they contain.

SUBSETS OF ORDER ℓ 3.4.9: *The set $\mathcal{P}(A)_\ell$ is the set whose elements are the subsets B of A having order exactly ℓ . Since every subset B of A is an element of $\mathcal{P}(A)$, the set $\mathcal{P}(A)_\ell$ is a subset of $\mathcal{P}(A)$. More informally, we'll call $\mathcal{P}(A)_\ell$ the set of subsets of order ℓ of the set A .*

A few comments. Suppose that A has order m . Then the order of any subset B of A is between 0 and m . Thus, every subset B —every element of $\mathcal{P}(A)$ —lies in one of the sets $\mathcal{P}(A)_0, \mathcal{P}(A)_1, \mathcal{P}(A)_2, \dots, \mathcal{P}(A)_m$. On the other hand, these sets are disjoint: no subset B can be in two of them because then B would have 2 different orders. The sets $\mathcal{P}(A)_\ell$ thus cut up or *partition* $\mathcal{P}(A)$. We'll need this concept in probability so let's record it.

PARTITION OF S 3.4.10: *A collection of subsets of a set S forms a partition of S if every element of S lies in exactly one of the subsets. Another way to say this is to say that the sets are pairwise disjoint (no two can have a common element) and every element of S lies in one of them.*

The next point to note is that order of $\mathcal{P}(A)_\ell$ —the *number* of subsets B of A that themselves have order ℓ —depends only on the order m of A and the order ℓ . What if A' is another set with order m ? This simply means that we can pair up the *elements* a of A with the elements a' of A' . But then we can also pair up the *subsets* B of A with subsets B' of A' by the rule that $a' \in B'$ if and only if $a \in B$. This should be no surprise since $\mathcal{P}(A)$ and $\mathcal{P}(A')$ both have order 2^m by

3.4 Power sets

the **POWER SET COUNTING FORMULA 3.4.5**. However, this pairing of subsets has an extra nice property: every A -subset B with ℓ elements a gets paired with an A' -subset B' with ℓ elements a' . Hence the sets $\mathcal{P}(A)_\ell$ and $\mathcal{P}(A')_\ell$ consisting of all such ℓ -element subsets have the same order.

PROBLEM 3.4.11: List all the subsets of the sets $A = \{1, 2, 3\}$ and $A' = \{1', 2', 3'\}$ and check that the pairing $1 \leftrightarrow 1'$, $2 \leftrightarrow 2'$ and $3 \leftrightarrow 3'$ also pairs the sets $\mathcal{P}(A)_\ell$ and $\mathcal{P}(A')_\ell$ for $\ell = 0, 1, 2$ and 3 .

The point of these observations is that we can speak about the number of ℓ -element subsets of a set A with m elements without worrying about what the set A is.

BINOMIAL COEFFICIENTS AND COMBINATIONS 3.4.12: *The binomial coefficient $\binom{m}{\ell}$ —read “ m choose ℓ ”—is defined to be the number of subsets B of order ℓ in any (and every) set A of order m . The count $\binom{m}{\ell}$ is also often called a **combination**, written $C(m, \ell)$ or $mC\ell$ and read “combination $m \ell$ ” or “ m combination ℓ ”.*

The idea behind the use of the word “choose” above is that we are counting the ways of *choosing* a k -element subset from a set with m elements. This idea is often expressed, more loosely, as “*choosing* k -elements from amongst m ”. I put those quotes in the second version because, as we’ll see shortly, such choices come in several flavors, and the subset flavor, that we’re dealing with now is just the most important.

Mathematicians prefer the binomial coefficient notation to combinations and that’s what I’ll mainly use here. But there’s a good reason for knowing both: your calculator probably has the combinations version built in. On the TI-8x series, it’s found on the Math panel as nCr . To use it, you enter the order m , select the nCr row, enter the order ℓ and press ENTER. But we’ll soon see that when m and ℓ are not too big, it’s often faster to work out combinations by hand and the most common combinations are easy enough to simply remember.



3.4 Power sets

Let's try to find a formula for $\binom{m}{\ell}$ by bootstrapping our way up. Here I'll take $m = 3$ to keep the number of subsets small and we'll try to work our way up from the smallest ℓ — $\ell = 0$ —to the biggest $\ell = 3$. It doesn't matter what master set A of order 3 I take so let's take $A = \{a, b, c\}$. Then it's easy to write out the sets $\mathcal{P}(A)_\ell$ for reference.

$$\mathcal{P}(A)_0 = \{\emptyset\}$$

$$\mathcal{P}(A)_1 = \{\{a\}, \{b\}, \{c\}\}$$

$$\mathcal{P}(A)_2 = \{\{a, b\}, \{a, c\}, \{b, c\}\}$$

$$\mathcal{P}(A)_3 = \{\{a, b, c\}\}$$

Now, let's start trying to count $\mathcal{P}(A)_\ell$ by pure thought. I'll try to prediction how many there are for *any* m , then we'll check my prediction for $m = 3$. First, $\mathcal{P}(A)_0$ is pretty easy: it always has only 1 element, the empty set \emptyset for any A .

It's not much harder to count $\mathcal{P}(A)_1$. A 1-element subset is determined by its unique element and there are m of these if A has order m . Sure enough there are 3 here where $m = 3$. I'll record this as $\binom{m}{1} = \frac{m}{1}$; my reason for inserting the apparently pointless denominator 1 will become clear in a moment.

It's for $\ell = 2$ that you have to think a bit. A set with two elements can be obtained by adding 1 more element to a set with 1 element. The first point that calls for a bit of care is that there are only $m - 1$ choices for this second element. Why? If my 1-element subset is $\{a\}$, I better not choose to add a or I won't end up with 2 elements. This was the point of [EXAMPLE 3.1.5](#). So 3 choices for the 1-element set and 2 for the second element makes 6 subsets with 2-elements. Except, of course, there are only three.

What went wrong? We can see if we carry out my choosing procedure. It does yield 6 subsets: from $\{a\}$ we get $\{a, b\}$ and $\{a, c\}$, from $\{b\}$ we get $\{b, a\}$ and $\{b, c\}$, and from $\{c\}$ we get $\{c, a\}$ and $\{c, b\}$. The problem is that each of these sets appears twice, with the elements listed in the opposite order. So in addition to multiplying by $(m - 1)$ for my second element, I need to also divide by 2 to com-

3.4 Power sets

pensate for the duplication. This division depends only on the fact that my subsets have 2 elements so it'll apply for any m . So my revised prediction is that $\binom{m}{2} = \frac{m}{1} \frac{m-1}{2}$ —I am multiplying the number of 1-element subsets by the effective number of different 2-element subsets I get from each by adding an element. For $m = 4$ this predicts $\binom{4}{2} = \frac{4}{1} \frac{3}{2} = 6$.

PROBLEM 3.4.13: Check that $A' = \{a, b, c, d\}$ has exactly 6 subsets with 2 elements.

Note that, even though the fraction $\frac{3}{2}$ is not a whole number, the product of fractions $\frac{4}{1} \frac{3}{2}$ is—as it has better be, since it's the number $\binom{4}{2}$ of 2-element subsets of a 4-element set and that is a whole number by definition.

We're almost done as $\ell = 3$ poses no new problems. I get one 3-element subset from each 2 element subset by adding one of the $m - 2$ elements not already inside. But now each 3-element subset will come up, not twice, but *three* times—once with each of its 3 elements as the “latest addition”. Let's check: from $\{a, b\}$ we get one 3 element set $\{a, b, c\}$ (because $(3 - 2) = 1$), from $\{a, c\}$ we get $\{a, c, b\}$ and from $\{b, c\}$ we get $\{b, c, a\}$. And, sure enough, we've got the same 3 element subset 3 times, once with each element as the “latest addition”. So I predict $\binom{m}{3} = \frac{m}{1} \frac{m-1}{2} \frac{m-2}{3}$.

PROBLEM 3.4.14: Check that $A' = \{a, b, c, d\}$ has exactly $\frac{4}{1} \frac{3}{2} \frac{2}{3}$ subsets with 3 elements.

Now, I hope you see why I started out with a fraction. It makes the pattern pretty clear. It looks like $\binom{m}{\ell}$ is a product of ℓ very predictable fractions. You start out with $\frac{m}{1}$ and then successively subtract 1 from the numerator and add 1 to the denominator to get the next fraction. If I'm right, then $\binom{m}{4} = \frac{m}{1} \frac{m-1}{2} \frac{m-2}{3} \frac{m-3}{4}$.

One quick check is to plug in $m = 3$ when we get $\frac{3}{1} \frac{2}{2} \frac{1}{3} \frac{0}{4} = 0$ as we better since a set with 3 elements has no subsets with 4. A slightly more convincing one is to plug in $m = 4$ getting $\frac{4}{1} \frac{3}{2} \frac{2}{3} \frac{1}{4} = 1$ which is



3.4 Power sets

right, because the only 4 element subset of, say A' is A' itself. The best check of all is that the argument we've been using still applies. To each 3 element subset we can add any of the $m - 3$ elements not in it to get a 4 element subset, and doing so will produce each 4 element subset 4 times, once with each of its elements as the "latest addition".

COMBINATION FORMULA 3.4.15: *If $\ell > 0$, then*

$$\binom{m}{\ell} = C(m, \ell) = \frac{m}{1} \cdot \frac{m-1}{2} \cdot \frac{m-2}{3} \cdots \frac{m-\ell+2}{\ell-1} \cdot \frac{m-\ell+1}{\ell}$$

and $\binom{m}{0} = C(m, 0) = 1$.

What about that last case $\ell = 0$? Do we even need to bother with this? Yes: in fact, this case is needed frequently in applications. But it's easy to handle: there's only one set with 0 elements, the empty set \emptyset and it's a subset of every set A which is why $\binom{m}{0} = C(m, 0) = 1$ for every m .

In practice, it's much better to view the [COMBINATION FORMULA 3.4.15](#) this as a method than a formula. The reason is that it's much easier and much more reliable to learn and use the method than to do the same for the formula.

METHOD FOR COMPUTING COMBINATIONS 3.4.16:

Step 1: Start with the fraction $\frac{m}{1}$ with m the size of the master set A .

Step 2: Keep lowering the numerator by 1 and raising the denominator by 1 to get the next fraction to multiply by.

Step 3: Stop when the denominator of your current fraction is ℓ , the size of the subsets you want to count.

PROBLEM 3.4.17: Write out the $2^5 = 32$ subsets of the set $A = \{a, b, c, d, e\}$ and use [METHOD FOR COMPUTING COMBINATIONS 3.4.16](#) to check that the [COMBINATION FORMULA 3.4.15](#) correctly predicts the number of subsets each order ℓ .

3.4 Power sets

There's another beautiful method for finding binomial coefficients which has the very nice feature that it involves no multiplications or divisions, just additions. This is the famous triangle of the great 17th century French mathematician Blaise Pascal.

TABLE 3.4.18 shows how it starts. In addition to the values of the binomial coefficients or combinations which are shown in black, I have indicated, in binomial coefficient form the values of m (in red) and ℓ giving the value. The rows are indexed by the size m of the master set starting with $m = 0$. Thus the fourth row contains the combinations with $m = 3$ that we found above. In each row, the coefficient $\binom{m}{0}$ with $\ell = 0$ is located at the left and as you move to the right ℓ increases by 1 each time, ending with $\ell = m$.

												$\binom{0}{0}$:1																												
												$\binom{1}{0}$:1		$\binom{1}{1}$:1																										
												$\binom{2}{0}$:1		$\binom{2}{1}$:2		$\binom{2}{2}$:1																								
												$\binom{3}{0}$:1		$\binom{3}{1}$:3		$\binom{3}{2}$:3		$\binom{3}{3}$:1																						
												$\binom{4}{0}$:1		$\binom{4}{1}$:4		$\binom{4}{2}$:6		$\binom{4}{3}$:4		$\binom{4}{4}$:1																				
												$\binom{5}{0}$:1		$\binom{5}{1}$:5		$\binom{5}{2}$:10		$\binom{5}{3}$:10		$\binom{5}{4}$:5		$\binom{5}{5}$:1																		
												$\binom{6}{0}$:1		$\binom{6}{1}$:6		$\binom{6}{2}$:15		$\binom{6}{3}$:20		$\binom{6}{4}$:15		$\binom{6}{5}$:6		$\binom{6}{6}$:1																
												$\binom{7}{0}$:1		$\binom{7}{1}$:7		$\binom{7}{2}$:21		$\binom{7}{3}$:35		$\binom{7}{4}$:35		$\binom{7}{5}$:21		$\binom{7}{6}$:7		$\binom{7}{7}$:1														
												$\binom{8}{0}$:1		$\binom{8}{1}$:8		$\binom{8}{2}$:28		$\binom{8}{3}$:56		$\binom{8}{4}$:70		$\binom{8}{5}$:56		$\binom{8}{6}$:28		$\binom{8}{7}$:8		$\binom{8}{8}$:1												

TABLE 3.4.18: PASCAL'S TRIANGLE

PROBLEM 3.4.19:

- i) What values of m and ℓ give the two entries equal to 15 above?
- ii) Pascal's triangle predicts that $\binom{7}{3} = 35 = \binom{7}{4}$ and that $\binom{8}{4} = 70$. Check these values using the METHOD FOR COMPUTING COMBINATIONS 3.4.16.

How is Pascal's triangle built? By an incredibly simple rule. You get the value of each entry by summing the two entries immediately above it and to its left or right. This works even if one of the entries

3.4 Power sets

above is missing, as long as we interpret missing entries as having value 0. Let's check the $m = 7$ row (the eighth!) by this rule. It starts with a 1 which is the sum of a blank entry (0) above and to the left and a 1 above and to the right. Next we get a 7 as the sum of 1 and 6, then a 21 as the sum of 6 and 15, then a 35 as the sum of 15 and 20. The same pattern repeats, in the opposite order in the right half of this row.

PROBLEM 3.4.20:

- Use the building rule to write down the next two rows of the triangle, corresponding to $m = 9$ and $m = 10$.
- Check the entries giving the combinations $\binom{9}{4}$ and $\binom{10}{6}$ using the [METHOD FOR COMPUTING COMBINATIONS 3.4.16](#).

Why does Pascal's rule work? To see the reason, we first need to see what 2 binomial coefficients are used to compute $\binom{m}{\ell}$. Since both lie in the next row up, both involve a master set with $m - 1$ elements instead of m . The two subsets whose binomial coefficients are used have sizes $\ell - 1$ to the left and ℓ to the right. You can check this easily by counting in from the left in the row above: an example with $m = 8$ and $\ell = 4$ is given by [ii\) of PROBLEM 3.4.19](#).

Thus Pascal's rule amounts to the prediction that:

PASCAL'S IDENTITY 3.4.21:

$$\binom{m}{\ell} = \binom{m-1}{\ell} + \binom{m-1}{\ell-1}.$$

This can be checked by some slightly messy algebra with the three instance of the [COMBINATION FORMULA 3.4.15](#) in the identity. It's not so hard so I'll leave this approach to you as a challenge problem.

CHALLENGE 3.4.22: Give this algebraic proof. Hint: First, write down the right hand side $\binom{m-1}{\ell} + \binom{m-1}{\ell-1}$ using the [COMBINATION FORMULA 3.4.15](#). Then, put the two fractions involved over a common denominator and simplify the resulting numerator to obtain $\binom{m}{\ell}$

3.4 Power sets

But there's a wonderfully simple way to check [PASCAL'S IDENTITY 3.4.21](#) with no algebra at all. Let's use the set $A = \{1, 2, \dots, m-1, m\}$ as our master set with m elements, and let's denote by A' the subset $A' = \{1, 2, \dots, m-1\}$ of order $(m-1)$ with the element m removed. We simply observe that each of the $\binom{m}{\ell}$ subsets B of A with ℓ elements either contains the element m or it doesn't! In the former case, the set B' obtained by removing the element m from B is a subset of A' with $(\ell-1)$ elements: there are $\binom{m-1}{\ell-1}$ of these. In the latter, we can view B itself as a subset of A' with ℓ elements: there are $\binom{m-1}{\ell}$. Presto!

Pascal's triangle also makes clear one other fact about combinations that's often computationally and theoretically useful, and that is not evident from the [COMBINATION FORMULA 3.4.15](#). The fact that Pascal's triangle is symmetric about the vertical line that bisects it amounts to the identity:

SYMMETRY OF BINOMIAL COEFFICIENTS 3.4.23:

$$\binom{m}{\ell} = \binom{m}{m-\ell}.$$

We can also see this directly. Every ℓ element subset B of a set A with m elements determines a unique subset B' with $(m-1)$ elements and vice versa: the elements of B' are just the $(m-\ell)$ elements of A that are *not* elements of B .

PROBLEM 3.4.24:

- What would you *not* want to find $\binom{100}{97}$ using the [COMBINATION FORMULA 3.4.15](#) directly?
- Use the [SYMMETRY OF BINOMIAL COEFFICIENTS 3.4.23](#), to avoid this problem and find $\binom{100}{97}$ indirectly.

Finally, the triangle makes it clear that some common combinations—the two “outside” diagonals—are so simple we can just remember them.



SIMPLE BINOMIAL COEFFICIENTS 3.4.25: For any m , $\binom{m}{0} = \binom{m}{m} = 1$ and $\binom{m}{1} = \binom{m}{m-1} = m$.

3.5 Lists and permutations

In this section, we'll study lists, which were mentioned briefly in [SETS VERSUS LISTS 3.1.2](#), in full detail and learn how to count them using permutations. Let's start with the fancy definition. Then we'll pare it down to the bare-bones version we'll use in practice.

Lists: sets with an order

LIST-FORMAL DEFINITION 3.5.1: A list L of length ℓ is a finite set B or order ℓ together with a pairing of B with the set $\mathbf{l} = \{1, 2, 3, \dots, \ell - 1, \ell\}$.

We view the set \mathbf{l} is a sort of National Bureau of Standards reference set of order ℓ . The pairing between the list L and the standard set \mathbf{l} is some extra information. To see what this information amounts to, let's look at an example. This example will also make clear that there are *many* lists L with the *same* underlying set B .

EXAMPLE 3.5.2: Consider the list L for which the set $B = \{a, b, c, d\}$ and the pairing with $\mathbf{4}$ is $a \leftrightarrow 1, b \leftrightarrow 2, c \leftrightarrow 3$ and $d \leftrightarrow 4$.

What the pairing lets us do is speak about the first element of L —it's a the element paired with the number 1; or the fourth— d because it's paired with 4; or the second which is ...? Right, b because it's paired with 2. In other words, the pairing gives us a ordering or ranking the elements of the set A .

Conversely, a first-to-fourth ordering on the elements of B determines a pairing of B with $\mathbf{4}$. The first element of B pairs with 1, the

3.5 Lists and permutations

second with 2 and so on. There are *lots* of such orderings—24 to be precise as we’ll see in the next subsection. The ordering b, d, a, c corresponds to the pairing $b \leftrightarrow 1, d \leftrightarrow 2, a \leftrightarrow 3$ and $c \leftrightarrow 4$ and defines a *different* list L' based on the same set B , the ordering c, b, a, d defines a third list L'' and so on.

It’s obvious from this example that it’s much easier to think of B -lists in terms of orderings on B than in terms of pairing of B with \mathbf{I} . That’s the working definition we’re going to adopt. All we need to complete it is a notation that distinguishes lists from plain sets, which remember do *not* have any preferred order. We’ll use square brackets—[] to surround lists to make the difference clear.

LIST: WORKING DEFINITION 3.5.3: *A list L of length ℓ is a finite set B of order ℓ together with the of a first-to-last order on the elements of B . We denote such a list L by listing the elements of B between square brackets with the first element of the left and the last on the right.*

EXAMPLE 3.5.4: In the notation of the working definition, the lists of [EXAMPLE 3.5.2](#) are $L = [a, b, c, d]$, $L' = [b, d, a, c]$, and $L'' = [c, b, a, d]$.

PROBLEM 3.5.5: There are 6 possible ways of ordering a set with 3 elements. Write down the 6 lists L whose set $B = \{1, 2, 3\}$.

EXAMPLE 3.5.6: One word about orders. As we go forward, we’ll see many ways of specifying an order, and sometimes it’s not so obvious that an order is what’s involved. In fact, we’ll often view ourselves as having chosen an order—even if we have now written down this order—if we simply want to view our choices as being *altered* by any reordering. Here are a few examples to keep in mind, starting with some that clearly involve an ordering.

- i) Standings in a football league are an ordering of the teams in the league.
- ii) Ranking the applicants to a college puts them in a best-worst order.

3.5 Lists and permutations

iii) Shuffling a deck involves choosing an order (hopefully “random”) on the set of cards.

iv) Seating a television panel in a row of chairs puts a left-right order on the panelists.

Here’s a typical trickier one:

v) Assigning the 9 starters on a softball team to fielding positions orders them.

This example comes up in many guises: the positions can be replaced by many kinds roles (as officers of a club, as crew on a boat, ...). It’s also the sneakiest since in, say, the softball team example, there’s no unambiguous “first” fielding position. It’s not first base since in standard scoring systems the pitcher is assigned the number 1 and even this assignment is made only by convention. However, the choice of positions amounts to an order because if we shuffled the players amongst the positions we can apply the “Abbot and Costello” method (“Who’s on first?”) to tell that a reordering has occurred.

Lists, ordered subsets and sequences without repetition

If you’re alert, you’ve probably been wondering why I’ve been denoting the underlying set of my lists by the letter B (usually reserved for *subsets*) rather than by the letter A by which we usually denote sets. The answer lies in the applications of list to counting problems that we’ll need soon. Most sets of lists that we’ll need to count do *not* share the exact same underlying set B . However, their B s do have enough in common that it’s still easy to count them. The key example is described by:

LISTS OF LENGTH ℓ FROM A SET A 3.5.7: *A list L of length ℓ chosen from or, more simply, from a finite set A is a list whose underlying set B is a subset of A of order ℓ .*

3.5 Lists and permutations

In other words, choosing L from A involves both choosing a subset B of order ℓ from A and choosing an order on B . In practice, we usually view such choices in an apparently different but equivalent way.

CHOICES FROM A WITHOUT REPETITION 3.5.8: *Choosing a list L of length ℓ from A is the same as selecting a sequence of ℓ individual elements of A without repetitions.*

Successively selecting a sequence of ℓ individual elements of A *without repetitions* provides both the subset B of a list (since we do not allow ourselves to repeat the choice of any element, we are guaranteed to wind up with an ℓ element subset when we're done) and the order on B that makes it a list (namely, the order in which its elements are chosen).

EXAMPLE 3.5.9: Let's write down all the lists of length 3 chosen from the set $A = \{a, b, c, d\}$. The only tricky part is simultaneously avoiding omissions and repetitions. The easiest way to achieve this to try to choose the underlying set first, then fix the order on it. First, comes $\{a, b, c\}$ which gives the 6 lists

$$[a, b, c], [a, c, b], [b, a, c], [b, c, a], [c, a, b], [c, b, a].$$

Then we have successively $\{a, b, d\}$ giving the lists

$$[a, b, d], [a, d, b], [b, a, d], [b, d, a], [d, a, b], [d, b, a] :$$

$\{a, c, d\}$ giving the lists

$$[a, c, d], [a, d, c], [c, a, d], [c, d, a], [d, a, c], [d, c, a] :$$

and, $\{b, c, d\}$ giving the lists

$$[b, c, d], [b, d, c], [c, b, d], [c, d, b], [d, b, c], [d, c, b].$$

There are 24 lists in all.

Once again, I hope that this (very simple) example makes it clear that we don't want to count lists by listing them. We'll usually only be interested in knowing the number of lists—24 in our example. This

3.5 Lists and permutations

we could find by using [SIMPLE BINOMIAL COEFFICIENTS 3.4.25](#) which says that there are 4 subsets of order 3 and then using [PROBLEM 3.5.5](#) which says that there are 6 orders on each 3 element subset. Don't worry soon we'll have a single simple formula for counting lists of length ℓ from a set of m elements.

Before we turn to the problem of counting lists of length ℓ from A , I want to introduce another way of thinking of them that's important in applications.

It's very simple. Lists are also a special kind of sequence! Precisely, every list of length ℓ from A determines sequence of length ℓ in the alphabet A . We can think of the successively choosing the ℓ elements in the list as choosing the ℓ letters in a sequence. What makes the sequences that arise in this way as lists special is that they contain no repeated letters. This is immediate from the requirement that no element in a list may be repeated. Conversely, if a sequence of length ℓ has no repeated letters, then its letters form a subset B of A of order ℓ . The process of passing from the list to the sequence and back just involves deleting and inserting commas. To make the parallel clearer, I'll write the sequences in the same standard font as the lists (and not in typewriter font).

EXAMPLE 3.5.10: The lists $L = [a,b,c,d]$, $L' = [b,d,a,c]$, and $L'' = [c,b,a,d]$ of [EXAMPLE 3.5.2](#) can be viewed as the sequences $L = abcd$, $L' = bdac$, and $L'' = cbad$.

We sum up:

LISTS FROM A AS SEQUENCES 3.5.11: *The lists of length ℓ from a finite set A are exactly the sequences of length ℓ in the alphabet A in which no letter (element of A) is repeated.*

EXAMPLE 3.5.12: Let's write down all the lists of length 3 chosen from the set $A = \{a,b,c,d\}$ as sequences. This time I'll write them down directly in alphabetical order, as is more natural for such word-like sequences: abc , abd , acb , acd , adb , adc , bac , bad , bca , bcd ,

3.5 Lists and permutations

$bda, bdc, cab, cad, cba, cbd, cda, cdb, dab, dac, dba, dbc, dca, dcb$. You can check that we've got the same *set* of 24 lists as in [EXAMPLE 3.5.9](#). But since the lists/sequences are in a different order, we've got a different *list* of 24 lists! Yes, we can have lists of lists just as we can have sets of sets. The good news is that they won't appear very often.

PROBLEM 3.5.13:

- Show that there are 12 lists with 2 elements from the set $A = \{1, 2, 3, 4\}$.
- Write down these lists as bracketed pairs—like $[1, 3]$ —one for each underlying 2-element subset B of A .
- Write down these lists as 2 letter sequences—line 13— in increasing numerical order.

Before we start to count them, let's sum up. A list of length ℓ is a set B of order ℓ together with an order on that set. A list of length ℓ from a set A can be viewed in two ways. First, it's a list of length ℓ for which the underlying set B is an ℓ -element subset of A . In other words, it's an ℓ -element subset B of A plus an order on that subset. Second, a list of length ℓ from A is a sequence of length ℓ in the alphabet A that contains no repeated letters. The upshot is that lists from A are halfway between subsets and sequences and can be viewed either as subsets of A with an order or as sequences in the alphabet A with no repeated letters.

Permutations: counting lists

We now want to introduce a shorthand (or name) for the number of lists of length ℓ from a set A of order m . First, the same arguments that show that the number of ℓ -element subsets of a set A with m elements does not depend on *which* set A we work with apply to lists: the number of lists of length ℓ from a set A with m elements



3.5 Lists and permutations

depends only on m and ℓ . This means that the following definition makes sense.

PERMUTATIONS 3.5.14: *The permutation $P(m, \ell)$ or $mP\ell$ —read “permutation m ℓ ” or “ m permutation ℓ ”—is defined to be the number of subsets B of order ℓ in any or every set A of order m .*

Warning: The term permutation is used both as a shorthand for “numbers of lists” and as a synonym for “ordering”.

PERMUTATIONS OF A 3.5.15: *A choice of a fixed ordering on a set A is also called a permutation of A .*

In fact, word permutation originated in its sense as an ordering and only later was used as a shorthand for the function that will count lists for us. The connection, as we’ll see below is that the *number* of orderings or permutations of a set A of order m is the permutation count $P(m, m)$.

Finding a formula for permutations is easy, in fact, a lot easier than finding a formula for combinations. First note that, for any m , we only need to deal with ℓ between 0 and m since the length ℓ is also the order of a subset of A and that’s at most the order m of A .

As with combinations, we need to understand the trivial case $\ell = 0$. What’s a list of length 0? An empty list, just as a set with order 0 is the empty set. From the “subset plus ordering” point of view of **LIST-FORMAL DEFINITION 3.5.1**, it’s clear that there’s exactly 1 such list: the underlying set must be the empty set \emptyset and then the ordering is irrelevant. We’ll write \emptyset for this list when we need to refer to it. So $P(m, 0) = 1$ for any m .

To get a feel for what a permutation formula should look like, let’s look at the case of a set of order 4—say $A = \{a, b, c, d\}$. I’ll use the more compact “sequences with no repeated letter” representation of lists.

Below are the lists from A ordered first by length, and then alphabetically within each length:



3.5 Lists and permutations

\emptyset

a, b, c, d

$ab, ac, ad, ba, bc, bd, ca, cb, cd, da, db, dc$

$abc, abd, acb, acd, adb, adc, bac, bad, bca, bcd, bda, bdc, cab, cad, cba, cbd, cda, cdb, dab, dac, dba, dbc, dca, dc b$

$abcd, abdc, acbd, acdb, adbc, adcb, bacd, badc, bcad, bcda, bdac, bdca, cabd, cadb, cbad, cbda, cdab, cdba, dabc, dach, dbac, dbca, dcab, dcba$

There are 1, 4, 12, 24 and 24 respectively.

The case $\ell = 1$ is easy: we need to pick one letter from the alphabet A and there are 4 choices. Since there's only 1 letter $P(4, 1) = 4$. Likewise $P(M, 1) = m$ for any m , as there are m choices for the letter.

The case $\ell = 2$ is not much harder. I need to add a second letter to one on my 1 letter sequences. There are only $m - 1$ possibilities for this letter because, since I am making a list and cannot repeat letters, I am not allowed to use the first letter. For example, there are $4 - 1 = 3$ lists of length 2 that start with a : ab, ac and ad . So $P(4, 2) = 4 \cdot 3 = 12$. In general, each of the m lists of length 1 will yield $(m - 1)$ lists of length 2 so $P(m, 2) = m \cdot m - 1$.

PROBLEM 3.5.16:

- i) Check that there are $P(3, 2) = 3 \cdot 2 = 6$ lists of length 2 from the set $A = \{1, 2, 3\}$.
- ii) How many list of length 2 are there from the set $A = \{a, e, i, o, u\}$? List them to check your count.

We see why lists are easier to count than subsets: the fact that lists are ordered means that I can build a list one element at a time in only one way. In the discussion leading up to [COMBINATION FORMULA 3.4.15](#), each subset could arise from several sequences of choices and we had to keep track of these duplications to get the right count.

With this in mind, it's easy to see the pattern and the formula it leads to. To make a list of length 3, I need to add one letter to a



3.5 Lists and permutations

list of length 2. There are now $(m - 2)$ choices for this letter since I cannot reuse either of the 2 letters chosen so far. Thus $P(m, 3) = P(m, 2) \cdot (m - 2) = m \cdot (m - 1) \cdot (m - 2)$. This correctly predicts that $P(4, 3) = 4 \cdot 3 \cdot 2 = 24$.

PROBLEM 3.5.17: How many lists of length 3 are there from the set $A = \{1, 2, 3\}$? List them to check your count.

At each stage, I will find that $P(m, \ell + 1) = P(m, \ell) \cdot (m - \ell)$ —each list of length $\ell + 1$ comes from a unique list of length ℓ by adding one of the $(m - \ell)$ letters not used so far. This unwinds to the formula,

PERMUTATION FORMULA 3.5.18:

$$P(m, \ell) = m \cdot (m - 1) \cdot (m - 2) \cdot \dots \cdot (m - \ell + 1)$$

As with the **COMBINATION FORMULA 3.4.15**, it's much better to view this as a method rather than a formula. The reason is again that it's much easier and much more reliable to learn and use the method than to do the same for the formula.

METHOD FOR COMPUTING PERMUTATIONS 3.5.19:

If $\ell = 0$, there is just 1 list. If not,

Step 1: Start with the factor m , the size of the master set A .

Step 2: Keep lowering the last factor by 1 to get the next factor to multiply by.

Step 3: Stop when you have as many factors as the length ℓ of the lists you want to count.

There's an even simpler way to remember the permutation formula. Notice that it's simply the numerator in the combination formula. So if you learn the latter, you've automatically mastered the former. We'll come back to this point when we discuss factorials below.

PROBLEM 3.5.20:

i) How many lists of each length ℓ from 0 to 5 are there from a set A with 5 elements?



3.5 Lists and permutations

- ii) How many lists of length 6 are there from a set of order 6 and from a set of order 10?
- iii) How many lists of length 10 are there from a set of order 6 and from a set of order 10?

This problem should make it clear, both that it's easy to compute permutations and that it's hopeless to do so by writing down the lists involves. The numbers just get big too fast. When ℓ is small, it's easiest to just multiply out $P(m, \ell)$. But when ℓ is large, you may find it easier to let your calculator do this for you. Most calculators have the permutation function version built in. On the TI-8x series, it's found on the Math panel as nPr. To use it, you enter the order m , select the nPr row, enter the order ℓ and press ENTER. The next challenge shows, however, that even your calculator can't deal with permutations when ℓ starts to have two digits.

PROBLEM 3.5.21:

- i) Use your calculator to check that $P(30, 4) = 681210$ and that $P(30, 8) = 254602237500$
- ii) What happens when you ask your calculator for $P(30, 20)$? For $P(100, 50)$?

CHALLENGE 3.5.22: Find $P(100, 50)$. Hint: It's

30685187562549660372027304595294697392284597216\
84688959447786986982158958772355072000000000000

so don't waste too much time on this, unless you have access to a computer algebra system that handles large integers.

Counting orderings

Recall from [PERMUTATIONS OF A 3.5.15](#), that a permutation originally meant an ordering of a finite set A and that I claimed that the number of such orderings was $P(m, m)$, the number of lists of length m from A . The reason is that such a list is nothing more or less than



3.5 Lists and permutations

an ordering of the elements of A . This is easy to see from either of our viewpoints on lists from A . From the point of view of [LISTS OF LENGTH \$\ell\$ FROM A SET \$A\$](#) [3.5.7](#), we need to first pick an m element subset of B of A . But there's only one: B must equal A itself. Then we need to pick an ordering on the elements of B : since $B = A$, this is the same as an ordering on A . From the point of view of sequences, a list of length m from A is a sequence of length m with no repeated letters. Such a sequence must contain all the elements of A and the elements are then ordered by their position in the sequence. We record this, and introduce another standard notation for $P(m, m)$.

ORDERING AND FACTORIALS 3.5.23: *We define the factorial function $m!$ —read “ m factorial”—by $m! = P(m, m)$. In other words, $m!$ is just a shorthand for the special permutation $P(m, m)$. We interpret both $P(m, m)$ and $m!$, as above, as counting the number of orderings of a set with m elements. The [PERMUTATION FORMULA 3.5.18](#) for $P(m, m)$ gives the formula $m! = m \cdot (m - 1) \cdot (m - 2) \cdot \dots \cdot 2 \cdot 1$, expressing $m!$ as the descending product of the whole numbers from m to 1.*

We will not use factorial notation much in this course. There are two reasons. First, is the principle of least effort. we have no need for factorials, since they're just a shorthand for the permutation $P(m, m)$. Why learn a new formula when it's only a special case of a formula you already have and use? Second, while factorials are very useful in many algebraic manipulations—in calculus they very often appear in asymptotic expansions and in counting problems they often let us condense formulae—they very quickly become too large to work with, even with a calculator.

This makes formulae involving factorials false friends. Students who use factorial based formulae in counting problems often get overflow errors from their calculators, even when the final answer is not so large, because the calculator is unable to handle very large factorials that later “almost” cancel. It's particularly easy to fall into such traps



3.5 Lists and permutations

because factorials are again built into most calculators. On the TI-8x series, it's found on the Math panel as $n!$. To use it, you enter the order m , select the $n!$ row and press ENTER.

But let me emphasize: **Do not use factorials** in counting and probability problems. In this course, using them will never be necessary and doing so can only cause needless problems. Let's start by getting a feel for how fast factorials grow.

PROBLEM 3.5.24:

- i) Multiply out the PERMUTATION FORMULA 3.5.18 to find $6!$ and $10!$. Check your answer using the built-in permutation function in your calculator. Check it again using the built-in factorial function (but don't make a habit of using this function).
- ii) What happens when you ask your calculator for $20!$? for $100!$?

Now, let's look at the condensed factorial formulae for permutations and combinations. These have lots of theoretical applications. But, I stress that you do *not* need to learn them; indeed, you should never use them after you've finished reading this section.

FACTORIAL FORMULAE FOR PERMUTATIONS AND COMBINATIONS 3.5.25:

$$P(m, \ell) = \frac{m!}{(m - \ell)!} \quad \text{and} \quad C(m, \ell) = \frac{m!}{(m - \ell)! \cdot \ell!}$$

They are temptingly simple. instead of a whole list of factors we just have a couple of factorials. We'll see why these are false friends in a moment. Let's first verify these formulae. In doing so, we'll shed some light on the relation between $P(m, \ell)$ and $C(m, \ell)$. What's the difference between $P(m, \ell) = m \cdot (m - 1) \cdot (m - 2) \cdot \dots \cdot (m - \ell + 1)$ and $m! = P(m, m) = m \cdot (m - 1) \cdot (m - 2) \cdot \dots \cdot 2 \cdot 1$. The latter has the extra factors $(m - \ell) \cdot (m - \ell - 1) \cdot (m - \ell - 2) \cdot \dots \cdot 2 \cdot 1$: this is just formula for $(m - \ell)!$ so if we divide m by $(m - \ell)!$, we get $P(m, \ell)$. There is one point we can learn from this. What if $\ell = m$? we want $P(m, m) = \frac{m!}{(m - m)!} = \frac{m!}{0!}$. On the other hand, we want $P(m, m) = m!$. So we have learned that $0! = 1$.



3.5 Lists and permutations

Now let's ask what's the difference between $P(m, \ell)$ and $C(m, \ell)$? They have the same numerator so the difference is exactly given by the denominator of $C(m, \ell)$ which is $1 \cdot 2 \cdot \dots \cdot (\ell - 1) \cdot \ell$. But this is just the product giving $P(\ell, \ell) = \ell!$ written backwards. In other words $C(m, \ell) = \frac{P(m, \ell)}{\ell!}$.

This shows us two things. First, plugging in the factorial formula for $P(m, \ell)$ we check the factorial formula for $C(m, \ell)$: $C(m, \ell) = \frac{\frac{m!}{(m-\ell)!}}{\ell!} = \frac{m!}{(m-\ell)! \cdot \ell!}$. The second insight comes from recalling the definition of $\ell!$ as a shorthand for the number of ways of ordering a set of order ℓ . What's the difference between a subset of order ℓ and a list of length ℓ ? The list is a subset *plus* an ordering of it. There are $\ell!$ choices for such an ordering so every subset gives rise to $\ell!$ lists: in other words, $P(m, \ell) = C(m, \ell) \cdot \ell!$. We thus have a nice check on our earlier formulae.

Since we learned something from checking both formulae, what's so bad about them? The next problem, which is absolutely typical of the counting problems we'll need to deal with later in this chapter gives the answer.

The US Senate has three constitutionally constituted officers chosen from amongst its members: a President Pro Tempore, a Secretary and a Sergeant at Arms. Let's ask how much ways there are of selecting these officers.

How do we proceed? The key idea is that this choice amounts to specifying a list of length 3 (the three officers) from a set with order 100 (the Senate). Why a list and not a subset? The offices let us view the choice as being a choice of 3 senators with an ordering: see [EXAMPLE 3.5.6](#). So the answer is simply $P(100, 3)$.

PROBLEM 3.5.26: How many ways are there to choosing the three officers of the United States Senate?

i) Show that there are 970200 by evaluating $P(100, 3)$ using the [METHOD FOR COMPUTING PERMUTATIONS 3.5.19](#).

3.5 Lists and permutations

ii) Show that there are 970200 by evaluating $P(100,3)$ using the FACTORIAL FORMULAE FOR PERMUTATIONS AND COMBINATIONS 3.5.25.

Take your time with part ii). I'll wait. As long as you like ... OK, I rest my case.

What goes wrong with the factorial approach? From i), we see that the answer is $100 \cdot 99 \cdot 98$. The factorial formula asks us to compute this product as the fraction $\frac{100!}{97!}$ which is

In other words, the factorial formula asks you to first multiply together *all* the numbers from 100 down to 1, then multiply together *all* the numbers from 97 down to 1, and finally take the quotient.

Of course, in the last step, all 97 factors in the denominator cancel with the last 97 factors in the numerator leaving $100 \cdot 99 \cdot 98 = 970200$. But both the numerator and the denominator choke your calculator before it ever gets that far.

To sum up, the factorial formulae for permutations and combinations may look much simpler than the product formulae for but these are really enormously *less* efficient, downright unusable in most practical applications. The solution is easy: don't even try to use them.

PROBLEM 3.5.27:

- If you have 5 rings how many ways are there to put one on each finger of your left hand?
- In a conference with 8 teams, how many different seedings are possible for the conference basketball tournament?
- How many ways can you shuffle a deck of cards?

3.6 Shorthands for basic counts

We're now coming to a part of the course where the questions are simple, but the answers are not always obvious. The questions we'll be tackling are multiple choice questions with answers like "and", "or" or "not", and even simpler True/False questions with answers "Yes" and "No". Very simple questions indeed. But also questions that students very often get wrong! How can anyone get questions with such simple answers wrong? Easy: by guessing. So first a word of advice.

The two rules of guessing: Don't and "D'oh"

Guessing the answer to the simple questions in the rest of this chapter is one of the biggest sources of mistakes in this course. Such mistakes probably rank number two right after those with [SECTION 1.1](#). Yes, the questions are easy and there's no need to guess, but these questions just seem to bring out the Homer Simpson in many students, so that no matter how often guessing them has caused them to get a problem wrong and exclaim "D'oh!", the next guess and the next "D'oh!" are just around the corner. There's really nothing hard here. Just remember the

FIRST RULE OF GUESSING 3.6.1: *Don't!*

The reason this rule is so cut and dried is the

SECOND RULE OF GUESSING 3.6.2: *"D'oh!"*

In this course, Murphy's Law of Guessing says that *all* guesses are wrong. When you can't recall a definition or formula and you're tempted to guess, stop! Breathe deeply, count to three, and think! Go back to the text to look up the facts you need. If you're still confused, ask your instructor or a friend for help. Just remember the



3.6 Shorthands for basic counts

FIRST RULE OF GUESSING 3.6.1 and don't guess. Because even if there are only 2 possible answers the **SECOND RULE OF GUESSING 3.6.2** will apply. Guess and you'll be saying "D'oh!" before you know it. Proceed only when you *know*.

Shorthand Questions

In this short section, we're going to learn *how* to answer the easy, very important counting question:

SHORTHAND QUESTION 3.6.3: *"How many ways are there to make ℓ choices from a set of m possibilities?"*

The question is important because, when we are faced with a more complicated counting problems, and have to apply a "divide-and-conquer" strategy, all the simple pieces into which we'll divide the problem will be versions of this easy question (usually with a different m and ℓ in each). Why are these questions easy? Well, we're going to learn to answer them not with numbers, or even formulas, but just by naming the answer with what I'll call a **shorthand**.

In fact, we've already learned how to answer this question in the preceding sections and we've already learned the names or shorthands for the answers. The only tricky point is hidden in the plurals in the last sentences. The **SHORTHAND QUESTION 3.6.3** has not one, but *three* answers, each with its own shorthand. Put differently, the question comes in three flavors (each with its own shorthand answer). So what we need to understand is how to look at one of these questions and decide which of the three flavors we've got.

The three flavors—which again we've been studying for some time—correspond to three different possibilities for what we're trying to pin down with our choices. These are: sequences, lists and subsets.

The **sequence** flavor is of the **SHORTHAND QUESTION 3.6.3** is, "How many sequences of length ℓ can be chosen using an alphabet (or set)

3.6 Shorthands for basic counts

with m letters ?” and the sequence flavor of the answer is m^ℓ by [SEQUENCE COUNTING FORMULA 3.3.6](#).

The **list** flavor of the [SHORTHAND QUESTION 3.6.3](#) is, “How many lists of length ℓ can be chosen using a set with m elements ?” and the list flavor of the answer is $P(m, \ell)$ by [PERMUTATIONS 3.5.14](#).

This example shows what I mean by a **shorthand**. We *defined* the **permutation** function $P(m, \ell)$ to *be* the number of lists of length ℓ from a set of size m . In other words, $P(m, \ell)$ is nothing more or less than a *name* for the answer to the list flavor of the [SHORTHAND QUESTION 3.6.3](#). Eventually, we will need to use the formula ([PERMUTATION FORMULA 3.5.18](#)) that lets us say what *number* this shorthand stands for—or better yet, the [METHOD FOR COMPUTING PERMUTATIONS 3.5.19](#) that guides us in finding that number. But, this section is a *dégustation*—a French term for a wine tasting in which the labels are hidden. In a *dégustation*, you do not drink wine. Instead, you taste a mouthful of each to recognize its flavor—so you can identify it correctly—and then spit it out. So, in this section, our goal is not to swallow the questions by computing numerical answers, but to sniff and swirl until we can say what flavor each has and can identify the corresponding shorthand.

The **subset** flavor of the [SHORTHAND QUESTION 3.6.3](#) is, “How many subsets of order ℓ can be chosen using a set with m elements ?” and the subset flavor of the answer is $C(m, \ell)$ by [BINOMIAL COEFFICIENTS AND COMBINATIONS 3.4.12](#). Here again the **combination** shorthand is nothing more than a name for the answer to the subset flavor.

The two question method

So far, so good. How then *are* we supposed to know which flavor we’re dealing with when we see a [SHORTHAND QUESTION 3.6.3](#) ? The answer must clearly be to decide whether what we’re looking for are sequences, lists or sets. But, how do we tell *that* from a [SHORTHAND](#)



3.6 Shorthands for basic counts

QUESTION 3.6.3 ? Of course, you can't if you have no other information. That would be like asking a wine taster to identify a wine without tasting it, just by knowing that it's a wine.

However, whenever a **SHORTHAND QUESTION 3.6.3** comes up in practice, it is embedded in a larger context. By context, I just mean that we know something about the set of possibilities and about what exactly the choices are pinning down. Our goal here is to learn how to “taste” this background information and identify the flavor of the choices being made. We'll use a **TWO QUESTION METHOD 3.6.4**—and both of the questions are simple “Yes/No” ones.

How hard can that be? Not very, but, as I said above, these two questions are still tricky. Don't worry if you make mistakes at first. With practice, you'll develop a mental map familiar problems and you'll soon be able to answer the two questions and identify the three basic flavors almost instantly. But beware of becoming too casual. Experience is no substitute for attention. Harry Waugh, for decades considered by many the greatest wine expert in Britain, was once asked if he had ever mistaken a Bordeaux for a Burgundy—roughly equivalent to mistaking beef for chicken. “Not since lunch”, was his reply. So when using the **TWO QUESTION METHOD 3.6.4** to answer a **SHORTHAND QUESTION 3.6.3**, stay alert, think, don't guess! Here's the method.

TWO QUESTION METHOD 3.6.4: *To decide what shorthand gives the answer to “How many ways are there to make ℓ choices from a set of m possibilities?”, ask:*

R? *Are Repetitions allowed?*

O? *Does Order matter??*

First ask R? That is, ask whether you are allowed to select a previously chosen possibility. If this answer is “Yes”, then you are choosing sequences and the shorthand is m^ℓ .



3.6 Shorthands for basic counts

*If it's "No", ask O? That is, ask whether making the same ℓ choices in two different orders gives the same answer twice or two different answers. If this answer is "Yes", then you are choosing **lists** and the shorthand is $P(m, \ell)$, and, if it's "No", then you are choosing **subsets** and the shorthand is $C(m, \ell)$.*

Before we look at some examples, I want to make a few comments. First, the method assumes that you already know what the numbers m and ℓ are. And usually this *will* be obvious: m is the order of the set of possibilities from which each choice is made while ℓ is the number of times you are choosing from these possibilities. If you visualize yourself making a *single* choice and ask what your possibilities are, you'll have m . Then ℓ is the number of times you need to make such a choice. Seems impossible to go wrong. It may be impossible if you remain alert and think, but it's easy if you doze. (Remember Harry Waugh.)

WATCH YOUR ℓ s AND m s 3.6.5: Before applying the **TWO QUESTION METHOD 3.6.4**, first ask "What do I get when I make a choice?". That is, identify the set of possibilities from which each choice will be made, and find its order m . Then, ask "How many choices do I need to make?" to determine ℓ .

Second, there's an obvious gap in this method. When the answer to **R?** is "Yes", we never ask **O?** We just assume that we are choosing sequences. But we know that the order of the letters in a sequence matters (just think of "top" and "pot"—two different words), so in such a case, the answer to "Does **Order** matter?" is always "Yes". What if we had to make choices where the answer to **R?** is "Yes" and to **O?** is "No"? Don't we need to consider this possibility?

What we need to do is not hard to see. Shouldn't we just define an **abomination** to be a way of making ℓ choices from a set of m possibilities with repeated choices allowed and order not mattering (that is, **R?** "Yes" and **O?** "No")? And then, couldn't we define a function

3.6 Shorthands for basic counts

$A(m, \ell)$ that counts the number of such abominations? We could (and maybe we should) but we won't just now.

It even turns out that abominations can be expressed in terms of combinations. The bad news is that the argument needed to show this is a good bit harder than the counting we've done so far. That alone wouldn't be enough to get you off the hook, and, later on, in [THE m&m's PROBLEM](#), we'll explore how to do this as a capstone to our study of counting. The good news is that in the counting problems that come up naturally in probability spaces—in particular, those that we'll be looking at—abominations just don't arise very often.

So I'm simply going to pretty much avoid *any* such problems. In particular, I won't ask you any [SHORTHAND QUESTION 3.6.3](#) that has an abomination as its answer without giving you due warning. This has one benefit that will often come in handy for you.

R IMPLIES O 3.6.6: In any [SHORTHAND QUESTION 3.6.3](#) in this course, if the answer to “Are Repetitions allowed?” is “Yes”, then you may (and should) assume that the answer to “Does Order matter?” is “Yes” too.

Finally, there are a few cases when you can answer a [SHORTHAND QUESTION 3.6.3](#) without using the [TWO QUESTION METHOD 3.6.4](#) because the number of sequences, lists and subsets is the same—that is, all 3 shorthands give the same number of choices. When does it not matter whether repetitions are allowed? When repeated choices are not *possible*, that is, when there are 0 or 1 choices. When can we ignore whether order matters? When there's only *one possible order*, again, when there are 0 or 1 choices. When $\ell = 0$, the common answer is *not* 0 but 1: the empty sequence, the empty list or the empty set are ways of making 0 choices. When $\ell = 1$, the common answer is just the number of possibilities for a single choice—a one letter sequence is a one element list is a one element subset. There are m of these by definition.

3.6 Shorthands for basic counts

EASY SHORTHANDS 3.6.7: No matter how it is described, 1 choice (or 0) from a set of m elements can always be made in m (or 1) ways.

PROBLEM 3.6.8: Check that for any m ,

- i) $m^0 = P(m, 0) = C(m, 0) = 1$.
- ii) $m^1 = P(m, 1) = C(m, 1) = m$.

Now we're ready to get a feel for where the answer to the two questions lurk in common problems by working some examples and problems.

EXAMPLE 3.6.9: The Fordham squash club has 30 members. It is run by a 3 member Executive Committee consisting of a President, Treasurer and Ladderkeeper and a 3 member Scheduling Committee that schedules challenge matches between members. Each week the Ladderkeeper uses the results of challenge matches played during the previous week to update the club's ladder. This ladder ranks the members from best to worst—no ties. On the last weekend of each month, the top 6 players on the ladder play a seeded club match against the top 6 players on the Columbia club's ladder, with each Fordham player matched against the Columbia opponent with the same ladder rank.

Find the right shorthand answer to each of the following questions.

- i) How many different Executive Committees can the club have?
- ii) How many different Scheduling Committees can the club have?
- iii) How many possibilities are there for the top player on the ladder during the first 3 weeks of October?
- iv) How many possible teams are there for the October club match against Columbia?
- v) How many seedings are possible for the October club match against Columbia?
- vi) If we are interested only in which team won in each seeding, many different possible results from the October club match with Columbia?



3.6 Shorthands for basic counts

vii) How many choices are there for the October, November and December teams?

Each part of this problem is a **SHORTHAND QUESTION 3.6.3**. You'll learn a lot from it if you try to tackle it on your own *before* looking at the solutions below.

First see if you can say what the m and ℓ are in each part except the last (which is harder). Then see if you can apply the **TWO QUESTION METHOD 3.6.4** to determine whether we are counting sequences, lists or subsets in each part. Finally, write down the shorthand you think gives the requested count. Only then, check your answers by comparing them with the solutions below.

i) How many different Executive Committees can the club have?

Solution

The set of possibilities for the Executive Committee is the set of 30 club members, so $m = 30$. To select an Executive Committee we need to choose 3 members so $\ell = 3$. Thus, we are making 3 choices from a set of 30 possibilities.

Now it's time for the **TWO QUESTION METHOD 3.6.4**. We first ask: "Are Repetitions allowed?" What would making a repeated choice mean? That we selected the same member more than one time. If we did, we'd have fewer than 3 members on the Executive Committee. So the answer is "No". Notice where this answer was hiding, in the phrase "It is run by a 3 member Executive Committee...".

The best way I know to bring such information into the light and to make it obvious is to try to picture visually both possibilities. What would an Executive Committee look like where I has picked Susan 3 times? What would one look like where I had picked Susan, Jill and Tiffany? Which do I want? The second, where I see 3 members, not the first where I see only 1.

Since the answer to "Are Repetitions allowed?" was "No", we need to ask "Does Order matter?" Learning how to answer this

3.6 Shorthands for basic counts

simple “Yes/No” question reliably is probably one of the most difficult hurdles in the entire course for many students. There’s a huge temptation to guess. Don’t! Remember the [FIRST RULE OF GUESSING 3.6.1](#) or you’ll be remembering the [SECOND RULE OF GUESSING 3.6.2](#).

Once again, there’s really no better approach than to try to visualize the choices. Make a choice like Susan, Jill and Tiffany. See your 3 choices and ask, “Do I see all the information needed to pin down the executive committee?” Here, the answer is “No”. I see *what 3 people* are on the Committee, but I don’t see who is President, Treasurer or Ladderkeeper. How can I see who’s President—or, who holds the other two offices? Hang a mental sign saying “President” around Susan’s neck. And, of course, one saying “Treasurer” around Jill’s and one saying “Ladderkeeper” around Tiffany’s.

Now I *can* see who has what office. I pick the President (Susan) first, then the Treasurer (Jill) and last the Ladderkeeper (Tiffany). Now ask, does what I see change iff I make the *same* choices in a different order. Suppose I pick Tiffany, Jill and Susan, Then what I see is that Tiffany wears “President”, Jill “Treasurer” and Susan “Ladderkeeper”. That’s *not* what I saw before so order *does* matter here.

Notice that, if we hadn’t hung the signs around the Committee’s necks, we wouldn’t have been able to tell the difference. Without the signs of office, I’d just see the same three members (Jill, Susan and Tiffany to use alphabetical order this time) as before. So it was the 3 offices held by the members of this committee that made the order matter. Visualizing the choices allowed me to see how the offices do this.

Now we’re home. Since I found that **R?** was “No” and **O?** was “Yes”, we are counting **lists** and the shorthand for the count is a permutation, here $P(3, 3)$.

3.6 Shorthands for basic counts

Did that seem like a lot of huffing and puffing to get such a simple answer? Well, it was, but only because I was trying to lay out carefully the mental footwork needed to reliably arrive at the right shorthand—and especially, to correctly answer the questions “Are Repetitions allowed?” and “Does Order matter?”

I certainly don’t expect you to write this much. For now, I recommend summarizing the process. Here’s how I’d do this:

“We are choosing 3 members from a club of 30. Are Repetitions allowed? is ‘No’ because I want 3 members on the Executive and Does Order matter? is ‘Yes’ because the 3 offices order the members. So I am choosing lists and there are $P(30, 3)$ possibilities.”

We’ll see a lot of such problems in the course. Pretty soon, you’ll be able to think—or better, see—through a [SHORTHAND QUESTION 3.6.3](#) in your head. This is fine. But if you notice that you are making even a few mistakes in identifying shorthands, then go back to writing down the steps, as many as *you* need, as above.

- ii) How many different Scheduling Committees can the club have?

Solution

Why isn’t this the same as the previous question? We are still choosing 3 members from a club of 30 so $\ell = 3$ and $m = 30$. Once again **R?** is “No” because I need to see 3 different members to have “a 3 member Scheduling Committee”.

The difference becomes clear when we ask “Does Order matter?” Now, there are no signs to hand around the necks if Susan, Jill and Tiffany because the members of the Scheduling Committee do not have offices. So now **O?** is “No” too. We are choosing a subset, and the shorthand is the combination $C(30, 3)$.

- iii) How many possibilities are there for the top player on the ladder during the first 3 weeks of October?

Solution

We are still choosing 3 members from a club of 30 (albeit, not at the same time) so $\ell = 3$ and $m = 30$. But here **R?** is “Yes”.

3.6 Shorthands for basic counts

What if Erica is the Tiger Woods of the Squash Club and has been ranked number 1 for 67 consecutive weeks? Then, my top ranked member all 3 weeks will be Erica. So we are choosing sequences and the shorthand is the power 30^3 .

Because **R IMPLIES O** 3.6.6, we know order matters. But let's check this. What if my top ranked players were Andy the first week, Dan the second and Dan again the third. My sequence would be Andy/Dan/Dan or ADD for short. Can I really distinguish this from the same members in a different order, say Dan/Andy/Dan or DAD for short? Yes, because now Andy was ranked highest the second week, not, as before the first. This will always be true for the same reason that ADD is not the same word as DAD.

iv) How many possible teams are there for the October club match against Columbia?

Solution

We are still choosing members so $m = 30$ but now we want to choose a 6 member team so now $\ell = 6$. The answer to **R?** is “No” because we want to “see” 6 different players on the team. What is the answer to “Does Order matter?” In my mind's eye, can I tell the difference when I reorder the team members? No, because there are no signs to put around their necks. So **O?** is “No”, we are choosing subsets and the shorthand is the combination $C(30, 6)$.

v) How many seedings are possible for the October match against Columbia?

Solution

Again, what is different from the previous question? Only that the answer to **O?** is now “Yes”: the seeding give me signs to put around the necks of the team members, sign that carry, not office like ‘President’. But the rankings 1 through 6. Ranking the same players in a different order changes the seeding. So we are choosing lists, and the shorthand is the permutation $P(30, 6)$.

vi) If we are interested only in which team won in each seeding,



3.6 Shorthands for basic counts

many different possible results from the October club match with Columbia?

Solution

Time to wake up. We're finally choosing something besides members here. Just what though? Here our possibilities are much more limited: either Fordham won (W) or lost (L) each match so $m = 2$. How many such choices do we make? One for each of the 6 seedings, so $\ell = 6$.

OK, so "Are Repetitions allowed?" If Fordham won the #1 match, can it also win at #2? Sure, it could even sweep so **R**? is "Yes".

But this was one case where we could see this answer without much visualization. When, as here, there are *more choices than possibilities* ($6 > 2$), either repeated choices must be allowed or *no* choices are possible. So we choosing sequences and the shorthand is a power. Which power 6^2 or 2^6 ? Once again, don't guess, think! We want m^ℓ or 2^6 . It may seem incredible that anyone could mix this up, but I *know* that if I ask this question to a class I *will* see some 36s mixed in with the 64s.

vii) How many choices are there for the October, November and December teams?

Solution

This is a bit harder. What are the possibilities here? Once again, they are not members. This time they are 9 member teams selected from the club. The number m is therefore the number of possible *teams*. That's exactly the question in part [iv](#)), so the shorthand is $C(30,6)$. Borrowing the evaluation of this shorthand from the next problem, we find that $m = 593775$.

How many choices are we making? One for each of October, November and December so $\ell = 3$. Are Repetitions allowed? Can we field the same team in November as we did in October? No reason why not—this just means the same 6 players are at the top of the ladder. So we are choosing sequences and the shorthand

3.6 Shorthands for basic counts

is the power 593775^3

Whew! But, you may think we haven't really answered the questions in [EXAMPLE 3.6.9](#) by giving the shorthand count for each. Not quite, but almost as the next problem should convince you. Getting from the shorthand count to a number is an easy plug-and-chug process using the formulae of the preceding sections. All the heavy lifting was done in finding the *right* shorthands in [EXAMPLE 3.6.9](#).

PROBLEM 3.6.10: Determine the number of possible choices for each of the parts of [EXAMPLE 3.6.9](#) by calculating the value of the shorthand in the solution.

Partial Solution

I am going to work out just a couple of parts.

iv) Here we want $C(30, 3)$. For this we use the [METHOD FOR COMPUTING PERMUTATIONS 3.5.19](#), getting $30 \cdot 29 \cdot 28 = 24360$.

i) Here we want $C(30, 3)$. For this we use the [METHOD FOR COMPUTING COMBINATIONS 3.4.16](#), getting $\frac{30}{1} \frac{29}{2} \frac{28}{3} = 4060$.

vii) This is a big number but the easiest of all: we want $593775^3 = 209346509902359375$. Your calculator might have to give this in scientific notation as something like $2.093465099 \cdot 10^{17}$.

This illustrates one advantage of shorthand answers. At least 593775^3 reminds us that we were cubing a number—the number 593775 of teams. Recalling that 593775 was $C(30, 6)$ we could also write this answer as $C(30, 6)^3$ which recalls both steps to finding it. The base $C(30, 6)$ is the number of teams and the exponent 3 is the number of months we choose such a team.

Here are some problems for you to practice with. A number of these introduce you to topics we'll work with extensively when we study probability. Each part involves a [SHORTHAND QUESTION 3.6.3](#), but as we move along, the phrasing of the questions will gradually become more informal. As it does, it will become increasingly important to first [WATCH YOUR \$\ell\$ s AND \$m\$ s 3.6.5](#) and only then apply the [TWO QUESTION METHOD 3.6.4](#). I'd like you to state your answer as a power,



3.6 Shorthands for basic counts

permutation or combination, as I did in [EXAMPLE 3.6.9](#). Then you can check, by evaluating and comparing the value you have to one I give in brackets. If you get one of the other answers, you’ve answered one of the questions “Are Repetitions allowed?” or “Does Order matter?” incorrectly. Go back and rethink your answer; remember to try to visualize the choices being made. If you get none of the 3 answers provided, you’ve probably got the wrong m or l . Go back and rethink the sets of possibilities and the number of choices being made from it.

Two standard dice are shown below. It won’t matter to us but the numbers on each pair of opposite faces sum to 7—so the 1 is opposite the 6 and so on. Usually both dice are the same color, but it’ll turn out to make counting problems with dice much easier if we image each die (die is the singular of dice) to have its own color. When




FIGURE 3.6.11: Two standard dice, blue and red

we roll dice, all we’re usually interested in is what the number on the top face of each die is—here the blue die has “come up” 4 and the red 6.

3.6 Shorthands for basic counts

PROBLEM 3.6.12:

- i) How many different ways can 1 die come up?

Solution

There are 6 ways, one for each face of the die. This is so obvious that you may wonder “Why even bother asking?”. Because, asking this kind of “obvious” 1-choice question is the best way to make sure we understand what the number m of possibilities in any problem is. So it’s a good habit to mentally ask yourself this question at the start of any problem, even if I don’t, to make sure you **WATCH YOUR ℓ s AND m s** 3.6.5. Here, for example, when we roll dice, we’re choosing numbers from the set of 6 possibilities on the faces of each die, so $m = 6$.

- ii) How many different ways can 2 dice come up? Hint: Now $l = 2$ and we need to apply the **TWO QUESTION METHOD** 3.6.4. Mentally picture the blue die and the red die. [36 not 30 or 15]
- iii) How many different ways can 5 dice come up? [7776]
- iv) How many different ways can 2 dice come up with different numbers on the red die and blue die? [30]
- v) How many different ways can 5 dice come up with different numbers on each die? [720]
- vi) How many different ways can 7 dice come up with different numbers on each die? [0]

The last question makes a very obvious but still very useful point.

IF $\ell > m$, IT’S SEQUENCES OR NOTHING 3.6.13: If a **SHORTHAND QUESTION** 3.6.3 has ℓ bigger than m , then either *no* such choices are possible or the answer “Are Repetitions allowed?” is “Yes” and we are choosing sequences.

This point also applies in the coin tossing problems that follow. When we toss coins, all we’re usually interested in is whether the coins lands “heads up” or “tails up”—H or T for short. Unlike the dice, which we colored to be able to keep straight, we’re going to try



3.6 Shorthands for basic counts

to work with a single coin. We'll be a bit sloppy and speak of "tossing 2 (or 42) coins" interchangeably with "tossing a single coin 2 (or 42) times".

To keep things straight, we'll number the tosses so we can say what side came up on the third toss, or the fifth or the hundredth (or on the third, fifth, or hundredth coin, if we're thinking of many coins tossed once rather than one tossed many times).

PROBLEM 3.6.14:

- i) How many different ways can we toss a coin 3 times? Hint: The numbers ℓ and m are either 2 and 3 or 3 and 2. Which? [8 not 0]
- ii) How many different ways can we toss a coin 5 times? [32]
- iii) How many different ways can we toss a 20 coins? [1048576]

If you arrived at the right answers in the preceding problem, then you saw that what we are choosing when we toss a coin ℓ times is a sequence of length ℓ in the alphabet $\{H, T\}$ (i.e. the set of $m = 2$ possibilities "heads" or "tails"). There are 2^ℓ such sequences by [SEQUENCE COUNTING FORMULA 3.3.6](#).

In the next problems, we'll use this count to bootstrap our way into thinking in more detail about what happens for a fixed value of ℓ . I'll first work an example with $\ell = 7$ —that is, we want to think about tossing a coin 7 times. Please think for a moment about the first question below before looking at my solution

EXAMPLE 3.6.15:

- i) If we toss a coin 7 times, how many ways can we have 6 heads and 1 tail?
- ii) If we toss a coin 7 times, how many ways can we have 5 heads and 2 tails?
- iii) If we toss a coin 7 times, how many ways can we have 3 heads and 4 tails?
- iv) If we toss a coin 7 times, how many ways can we have 2 heads and 5 tails?



3.6 Shorthands for basic counts

- v) If we toss a coin 7 times, how many ways can we have no tails?
Solution

The reason I asked you to think about this problem a bit first is that illustrates the warning to **WATCH YOUR ℓ s AND m s 3.6.5**. The only trick in this question is deciding what the set of possibilities is, and hence what m should be. We saw above that, when we toss a coin 7 times, we are choosing from the set A^7 7-letter sequence in the alphabet $A = \{H, T\}$ and there are $2^7 = 128$ such choices. So are our possibilities 7-letter sequences and do we want $m = 128$ here? Definitely, *not*!

We do want to *wind up* with 7-letter sequences because, in the end, we are talking about ways of tossing a coin 7 times. But we can't just go picking any old 7-letter sequence, because if we do we'll have no control over how many Hs and Ts it contains. In part i), we want sequences with just 1 T like H H H H T H H and very few 7-letter sequences have this property. What if we picked H T T H T T H or T H T T T H T? We want these in parts iii) and iv), *not* in part i). What we need is some way to pick 7-letter sequences that *only* picks those with a fixed number of Hs and a fixed number of Ts.

The key observation to make is that we can *forget all about the Hs*. If we know what letters in the sequence are Ts—say the Ts are T , then we know the sequence—the *rest* of the letters have to be Hs so the sequence is H H H H T H H. If the Ts are T T T T or T T T T T, what's the sequence?

In other words, what we need to choose is where the Ts go, that is, which of the 7 *positions* in the sequence are to occupied by Ts. So our set of possibilities is the set of 7 possible positions for a T, and hence $m = 7$. How do we ensure that we get a fixed number of Ts like 1 or 2 or 4? By choosing exactly, that *many* positions to hold Ts! In other words, the number of Ts tells us the number ℓ of choices to make.



3.6 Shorthands for basic counts

Now, and only now, are we ready for the **TWO QUESTION METHOD 3.6.4**. Are Repetitions allowed? Try to answer this question before reading on. Think, don't guess. Ask (and visualize) what it would mean to repeat a choice.

Did you say yes? When we were choosing the sequences, the answer was “Yes”: order matters for sequences; think of words like “moo”. But we're *not* choosing sequences now. We're choosing positions. Once you pick the fourth position, you're saying you going to put a T in it—____T____. You may need to lay down some more Ts but you can no longer put them in position four; it's taken. (This illustrates well what I mean when I say visualize). So here, repeated choices are *not allowed*.

OK, so we need to ask “Does Order matter?”. Once again, try to answer this question before reading on. Think, don't guess. Visualize a choice with two Ts: say you first choose position 4, then position 6: you'd see ____T__T__. What would you see if you choose the Ts in the other order? Right, the same thing. So although order *did* matter when we were choosing the sequences of Hs and Ts (remember “top” and “pot”) it doesn't here where we are choosing not sequences but positions.

To summarize, our set of possibilities is the set of $m = 7$ positions in a 7 letter sequence of Hs and Ts. Sequences with exactly ℓ Ts correspond to subsets of the set of 7 positions (because **R?** and **O?** are both “No”) so the shorthand is the combination $C(7, \ell)$. For $\ell = 6; 5; 3; 2$; and 0, there are $C(7, 6)$; $C(7, 5)$; $C(7, 3)$; $C(7, 2)$; and $C(7, 0)$ such choices or 7; 21; 35; 21; and 1.

Is the appearance of the number 21 twice in this list of answers an accident? On the one hand, the fact that $C(7, 5) = C(7, 2)$ follows from **SYMMETRY OF BINOMIAL COEFFICIENTS 3.4.23** because $5 + 2 = 7$. But it's easy to see this directly. Every sequence with exactly 2 Ts has exactly 5 Hs. We can associate to it a sequence

3.6 Shorthands for basic counts

with exactly 5 Ts has exactly 2 Hs by just interchanging the Hs and Ts. So there are the same number of each. Alternatively, any choice of 2 positions also gives a choice of 5 positions—the 5 *not* chosen—and vice-versa.

PROBLEM 3.6.16:

- i) If we toss a coin 10 times, how many ways can we have exactly 9 heads?
- ii) If we toss a coin 10 times, how many ways can we have exactly 5 heads?
- iii) If we toss a coin 10 times, how many ways can we have exactly 2 heads?
- iv) If we toss a coin 10 times, how many ways can we have no heads?

PROBLEM 3.6.17: In many card games, each player is dealt a hand from a standard deck (as in [FIGURE 3.3.11](#)) and then arranges his or her hand to group the cards of the same value (e.g. Go Fish) or suit (e.g. Bridge) or both (e.g. Poker). In such a game:

- i) How many 4-card hands are possible? [270725 not 7311616 or 6497400]
- ii) How many 5-card hands are possible?
- iii) How many 13-card hands are possible?

PROBLEM 3.6.18: How many different answer sheets are possible for a multiple choice test in which:

- i) there are 10 true false questions. [1024 not 100 or 0]
- ii) there are 20 true false questions.
- iii) there are 5 multiple choice questions, each with answers labelled A, B, C, and D and E.
- iv) there are 8 multiple choice questions, each with answers labelled A, B, C, and D and E.
- v) there are 5 multiple choice questions, each with answers labelled A, B, C, D, E, F, G and H.

3.6 Shorthands for basic counts

Now, let's jump ahead a bit and do a few problems that involve counting answers to more than one [SHORTHAND QUESTION 3.6.3](#). The entire next section is devoted to developing a method for answering counting questions that do not fit the basic [SHORTHAND QUESTION 3.6.3](#) model, by a "Divide and Conquer" approach that involves breaking down such questions into pieces each of which is a [SHORTHAND QUESTION 3.6.3](#), then reassembling the counts for these pieces into a final overall count.

Here we'll just warm up by working a few easy counting problems that ask us to count a set S that *doesn't* fit the [SHORTHAND QUESTION 3.6.3](#) model. The problems will be easy because we'll assume that:

- i) It's obvious that the problem involves two or more pieces, and it's obvious what each piece is.
- ii) Each piece counts the answers to a [SHORTHAND QUESTION 3.6.3](#).
- iii) The choice specified in the problem amounts to making an arbitrary choice for each of the pieces.

You'll just have to take [i](#)) on faith for now.

To see how the other assumptions make things easy, let's assume that there are just 2 pieces. What [ii](#)) ensures is that each piece counts a set of sequences, lists or subsets. Let's call the set counted by the first [SHORTHAND QUESTION 3.6.3](#) A and the set counted by the second B . Then, what [iii](#)) tell us is that the set S we want to count in the problem is just the [PRODUCT OF TWO SETS 3.3.8](#) $A \times B$.

But then the [PRODUCT OF TWO SETS 3.3.8](#) tells us that the order of S is just the product of the orders of A and B . In other words, the number of choices in our problem is just the product of answers to the two [SIMPLE QUESTIONS](#). In fact, we've already worked this kind of counting problems in [SECTION 3.3](#)—see [EXAMPLE 3.3.10](#) and [PROBLEM 3.3.12](#).

That last problem tells you what to do if there are 3 pieces in the problem, instead of just 2: the answer is just the product of the

3.6 Shorthands for basic counts

answers to all 3 **SIMPLE QUESTIONS**. In fact, it really doesn't matter how many pieces there are: you get the count you're after by just multiplying together the count for each piece by the **GENERAL PRODUCT SET COUNTING RULE 3.3.14**. Here's a model 2 piece example.

EXAMPLE 3.6.19: To have a valid MegaMillions™ ticket, you must fill out a form like that shown in **FIGURE 3.6.20**. How many different

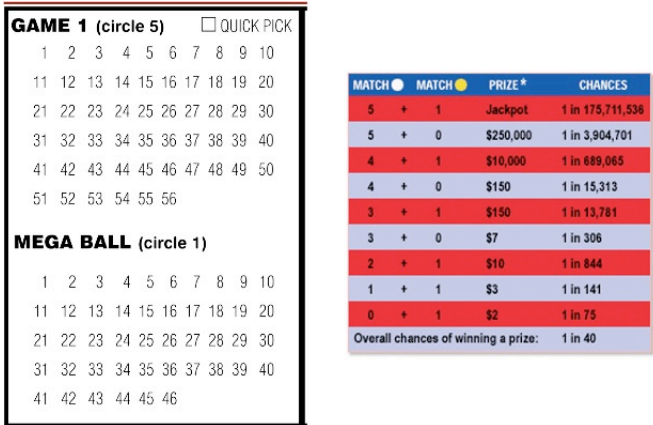


FIGURE 3.6.20: A MegaMillions™ form and odds

valid tickets are there?

As I promised, it's obvious what the pieces are. We need to chose:

- i) the 5 circled numbers between 1 and 56 on the upper part of the form; and,
- ii) the 1 circled number between 1 and 46 on the lower part of the form.

The upper part involves making 5 choices from a set of 56. Repetitions are not allowed ("circle 5") and order does not matter (since the machine that reads our slip will only know what 5 numbers we picked, not the order we picked them in). So the first count is the combination $C(56, 5) = 3819816$. The lower part involves the **EASY SHORTHANDS 3.6.7** of making 1 choice from 46: 46. Both

3.6 Shorthands for basic counts

choices are arbitrary so the number of valid tickets is $3819816 \times 46 = 175711536$.

Looking ahead, not just to the next section, but to the next chapter, you can see an example of why we want to be able to find such counts. We just calculated the chances that a random MegaMillions™ ticket will have the 6 winning numbers: it's 1 in the number of valid tickets as shown in the chart on the right.

We're not yet ready to calculate any of the other chances on the right, but it's not too early to think bit about them. What's my chance of getting the 5 winning numbers on the upper part of my ticket? There are 3819816 choices for these numbers, but the form at the right says I only have a 1 in 3904701 chance of winning the prize for matching these numbers and not the lower (yellow megaball) number. Why don't these match up? Hint: The form is wrong. You really only have a 1 in 3904700.8 chance of winning this prize and $3904700.8 = \frac{175711536}{45}$.

PROBLEM 3.6.21: Suppose that to play GigaMillions you have circle 6 numbers from 1 to 56 and that to win you have to match 6 numbers drawn from 1 to 56. Are your chances of winning GigaMillions better or worse than of winning MegaMillions™?

Now here are a few more easy multipiece problems for you to try.

PROBLEM 3.6.22: The Fordham Math Club has decided to select a committee to organize its Christmas Party.

- i) If there are 25 members in the club and 6 members on the committee, how many different committees can be chosen?
- ii) If there are 16 members in the club and 4 members on the committee, how many different committees can be chosen?
- iii) If there are 9 members in the club and 2 members on the committee, how many different committees can be chosen?
- iv) Suppose there are 16 women and 9 men in the club. How many of the party committees will consist of 4 women and 2 men? Hint: Look at the two preceding parts.



3.6 Shorthands for basic counts

v) Suppose there are 13 women and 12 men in the club. How many of the party committees will consist of 3 men and 3 women?

In the next problem, you'll need to multiply 3—and even 4—numbers together to get the answers to some of the parts.

PROBLEM 3.6.23: An Ontario license plate consists of 3 letters from the Latin alphabet followed by a 3 digit number (each digit can be 0-9).

- i) How many valid Ontario license plates are there?
- ii) How many valid Ontario license plates end in a 0? Hint: Instead of choosing all 3 digits at once, choose the first two and the last separately.
- iii) How many valid Ontario license plates start with a vowel (a, e, i, o, or u)? Hint: Choose the first letter separately.
- iv) How many valid Ontario license plates start with a vowel and end with a 0?

Poker is a classic source of counting problems—with real-world applications. Here we'll do a few easy counts to begin to familiarize ourselves with the game's terminology.

PROBLEM 3.6.24: In [PROBLEM 3.6.17](#), we saw that there were $C(52, 5) = 2598960$ ways to choose a 5 card Poker hand (that is R? and O? are both “No”).

- i) How many poker hands consist entirely of hearts? Hint: What's the only difference between this and choosing an arbitrary poker hand?
- ii) How many poker hands consist entirely of cards from a single suit? (Such a hands is called a **flush** (with a few exceptions—see [POKER RANKINGS 3.8.58](#).)
- iii) How many poker hands contain 4 Aces? Hint: How many non-Aces are there in a deck?
- iv) How many poker hands contain 4 cards of the same value (or **four-of-a-kind**)?



3.6 Shorthands for basic counts

Strictly speaking, I haven't explained how to do the last two parts of [PROBLEM 3.6.24](#). But I'll bet you didn't have any difficulty with them. You'll have the same experience in many counting problems. You may not have a formal method for making the requested count, but if you just follow your nose and ask what choices do I need to make and how many ways are there to make each, the answer is easy to find. The trick is to find the right blend of arrogance ("I can do this in my sleep") and alertness: even in easy counting problems, you just have to doze for a moment to make a mistake like mixing up m and ℓ or getting the answer to one of the questions **R?** or **O?** wrong. Having methods and rules helps in counting problems but is often not necessary.

We'll write down the rule you used to answer [iii](#))—since there are 52 cards in a deck and 4 Aces, there are $52 - 4 = 48$ non-Aces—in [COMPLEMENT FORMULA FOR ORDERS 3.7.24](#).

Part [iv](#)) is a little different. There are 13 values for the 4-of-a-kind and then 48 possibilities for the fifth card or **kicker** so multiplying we get 624 possible hands. We didn't even need any shorthands here because both pieces involved the easy case of just a single choice and multiplying the counts might seem like a simple application of [PRODUCT OF TWO SETS 3.3.8](#). But what we are counting here is not a product set because each of the 13 values determines a *different* set of 48 kickers: in a product $A \times B$, the sets B of bs is the same for every a . No worries. All that we really need to justify multiplying the two counts is that the second *count* (the 48) is the same regardless of which choice (the value) is made in the first. We'll see this in [MULTIPLICATION PRINCIPLE 3.7.1](#).

PROBLEM 3.6.25: Every day after school, you order a 3 scoop cone in an ice cream store offers 31 flavors of ice cream.

i) How many days you can go without ordering the same cone twice if you always choose 3 different flavors and you consider two cones "the same" if:



3.7 And, Or, and Not: the three hardest words

- a. they have the same 3 flavors in the same top-middle-bottom order?
 - b. they have the same 3 flavors in any order?
- ii) How many days you can go without ordering the same cone twice if allows yourself to have multiple scoops of the same flavor and you consider two cones “the same” if:
- a. they have the same 3 flavors in the same top-middle-bottom order?
 - b. they have the same 3 flavors in any order? Hint: This part is a bit trickier because none of the three standard shorthands apply. We are choosing 3 flavors from 31 but **R?** is “Yes” and **O?** is “No”: this is a **abomination**. But you answer the question by applying what we’ll call the **DIVIDE AND CONQUER COUNTING STRATEGY 3.8.1**. Every such cone either has 3 scoops of a single flavor or 2 scoops of a first flavor and 1 of a second, or 3 scoops of different flavors. Since these possibilities are disjoint, we’ll see that the number we’re after is just the sum of the counts for the 3 types. Use this to show that the answer is 5456.

3.7 And, Or, and Not: the three hardest words

In the rest of this section, we’re going to learn about the operations that can be used to apply the divide and conquer strategy to simplify counting sets. These operations are deceptively easy. They seem easy because they are described by very common, non-technical English words like “and”, “or” and “not”. What makes them tricky is that these words will have a very precise mathematical sense. An extra complication, especially with “and”, and “or” is that there are two *different* precise meanings, and you have to pay careful attention to the context of each problem to know which is the “right ” precise meaning in that problem. Finally, you need to train yourself to be



3.7 And, Or, and Not: the three hardest words

constantly *alert* not to confuse these precise mathematical meanings with the informal, everyday uses of these words.

andthen: extensive and

Let's start by looking at **and**. This may be the trickiest word of the trio, partly because it seems like the easiest. Like the other simple words, it is a conjunction—from the Latin to “join together”—and we use it to join together two sets of choices. The real problem is that “and” has come to stand for two important, and very different, *ways* of joining together two sets.

We've already encountered the first meaning. Let's think about what's we need to do to “Pick a card! Any card.” from a standard deck D as introduced in [EXAMPLE 3.3.10](#). That's easy, we have to give the suit *and* the value of the card. We can pick any of the 4 suits from the set S (without worrying about what value we'll later select) and then pick any of the 13 values from the set V (without worrying about what suit we just selected).

Here we're using “and” to connect two choices that are made in sequence. The second choice is an *extension* of the first and there is *no restriction* on either choice. We'll call this the **extensive** or **andthen** meaning of “and”. Although I'll use **andthen** in informal sentences like this one, we'll always “print” it in the typewriter font shown to emphasize its mathematical meaning. This will be the more common meaning of **and** in the rest of this chapter, but not in the next.

[EXAMPLE 3.3.10](#) illustrates perfectly that what we are doing in making an **andthen** sequence of choices is choosing an element of a product of sets. The deck D is the product set $S \times V$ and we specify a card like the four of hearts or the king of spades by specifying its suit (**♥** or **♠**) and its value (4 or K) to get the ordered pairs (**♥, 4**) or (**♠, K**). By the [PRODUCT SET COUNTING FORMULA 3.3.9](#), we can make this choice in $4 \times 13 = 52$ ways.



3.7 And, Or, and Not: the three hardest words

More generally, making *and then* choices involves choosing in succession an element for each of several sets. Each choice extends the previous ones and is unrestricted by them so our choices correspond to elements of the product of individual sets. Again, the [GENERAL PRODUCT SET COUNTING RULE 3.3.14](#) says that the number of ways of making such a choice is just the product of the orders of the individual sets.

It's common, in counting such *and then* choices, not to mention the product set explicitly and to re-express the [GENERAL PRODUCT SET COUNTING RULE 3.3.14](#) more simply as:

MULTIPLICATION PRINCIPLE 3.7.1: *The number of ways of making a choice that involves an *and then* sequence of component choices is the product of the number of ways of choosing each component.*

For example, there are 4 suits and 13 values so there are $4 \cdot 13 = 52$ ways to choose a card from a standard deck.

A number of the counting formulae that we already know can be viewed as special cases of this principle. For example, the [SEQUENCE COUNTING FORMULA 3.3.6](#) counts the set A^ℓ of sequences of length ℓ chosen from an alphabet A with m elements. Choosing such a sequence amounts to a sequence of ℓ choices of a single letter, for which there are m possibilities each time. We get a product with ℓ factors, each equal to m , and that's just m^ℓ . In this example, we're taking repeatedly taking the product of a set with itself.

We've been requiring the set of possible second choices to be the same regardless of what first choice we made. I did not make this explicit in the [MULTIPLICATION PRINCIPLE 3.7.1](#) because, by being a bit more flexible, we can dispense with this restriction. If we are just trying to count choices, we need only ask that the *number* of possible second choices always be the same. If for each of a first choices there are b second choices, then the total number of pairs of choices will be a sum of a terms, one for each first choice, all of which equal the

3.7 And, Or, and Not: the three hardest words

number b of second choices:

$$\underbrace{b + b + \cdots + b}_{a \text{ terms}} = a \cdot b$$

The **PERMUTATION FORMULA 3.5.18** is an example of a formula that fits this more flexible model. For example, in counting lists of length 2 from a set A with m elements, we have m choices for the first element a . For each of these choices, we have a different *set* of possible second choices: we can choose any element of A except a . But that always leaves us $m - 1$ second choices so we recover $P(m, 2) = m \cdot (m - 1)$.

PROBLEM 3.7.2: Use a similar argument to show that $P(m, 3) = m \cdot (m - 1) \cdot (m - 2)$.

The **andthen** meaning of “and” causes few problems. It’s already familiar in many context and, in counting problems, we tend to naturally perform the multiplications is tells us to.

andalso: restrictive and

To introduce the second meaning of “and”, let’s consider a simple example. We’ll let A be the set of 13 states that were part of the original, pre-USA colonies, and let B be the set of 11 states that were members of the Confederacy. For those who cut Social Studies a bit too often, the states in A are New Hampshire, Massachusetts, Rhode Island, Connecticut, New York, New Jersey, Pennsylvania, Delaware, Maryland, Virginia, North Carolina, South Carolina and Georgia. Those in B are Virginia, North Carolina, South Carolina, Georgia, Florida, Tennessee, Alabama, Mississippi, Louisiana, Arkansas and Texas. Answer the two counting questions below.

- How many ways are there to choose an original colony and a Confederate member?
- How many states were original colonies and Confederate members?



3.7 And, Or, and Not: the three hardest words

First off, are these the same question. Definitely *not*!

The first question is asking us to pick *two* states in sequence, first one in A and then a second in B . We are asked to pick a original colony and then a Confederate state, so what we are selecting is an element of the product set $A \times B$. Typical possibilities for the pair of choices are (Rhode Island, North Carolina), (Virginia, Virginia), and (New York, Texas). By [SECTION 3.3](#) or the [MULTIPLICATION PRINCIPLE 3.7.1](#), we can make this choice in $13 \times 11 = 143$ ways.

The second question is only asking us to pick *one* state, but it gives us *two restrictions* on that choice: the state we pick must have been *both* an original colony *and also* a member of the Confederacy. We'll call this use of "and" the *restrictive "and"* or *andalso*. These two restrictions specify the 4 states Virginia, North Carolina, South Carolina and Georgia. In terms of the two sets of states, we're asking what elements (states) belong to *both* of the sets A and B . This restrictive andalso meaning of and has not yet come up mathematically in this course. As with andthen, andalso will always be "printed" in the typewriter font shown to emphasize its mathematical meaning. Before going further, let's look more closely this kind of choice.

INTERSECTION OF SETS 3.7.3: *We write $A \cap B$ and read "A intersect B" or "A and B" for the intersection of two sets A and B, the collection of elements the two sets have in common. To specify $A \cap B$ as a set, we give its membership test: the elements of the intersection $A \cap B$ are exactly those objects that are both in A and also in B. That is, to pass the test $x \in A \cap B$, x must pass the two admission tests $x \in A$ and $x \in B$.*

It can happen that there is no element that belongs to both A and B . We encountered examples of this in [SECTION 3.2](#) and we even named this special, but important case, saying that A and B are [DISJOINT SETS 3.2.5](#).



3.7 And, Or, and Not: the three hardest words

DISJOINT MEANS EMPTY INTERSECTION 3.7.4: *To say that two sets A and B are **disjoint**—no x is a member of both A and B —is the same as saying that $A \cap B = \emptyset$.*

It can also happen that every element of B is also an element of A and hence that $A \cap B = B$. We have already dealt extensively with this possibility: it happens exactly when B is a subset of A . Let's just record this for future reference as it can be useful in both directions.

SUBSETS AND INTERSECTIONS 3.7.5: *The set B is a subset of A — $B \subset A$ —if and only if $A \cap B = B$.*

Let's try a few problems to familiarize ourselves with intersections.

PROBLEM 3.7.6: Let's let D^2 be the set of 36 possible possible ways of rolling a blue and a red die that, in [PROBLEM 3.2.8](#), we viewed as ordered pairs of numbers from 1 to 6. In that problem we defined a number of subsets of D^2 :

ES is the set of pairs with an "Even Sum".

$S7$ is the set of pairs with "Sum 7".

$S4$ is the set of pairs with "Sum 4".

BE is the set of pairs which are "Both Even".

BO is the set of pairs which are "Both Odd".

OE is the set of pairs with "one Odd, one Even".

FO is the set of pairs with "First number Odd".

List the elements of the following intersections. I'll work the first few parts to get you started.

i) $S4 \cap BE$

Solution

The only even number less than 4 on the dice is 2, so to get a total of 4 with both even we must roll a (2,2): so $S4 \cap BE = \{(2,2)\}$.

ii) $S4 \cap BO$

Solution

3.7 And, Or, and Not: the three hardest words

The odd numbers less than 4 on the dice are 1 and 3, so to get a total of 4 with both odd we must roll either a (1, 3) or a (3, 1): so $S4 \cap BO = \{(1, 3), (3, 1)\}$.

iii) $S4 \cap FO$

Solution

If the total is 4 and the first die is odd, then so is the second. So even though $BO \neq FO$, $S4 \cap BO = S4 \cap FO = \{(1, 3), (3, 1)\}$

iv) $S7 \cap BE$

Solution

If both die are even, so is their total so the sum can never equal the odd total. This means that $S7 \cap BO = \emptyset$, or that $S7$ and BO are *disjoint*

v) $S7 \cap BO$

vi) $S7 \cap OE$

vii) $S7 \cap FO$

viii) $FO \cap OE$

ix) $ES \cap BE$

x) $ES \cap BO$

xi) $ES \cap FO$

One point that should already be clear from looking at the answers to this first problem is that there's no simple arithmetical formula for finding the number $\#((A \cap B))$ of elements in the intersection of A and B from the orders of A and B individually. We'll often be interested in this order, but to find it, we usually need to first describe directly its elements (that is, give its admission test) without reference to A and B .

In [PROBLEM 3.7.6](#), all the subsets whose intersections we wanted to find were themselves *subsets* of a single “master” or “universal” set (the set D^2 here). In the examples we'll be interested in, particularly when we study probability, we'll almost always be given such a master set.

3.7 And, Or, and Not: the three hardest words

UNIVERSAL SET 3.7.7: *When, in a problem or example, we agree to consider only sets that are subsets of a fixed set S , we call S the universal set or universe S . Many books use the letter U for universal sets. We'll use S because our main applications involve sample spaces in probability problems where it's standard to denote by S .*

Here's some more problems that confirm both these observations.

PROBLEM 3.7.8: Consider the universal set S of ways of tossing a coin 5 times. In [PROBLEM 3.6.14](#), we saw that this set could be thought of as sequences of length 5 in the letters H and T and that it has 32 elements. Consider the following subsets of S :

4H the set of sequences with *exactly* 4 Hs.

3+H the set of sequences with *at least* 3 Hs.

2+T the set of sequences with *at least* 2 Ts.

EVEN-T the set of sequences with an even number of Ts.

ODD-T the set of sequences with an odd number of Ts.

LAST-TT the set of sequences with last two letters TT.

FIRST-H the set of sequences with first letter an H.

Describe *as directly as possible*, but *without* listing the elements, the intersections of each pair of sets in this list.

This is a bit trickier than [PROBLEM 3.7.6](#) because of the injunction against listing elements. Why *not* list them? Once again, because soon we'll want to be able to work with much better sets whose intersections have too many elements to list. The *only* way to deal with such intersections is to find a way to state their admissions tests that does not require listing elements.

Partial Solution

Let's find the intersections with 4H and EVEN-T.

First 4H. All but one of these are at the easy extremes where one set in the pair is a subset of the other, or the two sets are disjoint. Every sequence with exactly 4 Hs has "at least 3 Hs" so $4H \subset 3+H$ and by [SUBSETS AND INTERSECTIONS 3.7.5](#), $4H \cap 3+H = 4H$. If a sequence has exactly 4 Hs it has exactly 1 T. This tells

3.7 And, Or, and Not: the three hardest words

us what the next 4 intersections are. First, no element of 4H is in 2+T, EVEN-T or LAST-TT and every element of 4H is in ODD-T. In other words, 4H is disjoint from both 2+T, EVEN-T and LAST-TT and 4H is a subset of ODD-T. So, $4H \cap 2+T = 4H \cap \text{EVEN-T} = 4H \cap \text{LAST-TT} = \emptyset$ and $4H \cap \text{ODD-T} = 4H$.

The one interesting case is $4H \cap \text{FIRST-H}$. Some elements of FIRST-H are in 4H but others aren't: for example, HTHHH is but HTTTT is not. Likewise some elements of 4H are in FIRST-H but others aren't: for example, HHTHH is but THHHH is not. So the intersection is not equal to either of the two sets, not is it the empty set. There are lots of ways to say what the intersection is such as “all the sequences with 4 Hs *except* THHHH” or “an H followed by a string of length 4 with exactly 3 Hs”. But these are just roundabout ways of saying that we're talking about $4H \cap \text{FIRST-H}$. So here there is not better answer than just to describe the intersection as the intersection.

This answer may seem a bit silly but it conceals an important point. Intersection is a useful concept because many sets arise naturally in applications as the intersection of other, larger sets. When they do, there's often no easy way to describe them—that is, give their admission test—except *as* an intersection. So we need to be able to recognize both when an intersection has an easier description and when it does not. You'll see a number of other examples of this type in working the rest of the problem.

So let's go on to the intersections with EVEN-T. Here, there are fewer easy cases. The set EVEN-T is disjoint from 4H and ODD-T so these intersections are empty. But EVEN-T neither contains nor is a subset of any of the other sets listed. Now, however, we can sometimes describe the intersection more directly than as the intersections. For example, a sequence with an even number of Ts and at least 3 Hs has either 3 Hs and 2 Ts or 5 Hs and 0 Ts. So we can describe $\text{EVEN-T} \cap 3+H$ as “sequences with exactly

3.7 And, Or, and Not: the three hardest words

3 or 5 Hs” or as “sequences with exactly 2 or 0 Ts”. Likewise we can describe $\text{EVEN-T} \cap 2+\text{T}$ as “sequences with exactly 2 or 4 Ts” or as “sequences with exactly 3 or 1 Hs”. But the intersections $\text{EVEN-T} \cap \text{LAST-TT}$ and $\text{EVEN-T} \cap \text{FIRST-H}$ are best described as intersections.

Just as we can make sequences of more than 2 **and** then choices, we can impose more than 2 **and** also restrictions—though the latter is less common. For example, let’s let C be the set of states with two-word names (like “New Hampshire” or “South Carolina”). Then we can ask what states were in the original thirteen colonies, were members of the confederacy and have two word names. That’s usually how we’d ask, but for emphasis, we could restate this as, what states were in the original thirteen colonies, **and** also were members of the confederacy, **and** also have two word names. This is the triple **intersection** of the sets A , B and C which we write $A \cap B \cap C$.

PROBLEM 3.7.9: List the elements of $A \cap B \cap C$.

OK, so how do we distinguish between the two “ands”. It’s not really so hard. If what is being described involves making more than 1 choice, you are dealing with **and** then and are (usually) describing the product of the sets giving each choice. If it involves making only 1 choice, but you are given more than 1 conditions or *restriction* that the choice must satisfy, then you are dealing with **and** also and you are being asked for intersection of the sets described by each condition. In practice, if you simply remain alert to the fact that “and” has two meanings, you’ll have no trouble keeping the two straight. It’s when you don’t ask which “and” you’re faced with that mistakes occur.

either or both and or else

To introduce the mathematical meaning of “or”, let’s ask another question about the states: “What states were original colonies *or*

3.7 And, Or, and Not: the three hardest words

Confederate members?” First, like an *and* also question, the answer is a single subset of the 50 States, even though the question involves the two related sets of original colonies and Confederate members.

To nail down the relation between the sets in the question and the set that’s the answer, let’s rephrase the question: “What states were *either* original colonies *or* Confederate members?” This makes clear the difference between an *either* or both and an *and* also question. “And then” asks what elements belong to *both* sets. “Either or” asks what elements belong to *one of the two*. Unfortunately, the condition “one of the two” can be interpreted in two different ways. Does it mean “exactly one of the two sets” or does it mean “at least one of the two sets”? These are different because elements in *both* sets—that is, elements of the intersection—are included by the latter but are excluded by the former: an element of both of 2 sets is in “at least one” but not in “exactly one”.

Mathematicians use “or” to mean “at least one”. Period. This is a convention, somewhat arbitrary but universally followed, like the order-of-operations convention [POGEMDAS 1.1.5](#). Well, not quite universally followed. Many students persistently insist on using “or” to mean “exactly one”: if you do, you can expect to get a lot of answers wrong. So please try to get the difference straight right now and learn the (conventionally) correct meaning: or includes both!

Since the *either* or both meaning of or is new, we also need a way to denote it. Once again, the typewriter font is used to emphasize its mathematical meaning even in ordinary text.

***either* or both UNION OF SETS 3.7.10:** We write $A \cup B$ and read “A union B” or “A or B” for the *union* of two sets A and B, the collection of elements that are *either* in A, or in B or in both. To specify $A \cup B$ as a set, we give its membership test: to pass the test $x \in A \cup B$, x must pass at least one of the admission tests $x \in A$ and $x \in B$. If x happens to pass both—that is $x \in A \cap B$, then x is in $A \cup B$. More informally, saying x is in $A \cup B$ means that x is *either* in A or in B or in both.



3.7 And, Or, and Not: the three hardest words

So the answer to the question, “What states were original colonies *or* Confederate members?” is: New Hampshire, Massachusetts, Rhode Island, Connecticut, New York, New Jersey, Pennsylvania, Delaware, Maryland, Virginia, North Carolina, South Carolina, Georgia, Florida, Tennessee, Alabama, Mississippi, Louisiana, Arkansas and Texas. In particular, the four states—Virginia, North Carolina, South Carolina, Georgia—in both the original thirteen and the Confederacy are elements of the union of these sets.

PROBLEM 3.7.11: Mathematicians have a name for the “exactly one” kind of or: it’s called an *exclusive or* or *xor*. We won’t use this kind of or in this course, but just to let you get it out of your system, find the exclusive or of the thirteen original colonies and the members of the Confederacy.

PROBLEM 3.7.12: This problem is a sequel to [PROBLEM 3.7.6](#) and uses the same sets of die throws defined there.

List the elements of the following unions. I’ll work the first few parts to get you started. You notice I list throws in “increasing order” (read left to right) because it makes it easier to avoid omissions.

i) $S_4 \cup BE$

Solution

Since

$$S_4 = \{(1, 3), (2, 2), (3, 1)\}$$

and

$$BE = \{(2, 2), (2, 4), (2, 6), (4, 2), (4, 4), (4, 6), (6, 2), (6, 4), (6, 6)\},$$

we find that

$$S_4 \cup BE = \{(1, 3), (2, 2), (2, 4), (2, 6), (3, 1), (4, 2), (4, 4), (4, 6), (6, 2), (6, 4), (6, 6)\}.$$

Once I had listed S_4 and BE this was pretty easy. I just merged the two lists. All I had to do was to avoid duplicating $(2, 2)$ which was in both sets (and avoid committing the “exactly one” mistake and leaving $(2, 2)$ out).

3.7 And, Or, and Not: the three hardest words

ii) $S4 \cup S7$

Solution

Using

$$S7 = \{(1, 6), (2, 5), (3, 4), (4, 3), (5, 2), (6, 1)\}$$

and the listing of $S4$ above gives

$$S4 \cup S7 = \{(1, 3), (1, 6), (2, 2), (2, 5), (3, 1), (3, 4), (4, 3), (5, 2), (6, 1)\}.$$

Here the two sets were disjoint so I didn't need to deal with any elements in both.

iii) $S4 \cup FO$

iv) $S7 \cup BE$

v) $BO \cup FO$

The answers to this first problem make a couple of points worth noting. First, if $A \subset B$ (or vice-versa), then $A \cup B = B$ (as in the part v), where $BO \cup FO = FO$).

SUBSETS AND UNIONS 3.7.13: *The set A is a subset of B if and only if $A \cup B = B$.*

The only other common case when there's a simpler way to describe a union of sets than as this union is when the union equals the **UNIVERSAL SET 3.7.7** of a problem. The next problem gives examples, some involving unions of more than 2 sets. We can take the union of any number of sets. An object that belongs to at least one of the sets belongs to their union (and that includes elements belonging to *more* than one of the sets).

PROBLEM 3.7.14: This problem is a sequel to **PROBLEM 3.7.12** and uses the same sets of dice throws defined there.

List the elements of the following unions.

i) $ES \cup OE$

ii) $BE \cup BO \cup OE$

Second, there's no simple arithmetical formula for finding the number $\#(A \cup B)$ of elements in the union of A and B from the orders

3.7 And, Or, and Not: the three hardest words

of A and B individually. We already saw the same thing for the order of the intersection $\#(A \cap B)$. There is no general formula relating just 3 of these orders. However, there is an extremely useful formula relating the orders of all *four* of these sets.

AND-OR FORMULA FOR ORDERS 3.7.15:

$$\#(A \cap B) + \#(A \cup B) = \#A + \#B$$

The easiest way to see this is to observe that counting a set is the same as summing up the number 1 once for each element: that is, we mentally write $\#A = \sum_{x \in A} 1_x$. The subscript x is just a label to remind us which element x we were counting when we added the corresponding 1 to the sum.

Thus, each of the four orders in the formula is a sum of terms 1_x for certain x , namely the elements of the corresponding subset. I claim that for every x , the number of terms 1_x on the left and right side are the same, and hence that the two sides are equal. The diagram below makes this easy to see.




FIGURE 3.7.16: Picturing the And-Or Formula

The set A is the upper rectangle, and the set B of the right one. The set $A \cap B$ is the overlap at the top right and the set $A \cup B$ consists everything in either rectangle.

Elements x in the top left quarter of the diagram lie in A and in $A \cup B$ but not in B or in $A \cap B$ so they contribute a *single* 1_x to each side. Elements x in the bottom left lie in B and in $A \cup B$ but not in A or in $A \cap B$ so they also contribute a *single* 1_x to each side. Elements



3.7 And, Or, and Not: the three hardest words

in the top right lie in both A and B and hence in both $A \cup B$ and $A \cap B$ so they contribute *two* 1_x 's to each side. Elements outside both rectangles are in none of the four sets and contribute *no* 1_x to either side. In all cases, the contribution of x to each side is the same as claimed.

We usually use the formula to solve for one of the 4 orders in it, rewriting it to express the unknown order in terms of the 3 known ones. The most common variant is $\#(A \cup B) = \#A + \#B - \#(A \cap B)$ which we use to find orders of unions.

A few examples will give you a better feel for this argument.

PROBLEM 3.7.17: This problem is a sequel to [PROBLEM 3.7.12](#) and uses the same sets of dice throws defined there.

Verify the [AND-OR FORMULA FOR ORDERS 3.7.15](#) for each of the pairs of sets below by counting elements to compute both sides.

i) $S4$ and BE

Solution

From [PROBLEM 3.7.6i](#)), $\#(S4 \cap BE) = 1$. From [PROBLEM 3.7.12i](#)), $\#(S4 \cup BE) = 11$. So $\#(S4 \cap BE) + \#(S4 \cup BE) = 12$. On the other hand, $\#S4 = 3$ and $\#BE = 9$ —again from the lists of elements in [PROBLEM 3.7.12i](#))—so $\#S4 + \#BE = 12$ too.

ii) $S7$ and BE

iii) OE and BE

iv) BO and BE

In parts [ii](#)) and [iv](#)), the two sets were disjoint and thus the (empty) intersection had order 0. Correspondingly, the order of the union of the two sets in each part was just the sum of their orders. This simpler, special case happens often enough that we want to recognize and take advantage of it. However, we also want to distinguish it carefully from the general case when the intersection is *not* empty for reasons I'll explain in a moment. We'll call the operation of taking the union of two sets *known to be disjoint* an *orElse* union and use



3.7 And, Or, and Not: the three hardest words

the typewriter font to emphasize its mathematical meaning even in ordinary text.

orelse/DISJOINT UNION OF SETS 3.7.18: *If we know that A and B are disjoint sets, then we write $A \dot{\cup} B$ and read “ A disjoint union B ” or “ A orelse B ” for their union. An object x is in $A \dot{\cup} B$ if x is in A orelse x is in B —since A and B are disjoint, x cannot be in both.*

If A and B are not disjoint, then their disjoint or orelse union is not defined.

ORELSE FORMULA FOR ORDERS 3.7.19: *If A and B are disjoint sets (so that $A \dot{\cup} B$ is defined) then*

$$\#(A \dot{\cup} B) = \#A + \#B.$$

You’ve probably already noticed that, when A and B are disjoint, the orelse union $A \dot{\cup} B$ and the eitherorboth union $A \cup B$ are the same set and hence have the same order. So why did I introduce the extra term orelse, the extra notation $A \dot{\cup} B$, and the extra formula $\#(A \dot{\cup} B) = \#A + \#B$.

For one very good reason. I’m hoping to head off at least some of the many errors that arise when students treat eitherorboth unions of sets that are *not* disjoint as if they were orelse unions. When $A \cap B = \emptyset$, it’s very tempting to view **AND-OR FORMULA FOR ORDERS 3.7.15** as saying that $\#(A \cup B) = \#A + \#B$. Don’t! Don’t learn this formula! Don’t ever even write this formula! Yes, it’s simpler to remember and to use, when it applies, but experience shows that knowing it is dangerous. Under pressure, you forget that the simpler form is only correct *when* the sets are disjoint, and you forget to *check* whether they are disjoint before applying it. Murphy’s law tells us when you do this, the two sets will *not* be disjoint and you’ll get the wrong answer.

My recommendation is to learn *only* the **AND-OR FORMULA FOR ORDERS 3.7.15**, with its builtin $\#(A \cap B)$, when using the formula and



3.7 And, Or, and Not: the three hardest words

to make sure you understand this order. After all, how much work is it really to add 0 on the left side when the two sets *are* disjoint? Unfortunately, few students seem to be able to stick with this recommendation, and they always seem to forget the intersection just when it's non-empty.

So I have bowed to the inevitable and written down the simpler **ORELSE FORMULA FOR ORDERS 3.7.19**, but I have used a special terminology (`orelse`) and notation ($\dot{\cup}$) for this formula, to remind you that $\#(A \dot{\cup} B) = \#A + \#B$ only applies *after* you have checked that A and B are disjoint.

ASK “WHICH OR?” 3.7.20: *Before computing the order of the union of two sets A and B , ask “Which or?”. Do I know the sets are disjoint, allowing me to use the **ORELSE FORMULA FOR ORDERS 3.7.19**, or do I need to use the **AND-OR FORMULA FOR ORDERS 3.7.15**?*

What about counting unions or intersections of more than 2 sets? Are there formulas for these? The answer is yes, but they quickly become much more complicated. Already when there are 3 sets, there are 8 terms in the formula. It's worth knowing such formulae exist but, fortunately, we won't have any need for them.

Here are some problems that illustrate the ways we can use the **AND-OR FORMULA FOR ORDERS 3.7.15** (and sometimes **ORELSE FORMULA FOR ORDERS 3.7.19** in its place).

PROBLEM 3.7.21: Majors in the Mathematics department must take courses in at least one of Algebra and Calculus but may take both.

i) If 24 math majors take Algebra, 28 take Calculus and 12 take both, how many majors are there?

Solution

If we let M be the set of majors, A be the set of students taking Algebra and C be the set of students taking Calculus, then we know that $M = A \cup C$ since majors must take “at least one” of these courses. We're given that $\#A = 24$, $\#C = 28$ and $\#(A \cap C) =$

3.7 And, Or, and Not: the three hardest words

12. So we can just rewrite **AND-OR FORMULA FOR ORDERS 3.7.15** and plug in to get

$$\#M = \#(A \cup C) = \#A + \#C - \#(A \cap C) = 24 + 28 - 12 = 40.$$

I introduced letters to denote the sets involved here to make it clear how we were using the **AND-OR FORMULA FOR ORDERS 3.7.15**, but this is not really necessary. I could have answered instead as follows. The number of majors is the number taking either Algebra or Calculus. This in turn is the number taking Algebra plus the number taking Calculus minus the number taking both, so there are $24 + 28 - 12 = 40$ Majors. Feel free either to write your answers in a more informal form as you become more familiar with the formula, or to continue using formal set notation if you find that easier.

- ii) If there are 44 math majors of whom 30 take Algebra and 30 take Calculus, how many majors take both?
- iii) If there are 38 math majors of whom 30 take Algebra and 20 take both Algebra and Calculus, how many majors take Calculus?
- iv) Suppose we know only that 32 majors take Algebra and 23 take Calculus. Decide whether each number below could equal the number of math majors or not. If it can, explain how. If it cannot, explain why not?
 - a. 30.
 - b. 40.
 - c. 50.
 - d. 60.

Partial Solution

Here we know (in the notation above) neither $\#(A \cap C)$ nor $\#(A \cup C)$ so we can't just solve for the latter using the **AND-OR FORMULA FOR ORDERS 3.7.15**. But we can solve for the former if we *assume* that $\#(A \cup C)$ takes on any of the given values and plug into $\#(A \cap C) = \#A + \#C - \#(A \cup C) = \#A + \#C - \#M$.

3.7 And, Or, and Not: the three hardest words

For example taking $\#M = 30$ would force $\#(A \cap C) = 32 + 23 - 30 = 25$. Now we need to decide if 25 is a possible value for $\#(A \cap C)$. It may seem that we have gained nothing. Not so. Since $A \cap C$ is a subset of both A and C it's order cannot be bigger than the order of either. Since $25 > 23$, we can't have $\#M = 30$. In this case, we could reached the same conclusion more easily—since $A \subset M$ we can't have $\#A$ bigger than $\#M$ —but in the other parts, you'll need to use the [AND-OR FORMULA FOR ORDERS 3.7.15](#).

At the end of the next subsection, we'll work some slightly harder problems that combine the [AND-OR FORMULA FOR ORDERS 3.7.15](#) with the [COMPLEMENT FORMULA FOR ORDERS 3.7.24](#).

butnot

This is much the easiest of the three hard words. We use it to describe a set by saying where it's elements do *not* lie. The only minor difficulty we need to address is where to stop. Suppose I'm considering several Harvard mathematics courses. I'd like to be able to talk about the set of students who are not taking Algebra. Let's denote the set of students who are taking algebra by A and the set of students who are not taking Algebra by B .

It certainly seems like this defines a perfectly good set B . Remember, from [OBJECTS AND ORACLES](#), that what that means is that I know how to tell whether any object x is an element of B . If x is math major Joe Blow who is taking Algebra, the answer is “No”. If x is math major Jane Doe who is not taking algebra, the answer is “Yes”. What if x is psychology major Wade Roe who is not taking Algebra? Now the answer is a bit fuzzy. Did we intend B to include only math majors? We have *not* described a set B unless we *also* give an answer to this question because if the answer is yes, then Wade is not in B and, if it's no, he is. Whatever we answer, we're still not out of the woods. What if x is Hei Yo who's a math major at the University of Tokyo?

3.7 And, Or, and Not: the three hardest words

The way to avoid this kind of ambiguity is to always say both “Yes” and “No” when defining a set by what it’s *not*. I’ll use `butnot` (in typewriter font) to highlight when we are using “not” in this mathematical sense.

COMPLEMENT OF A SUBSET 3.7.22: *If A is a subset of a set S , then we define the complement A^{cs} to be the set whose elements are the those elements of S that are not elements of A . In other words, $x \in A^{cs}$ means that $x \in S$ and $x \notin A$: in words, x is in S butnot in A .*

Very often, the set S will be a UNIVERSAL SET 3.7.7 that is clear from the context we are in—for example, when we work probability problems, S will almost always be what’s called the sample space. In such cases, we simplify and just write A^c instead of A^{cs} .

A subset A of S and its complement A^{cs} are related in several very special and useful ways.

- COMPLEMENT RELATIONS 3.7.23:** If A is a subset of a set S , then:
- i) $A \cap A^{cs} = \emptyset$: A and its complement are disjoint.
 - ii) $A \cup A^{cs} = S$: S is the union of A and its complement.
 - iii) $A \dot{\cup} A^{cs} = S$: S is the union of A and its complement.
 - iv) $(A^{cs})^{cs} = A$: A is the complement of its complement.

These are all easy to see from the tautological remark that: *every element of S is either in A or not in A and that no element of S is both in A and not in A . In terms of complements: every element of S is either in A or in A^{cs} and that no element of S is both in A and in A^{cs} .*

Parts i) and ii) just re-express the statement in italics in terms of union and intersection. Parts i) also says that A and A^{cs} are disjoint, so their disjoint union is defined and given by ii). Likewise, iv) follows: elements of $(A^{cs})^{cs}$ are by definition the elements of S not in A^{cs} and these are just the elements of S in A . Note that, since A and A^{cs} are disjoint, we can rewrite ii):

$$A \dot{\cup} A^{cs} = S.$$

3.7 And, Or, and Not: the three hardest words

Applying [OR ELSE FORMULA FOR ORDERS 3.7.19](#),

COMPLEMENT FORMULA FOR ORDERS 3.7.24: *For any $A \subset S$,*

$$\#(A^{cs}) = \#S - \#A.$$

PROBLEM 3.7.25: Derive [COMPLEMENT FORMULA FOR ORDERS 3.7.24](#), without using the [OR ELSE FORMULA FOR ORDERS 3.7.19](#), by combining [AND-OR FORMULA FOR ORDERS 3.7.15](#) and i) and ii) of [COMPLEMENT RELATIONS 3.7.23](#).

EXAMPLE 3.7.26: Even some politicians know the [COMPLEMENT FORMULA FOR ORDERS 3.7.24](#). Benjamin Disraeli, a 19th century British Prime Minister, once, as Leader of the Opposition, complained that their recent conduct showed that “Half the cabinet are asses”. Members of the cabinet protested loudly, leading the Speaker to demand that Disraeli retract his remark. Disraeli rose and said, “Mr Speaker, I withdraw. Half the Cabinet are not asses.”

PROBLEM 3.7.27: Explain why Disraeli, far from having retracted his insult, had managed to repeat it.

There’s one point about complements that’s best explained with an example.

- PROBLEM 3.7.28:** Consider a State Legislature with 235 members.
- i) If the Legislature contains 120 men, how many women does it contain?
 - ii) If the Legislature contains 90 golfers, how many tennis players does it contain?

Of course, the first answer is $235 - 120 = 115$ because the subset of women legislators is the complement of the subset of male legislators. And we can’t answer the second question, because there’s no such complementary relation between golfers and tennis players. Note that the word “not” is in neither question. We had to recognize that “women” is an antonym to “man” which is just a fancy way to say it describes the complement. Pay attention to such opposites.

3.7 And, Or, and Not: the three hardest words

It's easy no miss them, and find a question like the first above as puzzling as the second.

ANTONYMS DESCRIBE COMPLEMENTS 3.7.29: In relating subsets of a universal set be alert for complements that are described by antonyms, without using the word “not”.

The Four Quadrants and the Three Hard Words

Let's now introduce a classic schematic picture that we'll see many, many times in the rest of the course. The situation we want to model is a that of a universal set S (shown as the large square), and two sub-sets that I've called A (shown as the upper rectangle) and B (shown as the right rectangle). In addition to these sets, each of the four quadrants defines a subset of S that we'll often need to identify. These are marked in the venerable [FIGURE 3.7.30](#) which we'll call a **Q-diagram**.




FIGURE 3.7.30: A standard Q-diagram

We introduce a few terms. Although these are not standard, they make it easier to talk about such problems.

Q-SETS AND Q-PROBLEMS 3.7.31: We call [FIGURE 3.7.30](#), the Q-diagram (Q for quadrants) of a universal set S and two subsets A and B . A set C is a **Q-set** if C is a union of one or more of the quadrants in [FIGURE 3.7.30](#). A **Q-problem** is any problem in which we are given the order of several Q-sets and asked to compute the

3.7 And, Or, and Not: the three hardest words

other of one or more other Q -sets. In a proportional Q -problem, we are given and asked for, not $\#C$ but the proportion $\frac{\#C}{\#S}$ as fraction or percentage.

PROBLEM 3.7.32:

- i) Which quadrant is *not* in each of the Q -sets below?
 - a. $A \cup B$ or “A and also B”
 - b. $A \cup B^c$ or “A and not B”
 - c. $A^c \cup B$ or “not A and B”
 - d. $A^c \cup B^c$ or “not A and not B”
- ii) Both A and B are Q -sets containing exactly 2 quadrants. Find 4 more such 2-quadrant Q -sets.

PROBLEM 3.7.33: Here’s a little counting problem involving Q -sets.

- i) Show that there are exactly 16 different Q -sets? Hint: There’s exactly 1 Q -set for each subset of the set of 4 quadrants in [FIGURE 3.7.30](#).
- ii) Let \mathcal{Q} be the set consisting of the 14 non-trivial q -sets (that is, we leave out S itself and the empty set). How many 3 and 4 element subsets does \mathcal{Q} have?

First, some good news. Q -problems are generally very easy; that’s good because we’ll have to solve a great many in studying probability. The only difficulty Q -problems pose is that they come in a great many flavors, corresponding to the different Q -sets whose orders are given and asked for in these problems. Usually the we are given the order of either 3 or 4 non-trivial Q -sets. If you worked [PROBLEM 3.7.33](#), you’ll realize that that’s well over a thousand possible flavors. Even though most of these flavors never occur, we’ll still see lots of variants.

For this reason, many students like have a systematic way to answer such questions. I’m now going to explain my favorite system, the one that over the years my students have found easiest to use. On the other hand, there are many students who find it easy to solve



3.7 And, Or, and Not: the three hardest words

such problems using their native intelligence. If you are one of these students, you can just ignore this method that follows and solve these problems in whatever way is easiest for you.

The basic idea of the system is very easy. The order of Q -set is just the sum of the orders of the quadrants in it. That's because the intersection $Q \cap Q'$ of any two quadrants is empty, which means that the **AND-OR FORMULA FOR ORDERS 3.7.15** simplifies to $\#(Q \cup Q') = \#Q + \#Q'$. So we can immediately read off the order of any Q -set we may be asked about, once we know the orders of the 4 quadrants. These can always be found by solving some very easy equations extracted from the problem at hand. Proportional Q -problems work the same way except that we add up proportions rather than orders. Here's the method.

QUADRANT METHOD 3.7.34: *To solve a Q -problem:*

Step 1: Identify the sample space S in the problem and the subsets A and B : we'll see you're free to give these more mnemonic names.

Step 2: Each number given in the problem is the order (or proportion) of some Q -set. Identify the quadrants that make up this Q -set and write down the equation that expresses this number as a sum of orders (or proportions) of quadrants.

Step 3: Usually, you'll have 4 numbers and hence 4 equations. If you see only 3 numbers, you have a proportional Q -problem. Use the fact that the proportion $\frac{\#S}{\#S}$ equals 1 (or 100%) to get a 4th equation.

Step 4: Solve your equations. This usually involves only 2 simple operations. Substitute an known value, or subtract one equation from another to solve for an unknown value. Occasionally, you may need to subtract one equation from the sum of two others.

Step 5: Each number asked for in the problem is the order (or pro-



3.7 And, Or, and Not: the three hardest words

portion) of some Q-set. Identify the quadrants that make up this Q-set and find this number by adding up the orders (or proportions) of these quadrants.

This may sound a bit scary, since it asks you to solve 4 equations in 4 unknowns (the orders of the quadrants). Don't worry; solving these doesn't even call for any multiplication or division, just a bit of very easy addition and subtraction. In fact, the only place where you're likely to go wrong and need to pay attention is in the steps of matching quadrants to numbers and questions in the problem. And here the difficulty is all about interpreting the meaning of informal descriptions accurately. A few examples will make all this clear, and convince you these problems are not so bad.

EXAMPLE 3.7.35: Of 110 students in the Economics department, 35 have taken Calculus and 85 have taken Finance and 10 have taken both. How many have taken neither Calculus nor Finance? How many have taken exactly one of the two courses?

Solution

Step 1: The sample space is the set of economics students (E) and the subsets are those who have taken calculus (C) and Finance (F). Notice I used names that remind what these sets stand for, not the generic S , A and B . Ordinarily, you needn't repeat the diagram but let's set it down just this once:

E	$C \cap F^c$	$C \cap F$
	$C^c \cap F^c$	$C^c \cap F$

Step 2: What quadrant(s) does each number in the problem describe? The number 110 is the order of E which is the union of all 4 quadrants. The number 35 is the order of C which is the union of the two top quadrants and the number 85 is the

3.7 And, Or, and Not: the three hardest words

order of E which is the union of the two right quadrants. The number 10 is the order of bottom left quadrant, in neither C nor F . This translates to the equations:

$$\begin{aligned} \#(C \cap F) + \#(C \cap F^c) + \#(C^c \cap F) + \#(C^c \cap F^c) &= 110 \\ \#(C \cap F) + \#(C \cap F^c) &= 35 \\ \#(C \cap F) + \#(C^c \cap F) &= 85 \\ \#(C^c \cap F^c) &= 10 \end{aligned}$$

Step 3: We have the expected 4 equations in 4 orders.

Step 4: Now we solve equations. Here, substituting $\#(C^c \cap F^c) = 10$ in the first equation we get $\#(C \cap F) + \#(C \cap F^c) + \#(C^c \cap F) = 100$. Then subtracting the second equation from this gives $\#(C^c \cap F) = 65$. Now we substitute this value in the third equation to get $\#(C \cap F) = 20$. Substituting this in the second equation gives us $\#(C \cap F^c) = 15$. As a check, just plug your values back into all 4 equations and check that the two sides of each match.

Step 5: Notice that we haven't worried at all, up to this point, about what the question *asked*. But the work we have done now makes answering these questions easy. Students who took both courses are those in the upper right quadrant $C \cap F$ so there are 20 of these. Students who took exactly one of the two courses are those in top-left and bottom right quadrants, so the number of these is $\#(C \cap F^c) + \#(C^c \cap F) = 15 + 65 = 80$.

As I said, many students can see their way through the same steps without writing down explicit equations, and if you can you'll usually be able to shorten your solution a bit—though such informal solutions are more prone to errors. But many students find having a clear method in mind worth the slight extra effort. As I said, use whatever method works best for you. Even if you prefer informal at-

3.7 And, Or, and Not: the three hardest words

tacks, you can always write down equations if a particular problem stumps you.

EXAMPLE 3.7.36: A recent survey of American voters revealed that 60% were opposed to using Federal funds to bailing out Wall Street investment banks in 2008 but in favor of bailing out US automakers in 2009, that 44% were opposed to bailing out US automakers but in favor of bailing out the investment banks and that 33% supported both bailouts. How many opposed both bailing out the banks and bailing out the automakers?

Solution This sounds harder than [EXAMPLE 3.7.35](#), but when we write down equations we'll see it's simpler.

Step 1: The sample space is the set of voters surveyed (V) and the subsets are those who *supported* the bank bailout (B) and those who *supported* the automaker bailout (A). We'll have to be careful about complements here.

Step 2: Notice that here we have just 3 numbers and these numbers are percentages. Both these facts tell us that we have a proportional problem here and that our 4th equation will come from saying that the total in all 4 quadrants is 100%. The equations will turn out to be much simpler than in [EXAMPLE 3.7.35](#), even though the story here seems more complicated.

$\#(A \cap B)$	$+$	$\#(A \cap B^c)$	$+$	$\#(A^c \cap B)$	$+$	$\#(A^c \cap B^c)$	$=$	100
$\#(A \cap B)$							$=$	33
		$\#(A \cap B^c)$					$=$	60
				$\#(A^c \cap B)$			$=$	44

Step 3: We already found the “missing” 4th equation above.

Step 4: Here all we have to do is substitute the values for the three known quadrants to find the fourth: $33 + 60 + 44 + \#(A^c \cap B^c) = 100$ so $\#(A^c \cap B^c) = 37$.

Step 5: Those opposed to bailing out the banks are just those in B^c and as $B^c = (A \cap B^c) \cup (A^c \cap B^c)$, these were $\#B^c = \#(A \cap$

3.7 And, Or, and Not: the three hardest words

$B^c) + \#(A^c \cap B^c) = 60 + 37 = 97$ or 97% of voters. Likewise, those opposed to the automaker bailout were $\#A^c = \#(A^c \cap B) + \#(A^c \cap B^c) = 44 + 37 = 81$ or 81%.

Here are some more problems for you to practice with.

First some easy ones.

PROBLEM 3.7.37: A valid ballot in Florida must have exactly 1 punched chad. A precinct officer in the 2000 Presidential election counted 2240 votes, of which 1120 had a punched Bush chad, 1103 has a punched Gore chad and 230 had no chad punched. How many valid votes did he report for each candidate?

PROBLEM 3.7.38: The fraction of Canadian families with a pet polar bear is 0.05 and with a pet moose is 0.18. If the fraction with neither animal as a pet is 0.80, what fraction have a very nervous moose?

PROBLEM 3.7.39: If 22% of people are left-handed, 6% are left-handed and blonde, 35% are left-handed or blonde, and nobody is ambidextrous, find what percentage of the population is in each group below.

- i) right-handed.
- ii) blonde.
- iii) neither left-handed nor blonde.
- iv) right-handed and blonde.

Finally, one that's a bit harder.

PROBLEM 3.7.40: On the first midterm of the year, 22 students got As and on the second 26. There were 10 who got an A on neither test and the same number of students got an As on exactly 1 test as did on both. How many students were in the class?

When conjunctions collide

The aim of this short subsection is to call your attention to a couple of points about “and”, “or” and “not” that often trip up the unwary.



3.7 And, Or, and Not: the three hardest words

You'll only encounter a few of these traps in this course, but you will see all of them in later life, so I thought it worth making you aware of them while we're discussing the three hardest words.

First, it's very common to see "and" used in informal English where "or" is meant. Here are a couple of typical examples:

"Mathematics majors should take at least one of physics, computer science and statistics."

"Almost 5,000,000 people live in Houston, Dallas and San Antonio." Neither of these is likely to cause much confusion. In the first, the "at least one of" tells us that the writer meant "or" instead of "and". Of course, take this away and the difference is critical. If "Mathematics majors should take physics, computer science and statistics" there are three breadth requirements instead of one. The second sentence is nonsense as it stands. Nobody lives in "all three of" Houston, Dallas and San Antonio. But again this is harmless: because, when read literally, the sentence is *so wildly* false, we mentally change the "and" to the "or" the writer clearly must have meant. So my advice here is to make an effort, in your own writing, to avoid this kind of confusion and to be alert for it when reading what others write.

Things get a bit more serious when "and" and "or" collide. Consider: "Mathematics majors must take physics and computer science or statistics." "Southwest hopes to start flying into Houston or Dallas and San Antonio."

Both sentences are ambiguous. In the first, the requirement may be either both of physics and computer science, or just statistics; or maybe physics is required along with either computer science or statistics. In the second, does Southwest want to fly either to Houston or to both of Dallas and San Antonio (or to all three)? Or does it want to fly to at least one of Houston or Dallas, as well as to San Antonio? The issue in both cases is one of *order of operations*. Changing the order in which the semantic operations represented by "and" (both true) and "or" (at least one true) are carried out changes the

3.8 Counting by the “divide and conquer” method

meaning of the sentence. As for arithmetic operations, such ambiguities can be resolved either by applying a convention like [POGEMDAS 1.1.5](#) (the standard one is “and before or”) or by the semantic equivalent of parentheses. The latter is how I made the alternatives clear above and, here, it’s the *only* reliable way to make sure your meaning is clear. The moral when mixing “and”s and “or”s be sure to make the order you intend clear in what you write.

Even more caution is called for when “not” gets into the act. The following trap is the one you are most likely to fall into in **MATH⁴LIFE**.

“I didn’t take physics, computer science and statistics.”

Is the student in this sentence saying she didn’t take *all* 3 subjects? Or is she saying that she didn’t take *any* of them? Here it’s impossible to say with any certainty. If we write P , C and S for the set of students who have taken each subject, then those who have taken “physics, computer science and statistics” are the set $P \cap C \cap S$. To belong to this set, you need to have taken all 3. So the sentence as written means she did not take them all; that’s how you should read such sentences in this course. But, in everyday speech, what is usually meant is that the speaker took none of these subjects. “And” has been used where “or” was meant—“I didn’t take physics, computer science or statistics”. The difference from the sentence we started with is that nothing like the “at least one” before the list overrules the incorrect “and”. Had the speaker said “I didn’t take *any of* physics, computer science and statistics”, we’d be sure she meant “or”. So, when a negative or “not” is used, pay especial attention to “and” versus “or”.

3.8 Counting by the “divide and conquer” method

We have now developed all the technology needed to answer the counting problems that will arise in our study of probability. Still,



3.8 Counting by the “divide and conquer” method

if I started asking you the kinds of questions we’ll need to answer right now, I know that, while a few of you would find them a breeze, most of you would get very frustrated. Why? Well, the lucky few can just see how to reduce complicated counting problems to simple ones naturally. The aim of this section is to teach you a simple, very direct method for making such simplifications, and a set of **standard divisions** that you need to recognize to carry out this plan.

The method

There are three basic steps to our method.

DIVIDE AND CONQUER COUNTING STRATEGY 3.8.1: *To answer a counting or problem—that is a “How many ways?” or “How many choices?” question—try to apply the following three steps:*

Step 1: DIVIDE Try to break up the choice into simpler choices connected by andthen, or else and butnot. If a simpler choice asks for the answer to a SHORTHAND QUESTION 3.6.3, stop. If not, try to divide that choice into still simpler choices, again connected by andthen, or else and butnot. Repeat until each choice is the answer to a SHORTHAND QUESTION 3.6.3.

Step 2: SHORTHANDS Use the TWO QUESTION METHOD 3.6.4 to write down the shorthand count that answers each SHORTHAND QUESTION 3.6.3 you have isolated in the DIVIDE step. This step is easy; you just need to remember the FIRST RULE OF GUESSING 3.6.1 Don’t!

Step 3: REASSEMBLE Reassemble the shorthands, working bottom-up from the simplest choices back to the original choice, into a single multi-shorthand answer and evaluate this answer. For simpler choices that were connected by andthen, multiply the shorthands. For simpler choices that were connected by or else, add the shorthands. For two choices that were connected by butnot, take the difference of the counts.



3.8 Counting by the “divide and conquer” method

Before we start to work some examples to get a feel for this method, a few general comments. First, although the [DIVIDE AND CONQUER COUNTING STRATEGY 3.8.1](#) may appear a bit complex, most of the time you just follow your nose. We’ve already seen how the [TWO QUESTION METHOD 3.6.4](#) makes the SHORTHANDS step pretty cut and dried. I’ll just note one additional point about using SHORTHANDS.

SHORTHANDS REMEMBER 3.8.2: When applying the [DIVIDE AND CONQUER COUNTING STRATEGY 3.8.1](#), leave all answers to [SHORT-HAND QUESTION 3.6.3](#) in unevaluated shorthand form, until you need a final numerical answer.

For example, if you have two component choices with counts the shorthands $C(10, 3)$ and $C(8, 4)$ just leave them in this form rather than evaluating them as the numbers 120 and 84. If you need to make both these choices and hence multiply, leave the product as $C(10, 3) \cdot C(8, 4)$. Only when you need a final answer should evaluate to get $120 \cdot 70 = 8400$.

Why? In many problems, you’ll need answer variants of the problem you just solved with one or two small changes. Very often, the [SHORTHAND QUESTION 3.6.3](#) applies in the same way to this variant problem and you can just write down the answer by altering the shorthands to reflect the changes in the question. In our example, we might be asked to make the first 3 choices from 12 possibilities instead of 10. The shorthand answer $C(10, 3)$ remembers how the number 10 of possibilities entered the count and tells us that the new answer should be $C(12, 3)$. The answer 120 has forgotten the role of the number 10 and is of no help in writing down the new answer. Or, we might want to make 5 choices from amongst the 8 possibilities in the second count. Again, from the shorthand $C(8, 4)$, it’s easy to see that the new answer should be $C(8, 5)$ —the combination again remembers how the 4 was being used—and impossible from the number 84.

The REASSEMBLE step is also easy. When a choice involves making all

3.8 Counting by the “divide and conquer” method

of several simpler choices (we divided it into making the first simple choice, **and** then making the second, and so on), then the number of ways of making that choice is the product of the number of ways of making each of the simpler choices. This is just the **MULTIPLICATION PRINCIPLE 3.7.1**.

When a choice involves making **one** of several simpler choices (we divided it into making the first simple choice, **or** else making the second, and so on), then the number of ways of making that choice is the sum of the number of ways of making each of the simpler choices. Here we’re just using the **OR ELSE FORMULA FOR ORDERS 3.7.19**.

Finally, when a choice involves making one simpler choice **but not** a second simpler choice, then the number of ways of making that choice is the difference of the number of ways of making the first and second choices by the **COMPLEMENT FORMULA FOR ORDERS 3.7.24**.

To have a concise way of recalling these rules when working examples, we sum them up with the acronym **AMOANS 3.8.3**.

AMOANS 3.8.3: *To reassemble shorthand counts, just remember: All or And then : Multiply, One or Or else : Add, but Not : Subtract.*

This leaves the **DIVIDE** step. This calls for more thought than the other two, but, with practice, you’ll see that it, too, is straightforward. What you learn with practice is a toolkit of standard divisions that can be used to divide almost any problem into **SHORTHAND QUESTION 3.6.3** pieces. Only very rarely will we have to put on our thinking caps to see how to divide a problem. So all you really need to do to master the **DIVIDE** step is to work problems. The catch is that it’s impossible to say *how many* problems you’ll need to work before you get the knack. Some lucky students seem to be born knowing the how to apply the toolkit; others need to struggle through many problems. You’ll have to decide for yourself when you’ve caught on: you’ll know you have when you stop having to struggle and start to just “see” how to divide.



First Examples and the Most Common Patterns

Each of examples that follow introduces a new trick for dividing. Let’s start with a couple of the most common and then divisions. As we’re working through this, we’ll discover several ideas that come in handy in many problems.

First we’ll look at a variant of [PROBLEM 3.6.17](#). This dealt with card games in which each player is dealt a hand and then arranges his or her hand to group the cards of the same value or suit or both. We saw that, if the deck has m cards and a hand has ℓ , then we are choosing ℓ times from m possibilities and the shorthand for this count is $C(m, \ell)$. In applying the [TWO QUESTION METHOD 3.6.4](#), repetitions are not allowed (since we deal ℓ different cards) and order does not matter (it’s the same hand after we sort it). Now let’s ask some questions about bridge hands, where $m = 52$ and $\ell = 13$.

EXAMPLE 3.8.4: How many bridge hands contain 5 red and 8 black cards?

Solution

The important word in that question is the “and”. We need to pick 5 cards from the 26 red cards “and then” 8 cards from the 26 black cards. We have the and then flavor because we are making 2 choices (red cards and black cards, not one choice with 2 restrictions) and the two choices do not restrict each other in any way. We saw this type of division more informally in [PROBLEM 3.6.22](#).

That’s the divide step here: our choice divides into two [SHORT-HAND QUESTION 3.6.3](#) choices. Now we use the [TWO QUESTION METHOD 3.6.4](#) to find the right shorthand for each count. Let’s do the red cards first. Once again, our answers are **R?** “No” (we can’t deal any card more than once, regardless of color) and **O?** “No” (on what cards we are dealt matters, not what order they are dealt in) so the shorthand is $C(26, 5)$.

3.8 Counting by the “divide and conquer” method

PROBLEM 3.8.5: Show that there are $C(26, 8)$ choices for the 8 black cards.

All that’s left now is to reassemble. Since our two shorthand choices are linked by an andthen (or since we need to make *both*—in this case that’s *all*— choices), we multiply the shorthands and then evaluate to get the answer $C(26, 5) \cdot C(26, 8) = 65,780 \cdot 1,562,275 = 102,766,449,500$.

We can draw one very useful lesson from this example.

PRESERVATION OF SHORTHANDS 3.8.6: When we alter the parameters ℓ and m of a SHORTHAND QUESTION 3.6.3 without *otherwise* changing what we are choosing, the type of shorthand for the new count(s) and the old are the same.

In PROBLEM 3.6.17, we were choosing 13 cards from the 52 in the deck. In the previous example, we changed the ℓ and m (to 5 cards from the 26 red cards, and to 8 cards from the 26 black cards). But we were still choosing cards to make up a bridge hand; and since, in a bridge hand, repeated cards are not allowed and order does not matter, the shorthands for these new choices were again combinations.

PRESERVATION OF SHORTHANDS 3.8.6 can be a big timesaver, though you do need to exercise caution in applying it. Even if we’re talking bridge, if we need to choose suits and not cards, that “otherwise” comes into play and we’d need to start from scratch. But in many problems, you’re faced with a series of like choices differing *only* in their ℓ and m . When you are, there’s no need to apply the TWO QUESTION METHOD 3.6.4 more than once.

EXAMPLE 3.8.7: How many bridge hands contain 4 red and 9 black cards?

Solution

Here’s the first of many illustrations of the idea that SHORTHANDS REMEMBER 3.8.2 and make it easy to apply standard

3.8 Counting by the “divide and conquer” method

divisions—from the previous problem—to a new situation. All that’s changed from the previous question is that instead of 5 red and 8 black cards, we have 4 and 9. Exactly, the same reasoning tell us to divide into these two choices separately, find the shorthands $C(26, 4)$ and $C(26, 9)$ by [PRESERVATION OF SHORTHANDS 3.8.6](#), and reassemble by multiplying to get $C(26, 4) \cdot C(26, 9) = 14,950 \cdot 3,124,550 = 46,712,022,500$.

EXAMPLE 3.8.8: How many bridge hands contain 5 red and 7 black cards?

Solution

The answer is *not* $C(26, 5) \cdot C(26, 7)$. That number *is* the number of ways of choosing 5 red and 7 black cards, but these choices do not give us a bridge hand. There are only 12 cards, not 13 as required. The number of bridge hands with 12 cards, however, colored, is 0.

Stupid as this question is, it makes 2 good points. First, *always* stay on your toes when doing counting problems. It’s all too easy to just start calculating without thinking and the result is usually a brisk slap of palm to forehead. I won’t ask you trick questions like this one very often in this course, but I know that if I did put this on a midterm I’d see 43,270,084,000 a lot more often than 0.

Second, and more important, this problem underlines the fact that even though it looks like there are 2 new numbers on the previous part (the 5 and 8), there’s really only 1, because the context—we’re dealing bridge hands—demands that the numbers sum to 13.

It’s worth recording this last point.

CONSERVATION OF ℓ AND m 3.8.9: When we divide ℓ choices from a set of m possibilities using “andelse”, the number of choices in the pieces must total ℓ and the number of possibilities in the pieces must total m .



3.8 Counting by the “divide and conquer” method

EXAMPLE 3.8.10: How many bridge hands contain 3 red cards?

Solution

The answer is *not* $C(26, 3)$. That number is the number of ways of choosing 3 red cards, but we want a bridge hand and that calls for 13 cards. [CONSERVATION OF \$\ell\$ AND \$m\$ 3.8.9](#) tells us that we need to pick 10 more cards, and to end up with 3 red cards these will all need to be black. Having avoided this trap, we can make an andthen division, apply [PRESERVATION OF SHORTHANDS 3.8.6](#), and, since our [SHORTHANDS REMEMBER 3.8.2](#), multiply to get the count $C(26, 3) \cdot C(26, 10) = 2,600 \cdot 5,311,735 = 13,810,511,000$.

EXAMPLE 3.8.11: How many bridge hands contain 3 ♥, 4 ♦, 5 ♠ and 2 ♣?

Solution

The only new element in this problem is that instead of dividing it into 2 simpler choices by color, we need to divide into 4 choices by suit: we choose 3 ♥ andthen 4 ♦ andthen 5 ♠ andthen 2 ♣. Since there are 13 cards in each suit, the m will be 13 in each [SHORTHAND QUESTION 3.6.3](#). By [PRESERVATION OF SHORTHANDS 3.8.6](#), all the shorthands are combinations and [AMOANS 3.8.3](#) tells us to multiply these shorthands getting

$$C(13, 3) \cdot C(13, 4) \cdot C(13, 5) \cdot C(13, 2) = 286 \cdot 715 \cdot 1287 \cdot 78 = 20,527,933,140.$$

EXAMPLE 3.8.12: How many bridge hands contain 6 red cards, 5 ♠ and 2 ♣?

Solution

Here there are 3 simpler choices by (red, ♠ and ♣) and the only novelty is that the sets into which we have divided the 52 cards in the deck are not all of the same order. All that’s required is to keep track of these orders in our choices: m is 26 for the red cards but 13 for the ♠ and ♣. By [PRESERVATION OF SHORTHANDS 3.8.6](#), all the shorthands are combinations and [AMOANS 3.8.3](#) tells us to multiply these shorthands getting

3.8 Counting by the “divide and conquer” method

$$C(26,6) \cdot C(13,5) \cdot C(13,2) = 230230 \cdot 1287 \cdot 78 = 23,111,868,780.$$

Here are a few variants for you to try.

PROBLEM 3.8.13:

- i) Find the number of bridge hands with 6 red cards and 7 black cards.
- ii) Find the number of bridge hands with 4 ♥, 4 ♦, 3 ♠ and 3 ♣?
- iii) Find the number of bridge hands with 5 hearts, 4 ♣ and 4 ♠.
- iv) Find the number of bridge hands with 5 red cards and 4 ♣ and 4 ♠.
- v) Find the number of bridge hands with 5 ♣.
- vi) Find the number of bridge hands with 4 ♥ and 3 ♦.

There’s one point hidden in this problem that’s worth noting.

USE EMPTY PIECES AS PLACEHOLDERS 3.8.14: It’s often smart to include an empty piece in a choice—even when it’s not strictly needed to get the answer.

This, like **SHORTHANDS REMEMBER 3.8.2**, makes it easier to work with multiple variations of a problem by making answers with and without empty pieces look alike. A typical example where this can come in handy can be seen by comparing the answers to **EXAMPLE 3.8.11** and part **iii)** of **PROBLEM 3.8.13**. Choosing 3 ♥, 4 ♦, 5 ♠ and 2 ♣ leads to the shorthand $C(13,3) \cdot C(13,4) \cdot C(13,5) \cdot C(13,2)$. Choosing 5 ♥, 4 ♠ and 4 ♣ leads to the shorthand $C(13,5) \cdot C(13,4) \cdot C(13,4)$. The first has 4 factors and the second only 3, so we don’t see them as two variations on a common problem.

We can cure this by mentally changing the second question to ask for choices of 5 ♥, 0 ♦, 4 ♠ and 4 ♣. Those 0 diamonds are an empty piece that can be chosen in only $C(13,0) = 1$ way (and *not* 0!). Including this empty piece does not affect the numerical answer 657,946,575, but it *does* let us write the shorthand answer as

3.8 Counting by the “divide and conquer” method

$C(13, 5) \cdot C(13, 0) \cdot C(13, 4) \cdot C(13, 4)$ to match the first answer where there were no empty pieces.

PROBLEM 3.8.15:

- i) Show that there are 11,404,407,300 bridge hands with 4 ♥, 4 ♦, 3 ♠ and 2 ♣.
- ii) Show that there are 16,726,464,040 bridge hands with 4 ♥, 3 ♦, 3 ♠ and 3 ♣.

I hope you’re find all this easy enough to be a bit boring by now. Seems like we’ve done a lot of huffing and puffing to answer some pretty easy questions. Before you go on, please read over the previous problem (whether or not you attempted it—the answers given are correct) and use it to answer the following even easier question: “Which suit distribution is more common in bridge, 4-4-3-2 or 4-3-3-3?”

Do I hear any “Well, duh!”s? Good, because there are *more than twice as many* 4-4-3-2 hands as there are 4-3-3-3 hands. It’s surprises like this that make counting so much trickier than it at first seems. Just when you’re over your fear of climbing and ready to dash to the top of the ladder there’s a rung missing.

How can we simultaneously have half again as many hands distributed (4♥, 3♦, 3♠, 3♣) as (4♥, 4♦, 3♠, 2♣) (precisely, 1.47 times as many) and yet less than half as many 4-3-3-3 hands as 4-4-3-2 hands? The answer is that these ratios arise from answering two *different* questions.

The difference is subtle but (clearly) critical. In the first question, we specify the number of cards in each suit. In the second, we specify in the number of cards in 4 suits, but we do not specify *which suit has each number of cards*. We know how to answer the first question—we did so in [EXAMPLE 3.8.11](#) and [PROBLEM 3.8.15](#)—and in a moment we’ll see it’s not so hard to use this to answer the second question. But first, let me draw a moral:

3.8 Counting by the “divide and conquer” method

PHILLIP’S LAW 3.8.16: *Eternal alertness is the price of accurate counting².*

PHILLIP’S LAW 3.8.16 is the Murphy’s Law of counting problems. The moment your attention flags and you stop paying attention to small differences is when they’ll matter and your counts will be wrong. Go back and double check your answer to **PROBLEM 3.8.13.ii**)

OK, let’s count those distributions, say the 4-3-3-3 distribution as an first example. We’re almost there. Amongst all the bridge hands with this distribution are the 16,726,464,040 with 4 ♥, 3 ♦, 3 ♠ and 3 ♣. The others just have 4 cards in a different suit than ♥ and there are 4 choices for that 4-card suit. So choosing a 4-3-3-3 hands amounts to choosing a suit and then choosing a bridge hand with 4 cards in the chosen suit and 3 in each of the others. By **AMOANS 3.8.3**, the number of choices is $4 \cdot 16,726,464,040 = 66,905,856,160$.

With a bit more effort, the same approach handles the 4-4-3-2 distribution. There are 2 wrinkles. We need to choose not 1 but 2 suits to have 4-cards and then we also need to choose which of the remaining suits will have 3 cards and which 2. This last choice didn’t come up with the 4-3-3-3 distribution because we “number of cards” didn’t distinguish between the three 3 card suits.

First let’s choose the 4 card suits. This is a **SHORTHAND QUESTION 3.6.3**—we are choosing 2 suits from 4—so we need to apply the **TWO QUESTION METHOD 3.6.4**. The answers are no to both **R?** (we need 2 different 4 card suits) and to **O?** (for example, 4 ♥ and 4 ♠ gives the same distribution as 4 ♠ and 4 ♥), so the shorthand is $C(4, 2)$. Now there are 2 suits left from which we must choose the 3 card suit so the answer³ is 2. We multiply these, finding this time that

²This law is named for **Wendell Phillips** whose “Eternal vigilance is the price of liberty” it paraphrases.

³If we had been very picky and used the **TWO QUESTION METHOD 3.6.4**, we’d have found the shorthand $C(2, 1)$ by **PRESERVATION OF SHORTHANDS 3.8.6**. In future, when dealing with **EASY SHORTHANDS 3.6.7**, I’ll just give the count and let you dot the shorthand i’s.

3.8 Counting by the “divide and conquer” method

there are $C(4, 2) \cdot 2 = 6 \cdot 2 = 12$ choices for the suits⁴. So there are $12 \cdot 1,404,407,300 = 136,852,887,600$ 4-4-3-2 bridge hands.

Is there a systematic way to make the kinds of choices that just came up? The bad news is that the answer is yes. Worse, there’s a formula for the number of choices, a big ugly formula (If you want you can look it up and convince yourself. It’s called the [multinomial formula](#)). The good news is that I’m not going to explain the system to you; I’m not even going to write down the multinomial formula. We don’t really need it, as the example of the two distributions shows.

All we used to find those counts was our [DIVIDE AND CONQUER COUNTING STRATEGY 3.8.1](#). That’ll work equally well with the few other multinomial counts we’ll need to deal with. More generally, counting is a very complex craft and there are literally thousands of other formulas, like the multinomial formula, that can be used to “dry clean” specific problems. We’re not going to learn any of them. Whatever we need to count, we’ll be able to handle using what we’ve already learned and our [DIVIDE AND CONQUER COUNTING STRATEGY 3.8.1](#), and most of the time we’ll just be dumping the counts in and pressing “normal wash”. We’ll just need to remember [PHILLIP’S LAW 3.8.16](#) and be willing to do the odd piece of hand laundry.

PROBLEM 3.8.17:

- i) How many ways are there to choose a pair of suits (like the 2 red suits)?
- ii) How many ways are there to divide the suits into 2 pairs of suits (like the 2 red suits and the 2 black suits)?

Hint: The answers to the two parts are *not* the same.

One more word of advice about what’s involved in learning to count. The biggest benefit of getting lots of practice in counting is that you

⁴This, by the way, is another typical case in which the [MULTIPLICATION PRINCIPLE 3.7.1](#) can be used even though the second set from which we choose depends on the first choice made: if we chose ♠ and ♥ as our 4-card suits, the second choice is between ♣ and ♦; if we chose ♠ and ♦, it’s between ♣ and ♥, and so on. Again, I won’t mention this again but I encourage you to look out for it.

3.8 Counting by the “divide and conquer” method

develop the ability to recognize that an idea you’ve used in one context can be applied in another. Counting really becomes easy when you can just about always see that “this one is really just like ...”. You can’t achieve this by memorizing cases because, as the bridge hand examples we’ve already seen should make clear, even sticking close to a single topic, there are just too many variations and wrinkles to codify.

SEE SIMILARITIES, NOT DIFFERENCES 3.8.18: When a counting problem involves a new situation, focus on the ways that it’s *like* problems you already understand, not the ways that it’s different. Use these similarities as your guide in the **DIVIDE AND CONQUER COUNTING STRATEGY 3.8.1**. They’ll let you apply divisions, short-hands and other ideas you already know to the new situation.

Here are a few more bridge counts for you to practice with. While all the ideas are in the examples above, I have tried to make these problems a bit different and have provided pointers to the examples that share the same idea.

PROBLEM 3.8.19:

- i) How many bridge hands have no red cards? Hint: This is almost the same as **EXAMPLE 3.8.10**.
- ii) A bridge hand with no cards in a suit is said to have a void in that suit. How many bridge hands have a void in hearts?

PROBLEM 3.8.20: Let us suppose that there are 59 Democrats and 41 Republicans in the 100 member US Senate. The Senate Budget Committee has 5 members.

- i) How many ways can the budget committee be chosen?
- ii) How many ways can the budget committee be chosen if it must consist of 3 Democrats and 2 Republicans?

PROBLEM 3.8.21: Once again, let us suppose that there are 59 Democrats and 41 Republicans in the 100 member US Senate, but this time let us suppose that the Senate Budget Committee has a Chair and 4 other members.

3.8 Counting by the “divide and conquer” method

- i) How many ways can the budget committee be chosen?

Solution

What’s different from i) of [PROBLEM 3.8.20](#) where you should have got the answer $C(100, 5) = 75,287,520$? We can now distinguish Budget Committees with the same 5 members but different Chairs. To handle this, we need a new way of dividing. The solution is both obvious and widely applicable. If part of a choice is special in some way, make that part of the choice separately. Here, we need to choose the chair and then the other members separately.

There are 100 choices for the chair (by [EASY SHORTHANDS 3.6.7](#)). Then we need to choose the other 4 members but from 99 Senators (not 100) because we’ve already picked a Chair and the choice of committee members is one where **R?** (and **O?**) are both “No”. So there are $C(99, 4)$ choices for the members and in all $100 \cdot C(99, 4) = 376,437,600$. As a check, notice that each committee in [PROBLEM 3.8.20.i\)](#) turns into 5 committees here (each of the 5 members can be singled out as chair) and the answer indeed differ by a factor of 5.

- ii) How many ways can the budget committee be chosen if the chairman must be a Democrat?
- iii) How many ways can the budget committee be chosen if the chairman must be a Democrat and 2 of the other members must come from each party?

Next let’s look at some problems that show the most common types of `orelse` division. The most common way `orelse` divisions arise is when we know how to count choices given a specific value for a parameter (usually the ℓ in some [SHORTHAND QUESTION 3.6.3](#)) but the problem only gives us inequalities for that parameter.

DIVIDE INEQUALITIES USING `orelse` 3.8.22: If you encounter words involving inequalities—like “at least”, “at most”, “fewer than”, “no more than” and so on, replace the inequalities with with several

3.8 Counting by the “divide and conquer” method

exact values connected by `orElse`.

EXAMPLE 3.8.23: To prepare for the `orElse` questions below, let's first ask: How many ways can a 6 member committee be chosen if it must contain exactly ℓ Democrats?

Solution

In analogy with [EXAMPLE 3.8.10](#), choosing ℓ Democrats does not specify a complete committee, but we can fix this by adding the right number of Republicans. If, for example, we take $\ell = 2$, then we'd divide into choosing 2 of 59 Democrats ($C(59, 2)$ choices) and then $6 - 2 = 4$ of 41 Republicans ($C(41, 4)$ choices); multiplying we'd get $C(59, 2) \cdot C(41, 4)$.

The general case works identically: we'd divide into choosing ℓ of 59 Democrats ($C(59, \ell)$ choices) and then $6 - \ell$ of 41 Republicans ($C(41, 6 - \ell)$ choices). Multiplying we get:

$$C(59, \ell) \cdot C(41, 6 - \ell).$$

This is a typical example of the principle that [SHORTHANDS REMEMBER 3.8.2](#). It also illustrates the advice to [USE EMPTY PIECES AS PLACEHOLDERS 3.8.14](#). If ℓ happens to equal 0, then we've included an empty set of Democrats with count $C(59, 0)$ so we'll always have a “Democrat factor” on the left. The same applies (with no Republicans) when $\ell = 6$ i.e. $6 - \ell = 0$.

EXAMPLE 3.8.24: Let us suppose that there are 59 Democrats and 41 Republicans in the 100 member US Senate. The Senate Judiciary Committee has 6 members.

- i) How many ways can the budget committee be chosen if it must contain at least 3 Democrats?

Solution

The words “at least” in the question tell us that we need to divide using `orElse`. More specifically, [DIVIDE INEQUALITIES USING `orElse` 3.8.22](#) tells us that we should replace the inequality “at least 3 Democrats” with several simpler pieces involving an exact number of Democrats.

3.8 Counting by the “divide and conquer” method

Since there are 6 members on the Committee, at least 3 Democrats is equivalent to exactly 3 Democrats *or* else exactly 4 Democrats *or* else exactly 5 Democrats *or* else exactly 6 Democrats. There’s no real need to say “exactly” each time. I did it here to emphasize how we were implementing **DIVIDE INEQUALITIES USING *or* else** 3.8.22 but in future we’ll omit it.

Each of the counts with an exact number of Democrats is the answer to **EXAMPLE 3.8.23** for the corresponding value of ℓ . So we can just plug into the formula given in that example. Since they are connected by *or* else’s, **AMOANS 3.8.3** tells us to sum these contributions to get

$$(C(59, 3) \cdot C(41, 3)) + (C(59, 4) \cdot C(41, 2)) + (C(59, 5) \cdot C(41, 1) + C(59, 6)) \cdot (C(41, 0)) = 970,068,560$$

Note how the empty piece at the end is needed to maintain the pattern of the terms in the sum. The parentheses around the products are not needed (by **POGEMDAS 1.1.5**) by I included them to emphasize that we want to compute the number of committees with each distribution first, and then sum up.

There’s one other point we haven’t addressed. We’re only allowed to use *or* else to connect two sets that we *know* are disjoint. Of course, the sets above don’t intersect because you can’t have two different numbers of Democrats on a committee, so we’re in good shape here. The same observation applies quite generally whenever we rephrase inequality conditions as several exact values.

- ii) How many ways can the budget committee be chosen if it must contain at least 3 Democrats and at least 2 Republicans?

Solution

Here it’s easier to make the exact choices specify the number from both parties. The inequalities “at least 3 Democrats” and “at least 2 Republicans” are equivalent to (3 Democrats and 3 Republicans) *or* else (4 Democrats and 2 Republicans). Here we get $(C(59, 3) \cdot C(41, 3)) + (C(59, 4) \cdot C(41, 2)) = 719,749,260$.



3.8 Counting by the “divide and conquer” method

iii) How many ways can the budget committee be chosen if it must contain at least at least 1 Republican?

Solution

We could handle this like the previous two parts. At least 1 Republican is equivalent to exactly 1 Republican *or* exactly 2 Republicans *or* exactly 3 Republicans *or* exactly 4 Republicans *or* exactly 5 Republicans *or* exactly 5 Republicans *or* exactly 6 Republicans.

If you think this looks a bit tedious to evaluate, I agree and fortunately there’s a better way. Instead of saying the committees we want to include, let’s ask what committees we want to *exclude*. That’s easy: those with 0 Republicans. In other words we can express the committees we are trying to count as all choices of 6 members of the Senate *but not* choices of 6 Democrats. The former count is $C(100, 6)$ (just change the 5 to a 6 in part i) of [PROBLEM 3.8.20](#)) and the latter is $C(59, 6)$ (change the 100 Senators to 59 Democrats) and then [AMOANS 3.8.3](#) tells us to take the difference to obtain

$$C(100, 6) - C(59, 6) = 1,192,052,400 - 45,057,474 = 1,146,994,926.$$

PROBLEM 3.8.25: Check this answer by using the *or* *else* division above.

iv) How many ways can the budget committee be chosen if it must contain at least at least 2 Republicans?

v) Are there more committees with an even or an odd number of Republicans?

Solution

This problem is an example of a more general division using *or* *else*. The even and odd conditions do not involve inequalities, but we can nonetheless use *or* *else* to re-express them in terms of several simpler counts with an exact number of Republicans. An odd number of Republicans is equivalent to 1 Republican *or* 3 Republicans *or* 5 Republicans. This gives

3.8 Counting by the “divide and conquer” method

$$(C(59, 5) \cdot C(41, 1)) + (C(59, 3) \cdot C(41, 3)) + (C(59, 1) \cdot C(41, 5)) = 596,022,248$$

PROBLEM 3.8.26: Show that the number of committees containing an even number of Republicans is 596,030,152 in two ways. First, add up the numbers of committees with each even number of Republicans. Second, take a difference to check your first answer.

So there slightly more committees with an even number of Republicans than with an odd number.

Now a few problems dealing with coin tosses that illustrate the principle [SEE SIMILARITIES, NOT DIFFERENCES 3.8.18](#) because the ideas from the preceding problems work in them too.

PROBLEM 3.8.27: This problem uses two counts from [PROBLEM 3.6.14](#), [EXAMPLE 3.6.15](#) and [PROBLEM 3.6.16](#).

- We view a sequence of m tosses of a coin as a sequence of length m in the 2-letter alphabet H and T, and there are 2^m of these.
- To determine a sequence of m tosses of a coin that contains exactly ℓ Hs, we need to choose the ℓ -element subset of the m tosses on which a head comes up, and this can be done in $C(m, \ell)$ ways.

Use these fact to find the following counts:

- The number of ways of obtaining exactly 2 heads in 8 tosses.
- The number of ways of obtaining at most 2 heads in 8 tosses.
- The number of ways of obtaining at least 6 heads in 7 tosses.
- The number of ways of obtaining between 48 and 52 heads in 100 tosses.
- The number of ways of obtaining at least 1 head in 100 tosses.
- The number of ways of obtaining at most 98 heads in 100 tosses.

PROBLEM 3.8.28: Here are a couple of slightly harder variants:

- Is the number of ways of getting an even number of heads when you toss a coin 7 times, greater than, equal to, or less than the number of ways getting on odd number of heads?



3.8 Counting by the “divide and conquer” method

ii) Is the number of ways of getting an even number of heads when you toss a coin 8 times, greater than, equal to, or less than the number of ways getting on odd number of heads?

More Challenging Examples

CHALLENGE 3.8.29: Here are two counts that I pose as challenges.

i) Is the number of ways of getting an even number of heads when you toss a coin 777 times, greater than, equal to, or less than the number of ways getting on odd number of heads?

ii) Is the number of ways of getting an even number of heads when you toss a coin 888 times, greater than, equal to, or less than the number of ways getting on odd number of heads?

These are examples of problems that you’d be able to solve *if* somebody put a gun to your head and threatened to shoot you unless you worked out the answer, but that you’d never try to solve *unless* somebody put a gun to your head and threatened to shoot you unless you worked out the answer. To answer the problem—say for 777—we just need to add up the numbers $C(777, \ell)$ first for every odd ℓ between 1 and 777, and then for every even ℓ between 0 and 776, and finally compare the totals. In principle, these problem are no harder than those in [PROBLEM 3.8.28](#). But “Not”, to quote Dana Carvey in his Bush XLI mode, “gonna happen.”

However, the answers to [PROBLEM 3.8.28](#) suggest that the two numbers might be equal and, if so, maybe there’s an easier way. Maybe we can see they’re equal without ever having to total them up. For an odd m like 777, it’s quite easy to see this. For a clue, first answer these questions.

i) Compare the number of ways of ways are there of obtaining 7 heads in 777 and the number of ways of obtaining 770 heads in 777 tosses.

3.8 Counting by the “divide and conquer” method

ii) Compare the number of ways of ways are there of obtaining 8 heads in 777 and the number of ways of obtaining 769 heads in 777 tosses.

It’s no accident that the two counts are the same in each case. Equality is guaranteed by SYMMETRY OF BINOMIAL COEFFICIENTS 3.4.23 because $770 + 7 = 669 + 8 = 777$. In general we know that the number of ways of getting ℓ heads in 777 tosses equals the number of ways of getting $777 - \ell$ heads because $C(777, \ell) = C(777, 777 - \ell)$.

i) Explain why this tells us that the number of ways of getting an even number of heads when you toss a coin 777 times equals the number of ways getting on odd number of heads. Hint: If ℓ is even, what can you say about $777 - \ell$?

ii) Show that the number of ways of getting an even number of heads when you toss a coin m times equals the number of ways getting on odd number of heads whenever m is an *odd* number.

This approach breaks down when m is an even number like 888 because if ℓ is even—say $\ell = 10$, then so is $m - \ell$ —here $m - \ell = 878$. Fortunately, there’s an even easier way to hand both odd and even m . We can bootstrap our way up, one m at a time. The key remark is that:

Every sequence S of length m in the letters H and T can be written as $S'H$ or else $S'T$ for a unique sequence S' of length $(m - 1)$ in H and T.

The sequence S' is just what’s left after you cross out the last letter in S . What’s more, either S' is a sequence E of length $(m - 1)$ with an *even* number of Hs or else its a sequence O of length $(m - 1)$ with an *odd* number of Hs. So we can write each sequence S of length m as a sequence EH or else as a sequence ET or else as a sequence OH or else as a sequence OT .

Use this to show that if the number of sequences of length $(m - 1)$ with an even number of Hs equals the number with an odd number of Hs (that is, if the number of E s equals the number of O s), then the



3.8 Counting by the “divide and conquer” method

number of sequences of length m with an even number of Hs equals the number with an odd number of Hs. What do you conclude for $m = 888$?

CHALLENGE 3.8.29 is a good illustration of why mathematics is often described as the art of being “intelligently lazy”.

THINK TWICE, CALCULATE ONCE 3.8.30: This is the mathematicians version of the old carpenter’s maxim “Measure twice, cut once”. In mathematics, it’s not lumber that you might waste, it’s your own time and effort. Before committing to any potentially lengthy calculation, it’s always worth asking yourself if there isn’t an easier way.

If you make this a habit, you’ll be surprised how often a bit of thought can take the place of a lot of arithmetic. If no inspiration strikes, you can always get out your TI-8x, but if you won’t be struck by inspiration if you *don’t* look for it.

Here’s one more harder problem, followed by a few hints to guide you to a solution:

PROBLEM 3.8.31: Show that 32,427,298,180 bridge hands have a void in some suit?

In part ii) of **PROBLEM 3.8.19**, we computed the number of hands with a void in ♥. We need to choose 13 of the 39 non-♥ and then 0 of the 13 ♥. This gives $C(39,13) \cdot C(13,0) = 8,122,425,444$. A hand with a void needn’t have a void in ♥ but it must have a void in one of the 4 suits so why isn’t the answer just $4 \cdot 8,122,425,444 = 32,489,701,776$?

The answer is because the larger number counts some hands *more than once*. This kind of difficulty—counting all the things you want but counting some of them more than once—is called **overcounting**. It’s a source of a great many pitfalls in counting problems and it’s often quite a nuisance to get around it.

The cause is almost always using *or* else where *either* or *both* is correct: that is, adding several counts as if the sets they stood for

3.8 Counting by the “divide and conquer” method

were disjoint when they are not. Here, for example, the set of hands with a void in ♥ and the set with a void in ♠ are not disjoint. Any hand that contains only ♦ and ♣ (and there are a lot!) is in both sets and so is counted *twice* in that 32,489,701,776.

How do we get around such problems? We could avoid the overcounting if we knew the number of hands with a void in ♥ *and no void in any other suit*. Lets say that such a hand has an onevoid in ♥. Now the sets of hands with a onevoid in ♥ and the set with a onevoid in ♠ *are* disjoint.

However, this creates a new problem, **undercounting**. Now, we’ve entirely omitted to count hands that have only ♦ and ♣, or in general that have voids in more than one suit. The first part is easy.

PROBLEM 3.8.32: Show that $C(26, 13)$ bridge hands contain only ♦ and ♣

But can you see what will happen if we try to use this count to find the number of bridge hands that have voids in more than one suit? Right, we’ll overcount one-suiters like the hand that consists of all 13 ♣. This hand will be counted 3 times: once for containing only ♦ and ♣, once for containing only ♥ and ♣, and once for containing only ♠ and ♣.

The right solution is usually not to work, as we have just done, from the more common to the special. Instead we try to count the most special objects—here the one-suiters—first. Then we try to use work our way up step by step. Here’s how to carry this out for hands with voids.

PROBLEM 3.8.33: We start from the obvious remark: each suit gives rise to just 1 one-suiter, the hand containing the 13 cards in that suit.

i) Show that there are 4 hands that have a threevoid (that is, a void in exactly 3 of the 4 suits). Of course, these are just one-suiters.



3.8 Counting by the “divide and conquer” method

- ii) Use this to show that there are $C(26, 13) - 2$ bridge hands that contain only \heartsuit and \clubsuit but are *not* one-suiters. Hint: Use the NS in [AMOANS 3.8.3](#).
- iii) Show that there are $C(4, 2) \cdot (C(26, 13) - 2)$ two-suited bridge hands—that is, hands that contain cards from only 2 suits but that are *not* single suiters. Hint: How many ways are there to choose the 2 suits?
- iv) How many bridge hands have a two-void (that is a void in exactly 2 of the 4 suits)? Hint: You have already computed this answer above.
- v) Show that there are $C(3, 2) \cdot (C(26, 13) - 2)$ two-suited bridge hands that contain no \spadesuit . Hint: How many ways are there to choose the 2 suits if neither can be \spadesuit ?
- vi) Show that there are 3 one-suited bridge hands that contain no \spadesuit . Hint: How many ways are there to choose the 1 suit if it cannot be \spadesuit ?
- vii) How many bridge hands contain no \spadesuit and are one-suiters or else two-suiters? That or else shows what we gained by working from the bottom up: we avoided overcounting.
- viii) How many bridge hands contain no \spadesuit ?
- ix) How many bridge hands contain no \spadesuit and are neither one-suiters nor two-suiters? These hands contain cards from each of \heartsuit , \diamondsuit and \clubsuit but no spades. That is they have a one-void in spades. Hint: Reassemble the two previous answers.
- x) How many bridge hands have a one-void in some suit?
- xi) How many bridge hands have a one-void or else a two-void or else a three-void? Again, all this work was to be able to say or else correctly.
- xii) How many bridge hands have a void? Hint: Although the question does not say it, “have a void” means “have a void in at least 1 suit”.

Believe it or not, this is the *easy* way to count voids. Here’s the [hard way](#) as a check on our calculations.



Examples with Permutations and Powers

At this point, you may be wondering why we included the sections on [SECTION 3.3](#) and [SECTION 3.5](#), since the corresponding shorthands, powers and permutations, have hardly appeared. While it will remain true that combinations arise most often, here are some typical examples of questions that lead to the other shorthands.

First a very easy example, similar to [PROBLEM 3.6.23](#), that uses only the obvious, direct divisions.

EXAMPLE 3.8.34: A local phone number consist of 7 decimal digits from 0 to 9.

- i) How many “raw” local phone numbers are there?

Solution

We need to choose 7 digits from a set of 10 possibilities. The answer to “Are Repetitions allowed?” is “Yes”—taxi companies and pizzerias especially like such repetitions—so phone numbers are sequences and there are $10^7 = 10,000,000$ of them.

- ii) How many “raw” local phone numbers have no repeated digits?

Solution

We still need to choose 7 digits from a set of 10 possibilities but with the difference that now **R?** is “No”. So we need to ask **O?** to which the order is yes, either by visualizing how we use phone numbers (we always dial the digits in a phone number in the same order) or by recalling [R IMPLIES O 3.6.6](#). So such a phone number is a list and the shorthand is $P(10, 7) = 604,800$.

- iii) A “real” phone number cannot begin with a 0 or a 1 which are reserved as escapes to long distance services. How many “real” local phone numbers are there?

Solution

To handle the special conditions on the first digit, we pick it separately. So we want to first pick 1 digit but not 0 or 1, possible in $10 - 2 = 8$ ways, and then 6 more digits, possible in 10^6 ways by

3.8 Counting by the “divide and conquer” method

PRESERVATION OF SHORTHANDS 3.8.6 since we are again picking digits. In all, we get $8 \cdot 10^6 = 8,000,000$ “real” phone numbers.

PROBLEM 3.8.35: Here are some local phone number variants for you to work yourself.

- i) How many “real” local phone numbers have no repeated digits? Hint: The answer is *not* $8 \cdot P(10, 6)$.
- ii) How many “real” local phone numbers are even?
- iii) How many “raw” local phone numbers use only odd digits?
- iv) How many “real” local phone numbers use only odd digits?
- v) How many “real” local phone numbers use only distinct odd digits?

Now let’s roll a few dice. This is a lot like tossing coins except that each roll gives us 6 rather than 2 possibilities. However, while we picture tossing a *single* coin several times in succession, we usually want to think of rolling several dice at the *same* time. To keep the various rolls straight, we’ll imagine that each die has a different color, and when you visualize rolls in counting them, you should try to see dice of different colors. In these problems, [SEE SIMILARITIES, NOT DIFFERENCES 3.8.18](#) applies again and again.

EXAMPLE 3.8.36: In this example, we’ll consider rolls of 3 dice.

- i) How many different ways can the 3 dice come up?

Solution

We are choosing 3 numbers from the possibilities 1 through 6 and **R?** is “Yes” because there’s no reason two or more of the dice can’t come up the same way. So the shorthand is $6^3 = 216$.

This is reminiscent of [EXAMPLE 3.8.34](#)

PROBLEM 3.8.37: Show that the number of ways that m dice can come up is 6^m .

- ii) How many of these rolls have no repeated number?

Solution



3.8 Counting by the “divide and conquer” method

As in [EXAMPLE 3.8.34](#), we have now made the answer to **R?** a “No” but, by [R IMPLIES O 3.6.6](#), the answer to **O?** is still a “Yes” so there are $P(6, 3) = 120$ possibilities.

iii) On how many rolls do exactly 2 dice come up 6?

Solution

This reminds us of trying to count the number of coin tosses where a fixed number of heads appears. We have replaced coins by dice and H by 6 but in both cases we are looking for sequences with a fixed number of a specific letter. The idea that worked for Hs—ask what subset of the positions in the sequence (for coins, which tosses and for dice, which colors) the letter occupies—is again effective.

Just in that word “subset” [SEE SIMILARITIES, NOT DIFFERENCES 3.8.18](#) has paid off: by remembering that what we chose was a subset we can bypass the [TWO QUESTION METHOD 3.6.4](#) (though we can check that the answers to **R?** and **O?** *are* both “No”) and predict that the shorthand will be the combination $C(3, 2)$. The only difference is that, whereas in tossing coins, we knew that any toss that wasn’t an H had to be a T, here we only know that the dice that don’t come up 6 come up with a number from 1 to 5. We still have to choose that number which we can do in 5 ways. Since we need to know which tosses came up 6 and then what number came up on the other toss, there are $C(3, 2) \cdot 5 = 3 \cdot 5 = 15$ possibilities.

iv) On how many rolls does exactly 1 dice come up 6?

Solution

This is almost the same as the previous part. We need to choose a subset of 1 of the 3 tosses to come up 6 which we can do in $C(3, 1)$ ways, and then decide how many possibilities there are for the other 2 tosses. This last involves making 2 choices from the numbers 1 to 5.

Because the $\ell = 2$ here, we no longer have one of the [EASY SHORTHANDS 3.6.7](#) and must apply [TWO QUESTION METHOD](#)

3.8 Counting by the “divide and conquer” method

3.6.4. But all that’s changed from i) is that $\ell = 2$ and $m = 5$, so the answers remain “Yes” to both **R?** and **O?** and we get the shorthand 5^2 . So there are $C(3, 1) \cdot 5^2 = 3 \cdot 25 = 75$ such rolls. We can help ourselves see a pattern that will be useful in similar problems by going back and rewriting the answer to iii) as $C(3, 2) \cdot 5^1$.

PROBLEM 3.8.38: Show that the number of rolls of m dice when any given number (like the 6 above) comes up exactly ℓ times is $C(m, \ell) \cdot 5^{(m-\ell)}$.

v) On how many rolls of 3 dice do no dice come up 6 and on how many do all the dice come up 6?

Solution

These are just the cases $m = 3, \ell = 0$ and $m = 3, \ell = 3$ of the formula so we get $C(3, 0) \cdot 5^3 = 125$ and $C(3, 3) \cdot 5^0 = 1$ respectively.

As a check, the number of 6s on *any* roll of 3 dice is exactly 0 or else exactly 1 or else exactly 2 or else exactly 3 which predicts that $125 + 75 + 15 + 1$ should, as it does, total 216.

I hope this example has reminded you of the game **CHUCK-A-LUCK** (and its variants like Sic Bo). We can now justify, at least informally, the claim made there that, “On average, $\$ \frac{17}{216} \simeq 0.07870370370$ or roughly 7.87 cents for every dollar bet.”

Imagine that you played chuck-a-luck 216 times, betting on 6 each time—I hope its clear that what number you bet makes no difference—and that each of the 216 possible rolls came up exactly once. Of course, this would essentially never happen in real life but it models perfectly something we’ll study later called an **expected value**.

How much would you win or lose? Well, on 125 rolls you’d get no 6 and lose your \$1, on 75 you’d get one 6 and win \$1, on 15 you’d get two 6s and win \$2 and on 1 you’d get three 6s and win \$3. You’d net

3.8 Counting by the “divide and conquer” method

$\$(-125 + 75 + 2 \cdot 15 + 3 \cdot 1) = -\17 : in other words, you lose 17 of the 216 dollars you bet as claimed.

We can also check the argument that you’ll break even playing chuck-a-luck. This, remember, was based on the idea that you expect to see, on average, $\frac{1}{2}$ a 6 come up every time you roll 3 dice. The number of 6s we’d see in our 216 rolls is $75 + 2 \cdot 15 + 3 \cdot 1 = 108 = \frac{1}{2}216$ as predicted.

So what gives here? The answer is that you’d break even if, instead of getting your bet back plus \$1 for each 6 that came up, you received \$2 for each 6 that came up. Your bets \$216 would just match your winnings $\$2 \cdot (75 + 2 \cdot 15 + 3 \cdot 1) = \216 . Once it’s been explained, the difference is so obvious that it’s hard to remember that we couldn’t see at first. The proof that people can’t see the difference is that chuck-a-luck remains a widely popular betting game.

Here are a few easy dice problems for you to try.

PROBLEM 3.8.39:

- i) If we roll 5 dice, how many ways can we have exactly 4 dice come up 3?
- ii) If we roll 5 dice, how many ways can we have exactly 2 dice come up 3?
- iii) If we roll 5 dice, how many ways can we have no repeated number appear?

PROBLEM 3.8.40:

- i) If we roll 7 dice, how many ways can we have exactly 3 dice come up 2?
- ii) If we roll 7 dice, how many ways can we have exactly no die come up 2?
- iii) If we roll 7 dice, how many ways can we have no repeated number appear?

PROBLEM 3.8.41: The most common dice games involve 2 dice and the bets concern the total of the 2 numbers on the dice. Make a table



3.8 Counting by the “divide and conquer” method

showing the number of rolls that yield each of the possible totals from 2 to 12 Hint: Imagine that we have, as usual a red die and a blue die. If we want a total of t and the red die comes up r then the blue die *must* come up $t - r$: i.e. the red die determines the blue one. Thus at most 6 rolls can give any total, but there are usually fewer because we must also have $1 \leq t - r \leq 6$.

The next problem is conceptually straightforward, but the short-hands that arise are too big for your calculator. Some are too big even to write down so I have given floating point approximations.

EXAMPLE 3.8.42: Let’s assume that the Census Bureau’s estimate that there are 227,719,424 Americans citizens is exactly right and that each citizen has a distinct social security number. You are building a spreadsheet with 18,000 rows, each of which will contain data about one such American. In the first column of each row of your spreadsheet you enter a social security number chosen at random from the Census Bureau’s master list.

- i) How many different first columns can the spreadsheet have?

Solution

We are making 18,000 choices from 227,719,424. To determine the flavor, we need to ask “Are Repetitions allowed?” The clue in this question—one that be very common when we move on to probability—is the phrase “at random”. What this means is that each time a choice is made each of the 227,719,424 social security numbers is equally likely to be picked; in particular, numbers already chosen may be chosen again. So the first column of our spreadsheet is a sequence and the shorthand is $227719424^{18000} \simeq 0.158993449964 \times 10^{150434}$. This number has 150,434 digits and takes up 46 pages—[see for yourself](#)—so I won’t write it down here.

- ii) In how many of these first columns, are all 18,000 social security numbers in the first column different?

Solution



3.8 Counting by the “divide and conquer” method

Now the answer to **R?** is “No” but, by **R IMPLIES O 3.6.6**, the answer to **O?** is still a “Yes” so the shorthand is the permutation $P(227719424, 18000) \simeq 0.780603324655 \cdot 10^{150433}$. This also has over 150,433 digits but I did not put in a file.

iii) In how many of these first columns, is there at least one social security number in two or more different rows?

Solution

This asks for those first columns that are answers to **i)** but not to **ii)** so **AMOANS 3.8.3** tells us to take the difference $227719424^{18000} - P(227719424, 18000) \simeq 0.809331174981 \cdot 10^{150433}$.

iv) Which are more common, those first columns with all social security numbers different or those with at least once social security number duplicated?

Solution

We’re asking which is bigger, the answer to **ii)** or to **iii)**. There both big, but the they have the same number of digits (the same exponent) and the former starts with a 7 while the latter starts with an 8. So there are more with a duplicate somewhere.

Did you recognize this problem? These are the counts that underlie **AIG GIVES BACK: A FAIRY TALE WITH A MORAL**. The chance that *someone* will win two or more million dollar prizes must be more than 50% because the answer to part **iv)** is that columns with duplicates are more common. When we’ve defined equally likely outcome probabilities, we’ll see that that chance is the ratio of the number of first columns with duplicates to the total number of first columns. This gives

$$\frac{227719424^{18000} - P(227719424, 18000)}{227719424^{18000}} \simeq \frac{0.809331174981 \cdot 10^{150433}}{0.158993449964 \times 10^{150434}} \simeq \frac{0.809331174981}{1.58993449964} \simeq 0.50903428737.$$

as was claimed.

Here’s a variant you can calculate that deals with a group of 23 Army recruits all born in 1991. Before working it, would you guess that the

3.8 Counting by the “divide and conquer” method

chance that 2 recruits in such a group share the same birthday is less than 10% or more than 10%?

PROBLEM 3.8.43: At an Army induction ceremony on 2009, the presiding Sergeant has a chart of 23 recruits, all born in the 1991, showing their names (in alphabetical order) and their birthdays:

- i) How many different ways can birthday column of the chart be filled?
- ii) In how many, are all 23 birthdays different?
- iii) In how many, are there at least two recruits with the same birthday?
- iv) Which is more likely, that all 23 recruits will have different birthdays, or that at least one pair will have the same birthday?

Here are a few problems that introduce some common applications of counting orderings.

EXAMPLE 3.8.44: At a graduation ceremony, 9 dignitaries—5 women and 4 men—are to be seated in 9 chairs in the front row on the podium.

- i) How many different seating arrangements are possible?

Solution

We can think of this as making 9 choices—one for each chair—from the set of 9 dignitaries. Repetitions are not allowed (we can’t put the same person in more than one chair) and order matters (we can tell what chair is assigned to which person by, say, counting in from the left) so the shorthand is the permutation $P(9,9) = 362,880$. More directly, what we are doing is choosing an ordering of the 9 dignitaries (from left to right, by seat), and by [ORDERING AND FACTORIALS 3.5.23](#), this can be done in $P(9,9)$ ways.

- ii) How many arrangements are possible with 5 women all seated to the left of the 4 men?

Solution



3.8 Counting by the “divide and conquer” method

The condition divides this into two problems like the preceding one. Seat (or order) the 4 men in the 4 left seats, which can be done in $P(4, 4)$ ways, and then seat the 5 women in the 5 right seats, possible in $P(5, 5)$ ways. Reassembling there are $P(4, 4) \cdot P(5, 5) = 2,880$ ways.

- iii) How many arrangements are possible with 5 women all seated together and the 4 men all seated together?

Solution

This is almost the same as the preceding part, but we now also need to choose which sex to put on the left and which to put on the right. This amounts to ordering 2 sexes, possible in $P(2, 2) = 2$ ways, so there are $P(2, 2) \cdot P(4, 4) \cdot P(5, 5) = 5,760$ seatings.

- iv) How many arrangements are possible with 5 women are all seated together (but the 4 men are not necessarily all be seated together)?

Solution

Once again, what we need to figure out is how many ways there are to allocate blocks of seats to each sex. For each of these, we can fill the men's and women's blocks in $P(4, 4) \cdot P(5, 5)$ ways. Everything is determined once we know how many men are to the left of the block of 5 women? If, for example, there are 3, then the seating pattern is M M M W W W W W M. This number can be between 0 and 4 inclusive, so overall there are $5 \cdot P(4, 4) \cdot P(5, 5) = 17,400$ seatings

- v) How many arrangements are possible with the women and men alternated (no woman is seated next to a woman and no man next to a man)?

Solution

Here there's only a single pattern by sex—W M W M W M W M W—so there are again $P(4, 4) \cdot P(5, 5) = 2,880$ seatings.

PROBLEM 3.8.45: At a graduation ceremony, 9 dignitaries—5 women and 4 men—are to be seated in the front row on the podium. The middle seat is reserved for the University President. Answer the ques-

3.8 Counting by the “divide and conquer” method

tions in [EXAMPLE 3.8.44](#), assuming in addition that the middle seat is reserved for the University President and:

- a. the President is a woman.
- b. the President is a man.

PROBLEM 3.8.46: Suppose you have a small bookshelf that contains your 3 favorite books by each of your 4 favorite authors.

- i) How many ways are there to choose 4 of these books?
- ii) How many ways are there to choose your first through fourth favorite books?
- iii) How many ways are there to choose 1 book by each of the 4 authors?
- iv) How many ways are there to line up the books on the shelf?
- v) How many ways are there if the books by each author are kept together?

PROBLEM 3.8.47: A dining club consisting of 12 (heterosexual) married couples meets each week for an ethnic meal. The meal is cooked by a team of chefs consisting of 3 male and 4 female members.

- i) How many ways are there to choose the team of chefs?
- ii) How many ways are there if no couple is on the team?
- iii) How many ways are there if at least one couple is on the team?
- iv) How many ways are there exactly 2 couples are on the team?

PROBLEM 3.8.48: Your wardrobe contains 4 shirts, 4 pants and 4 jackets, one of each in each of the 4 colors red, white, blue and black. Each day you pick an outfit consisting of a shirt a pair of pants and a jacket.

- i) How many outfits can you choose?
- ii) How many outfits contain garments of only 1 color?
- iii) How many outfits contain garments of exactly 2 different colors?
- iv) How many outfits contain garments of exactly 3 different colors?

Hint: In the latter parts, first choose the colors.

A Classic Example: Poker Rankings

OK, now you’re ready for the classic test of whether you’ve learned how to count: poker. Remember that poker hands consist of 5 cards chosen from a standard deck.

PROBLEM 3.8.49: Rework [PROBLEM 3.6.17.ii](#)) and show that there are 2,598,960 poker hands.

The values are ranked, low to high, as [2, 3, 4, 5, 6, 7, 8, 9, 10, J, Q, K, A]. Adjacent values in this list are called consecutive, and for this purpose only, an A may rank either at the bottom (below a 2) or top (above a K). *consecutive values* in the list [2, 3, 4, 5, 6, 7, 8, 9, 10, J, Q, K, A]. That is an A can be the low or high value in a sequence of *consecutive values* but not a “middle” value.

Certain hands are special and have names to indicate this.

POKER TYPES 3.8.50:

straight The 5 cards have 5 *consecutive values*.

flush All 5 cards belong to the same suit.

straight flush The hand is both a straight (the 5 cards have 5 consecutive values) and a flush (all 5 cards belong to the same suit).

pair The hand contains 2 cards with the same value.

two pair The hand contains 2 sets of 2 cards with the same value.

three-of-a-kind The hand contains 3 cards with the same value.

four-of-a-kind The hand contains 4 cards with the same value.

full house The hand consists of 3 cards of one value and 2 cards of another (that is, contains three-of-a-kind *and* a pair).

high card Any hand not listed above—that is, all 5 cards have different values but these values are not consecutive and the cards do not all lie in the same suit.



3.8 Counting by the “divide and conquer” method

We’ve been a bit imprecise in this list, because, by convention, using one of these types to describe a hand implies that the hand is *not* of any more special type. Thus, calling a hand a straight rules out the possibility that it is a flush: if it were, we’d call it a straight flush. Saying that a hand has a pair rules out the possibility that it has three-of-a-kind.

The rules of poker rank these types (but *not* in the order I have listed them). Any hand of a higher ranked type beats any hand of a lower ranked type. If two hands have the same type, further rules that we won’t go into, are applied: these usually, but not always, break the tie. Counting enters into the discussion because the ranking of types of hands can be very concisely summarized by the rule:

THE LESS COMMON HAND WINS 3.8.51: *If there are fewer poker hands of one type than of a second, then the first type outranks the second.*

CHALLENGE 3.8.52: Use the principle that **THE LESS COMMON HAND WINS 3.8.51** to determine the ranking of types of poker hands.

Basically, all you need to do is to count the hands of each type and compare. I’ll suggest a good order to attack the types and give you some gentle help at a few tricky points along the way.

Let’s start by counting straights. The key thing to note here is that the low value in a straight must be one of [A, 2, 3, 4, 5, 6, 7, 8, 9, 10]. Once I know this value, I know all the others (if low value is 7, the others are 8, 9, 10, and J).

How many ways are there to choose 7, an 8, a 9, a 10, and a J? Hint: This amounts to choosing a suit 5 times. Use this to show that there are 10240 straights.

Unfortunately, that first answer is wrong. Some of those straights are actually straight flushes and shouldn’t be included in our count of “true” straights. How many?



3.8 Counting by the “divide and conquer” method

Show that the low card in a straight flush determines the entire hand. Use this to show that there are 40 straight flushes and 10200 straights.

Flushes are easier. You should have found in [PROBLEM 3.6.24.ii](#)) that there were 5148 poker hands containing 5 cards in the same suit.

Use this to find the number of flushes (which is *not* 5148—why?).

Likewise, You should have found in [PROBLEM 3.6.24.iv](#)) that there were 524 poker hands consisting of 4 cards of one value and a kicker.

Note that if a hand contains 2 (or more) cards of the same value, it cannot be either a straight or a flush. In counting the “of-a-kind” hands, we’ll have to worry quite a bit about not counting “more-of-a-kind” hands too, but this remark tells us that at least we can ignore straights and flushes. For example, we immediately find that the number of four-of-a-kind hands is 524.

Rather than approach the “of-a-kind” hands piecemeal, let me lay out a general strategy.

STRATEGY FOR COUNTING “OF-A-KIND” HANDS 3.8.53: To count “of-a-kind” hands, first pick the values that appear in the hand, then pick the cards of each value.

Let’s try full houses next. Here the [DIVIDE AND CONQUER COUNTING STRATEGY 3.8.1](#) says you need to pick 2 values from 13, one for the three-of-a-kind and one for the pair. What’s the right shorthand for this? This trips up a lot of students so draw some hands to illustrate what the each possible answer to **R?** and **O?** would mean.

Picking the cards of each value is always easy.

PROBLEM 3.8.54: To complete an x -of-a-kind whose value is known, you need to choose x of the 4 cards in that value. What’s the shorthand for this? Hint: if $x = 4$, the number of ways is 1.

PROBLEM 3.8.55: Show that there are 3744 full houses.



3.8 Counting by the “divide and conquer” method

Now we’re almost home. To choose a three-of-a-kind hand, you need to choose 3 values—1 for the three-of-a-kind and 2 for the kickers (unpaired cards). Certainly **R?** is “No” or we’d have a better hand, but **O?** is neither entirely “Yes” nor entirely “No”. Of the hands

J J J 9 7, 9 9 9 J 7 and 9 9 9 7 J

the first is different from the second and third but the last two are the same.

Explain why, in general, swapping the order of 2 or more x -of-kinds with the same x (like the 2 kickers) does *not* matter, but swapping an x -of-kind with a y -of-kind with $x \neq y$ (like a three-of-a-kind and kicker) *does* matter. Illustrate your argument with hands having two pairs and a kicker.

The solution is to pick the values with one common x together (since, for these, order does not matter, we get a combination) and then to pick the values with the next common x *from the remaining values*. Looking back to a full house, we pick 1 value from 13 for the three-of-a-kind and then 1 value from the remaining 12 for the pair, getting $13 \cdot 12$ choices. This is just the value the **PERMUTATION FORMULA 3.5.18** gives for the permutation $P(13, 2)$ we arrived at above. For the three-of-a-kind hands, we pick 1 from 13 for the three-of-a-kind and then 2 from 12 for the kickers.

PROBLEM 3.8.56: Use the guidelines above to show that:

- i) There are 54,912 three-of-a-kind hands.
- ii) There are 123,552 two pair hands.
- iii) There are 1,098,240 one pair hands.
- iv) There are 1,302,540 high card hands. Hint: Here there’s one extra wrinkle because, in this case, we have to choose 5 values for the 5 cards but not 5 consecutive values. Why?

PROBLEM 3.8.57: Every poker hand is one and only one of the types in . Use this and fact that there are 2,598,960 poker hands to check the counts we have derived for all the types



3.8 Counting by the “divide and conquer” method

POKER RANKINGS 3.8.58: The ranking of poker hands by type, from highest to lowest, and the number of each is:

straight flush	40
four-of-a-kind	524
full house	3744
flush	5108
straight	10200
three-of-a-kind	54,912
two pair	123,552
pair	1,098,240
high card	1,302,540

Here’s a final problem you can use to test yourself.

PROBLEM 3.8.59:

- i) List the special hands in a form of poker in which there are only 4 cards in a hand and determine the ranking of the types using the principle that **THE LESS COMMON HAND WINS 3.8.51**.
- ii) (Harder)Do the same for a form of poker in which there are 6 cards in a hand.

The m&m’s problem

In this subsection, I want to guide you to discover how to carry out one kind of count that we cannot apply our **DIVIDE AND CONQUER COUNTING STRATEGY 3.8.1** to. Our goal will be to answer the question:

THE m&m’s PROBLEM 3.8.60: *How many ways are there to color 56 m&m’s using the 6 colors  ,  ,  ,  ,  and *

This problem came up in class in 2002 when **m&m’s** held a contest to select a new color (purple won) and offered a prize of 100,000,000 yen (then about \$700,000) to the buyer of a bag containing only purple **m&m’s**. Discussion about the probability of finding such a bag “at

3.8 Counting by the “divide and conquer” method

random” (so small, it might as well be zero) led to a discussion of how many different bags of **m&m’s** there could be. Based on data on the **m&m’s** website, we expect to find about 56 **m&m’s** in a standard 1.69 ounce package of the milk chocolate variety; this type comes in 6 colors.


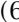

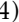




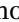

Let’s see why the **DIVIDE AND CONQUER COUNTING STRATEGY 3.8.1** fails here. Saying that we are coloring the **m&m’s** means that we are making 56 choices from 6 possibilities so, by **IF $\ell > m$, IT’S SEQUENCES OR NOTHING 3.6.13**, **R?** is “Yes”. On the other hand, **O?** must be “No” as we don’t change the colors in the bag by shaking it up before opening it. So what we’ve got here is what is dubbed an **abomination** in **THE TWO QUESTION METHOD**—and we don’t have any shorthand for these.

On the other hand, it’s clear that whatever the answer to **THE m&m’s PROBLEM 3.8.60** is, it’s a big number. We’re crazy if we think we can just count the bags on our fingers. If you don’t believe me, feel free to try. I’ll be waiting at the start of the next paragraph when you give up.

OK then. Big counts don’t scare us. We worked out a 46 page long count above. But we did so by plugging into a formula. So what we need is a formula for abominations, like the **COMBINATION FORMULA 3.4.15** or the **PERMUTATION FORMULA 3.5.18**. That is we want a formula for the number of abominations $A(m, \ell)$ —which we’ll think of as the number of ways to color m **m&m’s** (if you’ll forgive the pun) with ℓ colors—for any values of m and ℓ . If we had such a formula, we’d just plug in $\ell = 56$ and $m = 6$ and we’d have our answer.

We have already learned a very important lesson. Often the only way to answer a particular question you’re interested in is to understand how to answer a whole range of similar questions. Here the only way to find the value $A(6, 56)$ that we’re after is to find a formula for $A(m, \ell)$.

3.8 Counting by the “divide and conquer” method

Here’s where we hit a bit of a brick wall. The formulas I cited as models were all built up by making one choice at a time. That won’t work here. Suppose, to the contrary, that we had some choice-by-choice way of building up $A(m, \ell)$. To be definite, say we have arrived at $A(6, 14)$ having chosen 3 each of ,  and , 2 each of  and , and 1 . When we make the next choice, we need to multiply by 6 since we can choose any color. Suppose we choose . Then we no longer have an unordered set of blue **m&m’s**: we have 3 unordered and 4th. That we can handle, as with combinations, if we just divide by 4 to “forget” the position of the last . But if we’d chosen , we’d need to divide by 3 not 4 (the 3rd yellow is now special) and if we’d chosen , we’d need to divide by 2. In other words, the correction we need to divide by depends, not on something universal like the number 14 of choices we have made so far, by on something particular—how many of the first 1 choices were the same color as the 15th. So *no* formula can incorporate the necessary corrections.

Well, maybe we could try reducing m instead of ℓ . That is, we pick a first color—say blue—and try to divide the count in simpler counts, one for each possible number k of blue **m&m’s** in the bag. That works a bit better: we can at least say what these simpler counts are. After coloring k **m&m’s** blue, we’re left with $\ell - k$ **m&m’s** and with $m - 1$ colors. In other words, we’re left with the count $A(m - 1, \ell - k)$. This may, at first, seem reminiscent of say counting the permutation $P(m, \ell)$ where, after making the first choice we are left with the simpler problem of making $\ell - 1$ choices from the $m - 1$ possibilities still not used—that is with $P(m - 1, \ell - 1)$ choices.

But there’s an important difference. With the permutations, every first choice left us with the same simpler problem. With the abominations, every first choice leaves us with a different simpler problem. Instead of looking for the single abomination $A(m, \ell)$, we’re now looking for $\ell + 1$ *different* abominations $A(m - 1, \ell - k)$, one for each value of k between 0 and ℓ . All these simpler problems together are



3.8 Counting by the “divide and conquer” method

more complicated even if the numbers are smaller in each.

However, we have learned something. Let’s see just what and then how it can be used.

m&m’s RELATION 3.8.61: $A(m,\ell) = \sum_{k=0}^{\ell} A(m-1,\ell-k).$

This just says that we can color any number k of **m&m’s** blue, and then we have one fewer color to use on the remaining $\ell - k$ **m&m’s**. But is the formula that says this too messy to be of use? The answer is “Yes”, if the use we want to put it to is to get a formula for $A(m,\ell)$, but it’s “No”, if instead, we only want to build a table of *values* of $A(m,l)$.





Why would we want to do that? Well, our toolbox of ideas for counting doesn’t seem to be much use in computing $A(m,\ell)$. We need a new idea to understand abominations, and our best chance of getting that idea is to get our hands dirty with a bunch of values. This may seem strange, but it’s a basic idea—indeed, a ubiquitous one—in mathematical research. If something seems too complicated to understand, study some simple examples in the hope of seeing a pattern in them. So let’s make a table.

PROBLEM 3.8.62: The following questions ask you to check the values that have been entered in **TABLE 3.8.63**, and then to go on to complete the table.

- i) Explain why the counts $A(1,\ell)$ and $A(m,0)$ shown in blue are all 1.
- ii) Explain why $A(m,1) = m$.
- iii) Use the idea behind the **m&m’s RELATION 3.8.61** to obtain the values in the second column, $A(2,\ell) = \ell + 1$.
- iv) Below are picture of the 6 possible bags containing 2 **m&m’s** in the 3 colors **m**, **m** and **m** in “chromatic” order. Draw similar pictures that verify the counts $A(3,3) = 10$ and $A(3,4) = 15$.



3.8 Counting by the “divide and conquer” method

v) Find the value $A(4, 4)$ in two ways. First draw pictures of possible bags of 4 m&m’s in the colors , ,  and . Then, as a check, apply the m&m’s RELATION 3.8.61: this says that the value is obtained by summing the numbers in the $m = 3$ column from the 1 at the top to the 10 to the left of the $A(4, 4)$ cell.

$\ell \quad m$	1	2	3	4	5	6	7	8	9	10
0	1	1	1	1	1	1	1	1	1	1
1	1	2	3	4	5	6	7	8	9	10
2	1	3	6	10						
3	1	4	10							
4	1	5	15							
5	1	6								
6	1	7								
7	1	8								
8	1	9								
9	1	10								
10	1	11								

TABLE 3.8.63: TABLE OF VALUES OF $A(m, \ell)$

vi) Fill in the values in the blank cells in TABLE 3.8.63. The procedure, based on the m&m’s RELATION 3.8.61, for doing this is indicated in the 3 colored cells. To find the value in a cell, simply total the values that lie opposite or above it in the column immediately to its left. The easy way to fill in the table is to fill one column at a time, working from left to right: you can then find each entry by adding the value in the cell immediately above to the value in the cell immediately to the left.

vii) Look at the completed table. Where have all the numbers in this

3.8 Counting by the “divide and conquer” method

table arisen earlier in the course? What name did we give this table then?

That’s right. Our table of abominations is nothing more or less than a rotated version of Pascal’s triangle (TABLE 3.4.18). The entries in Pascal’s triangle are the combinations. So by computing some values, we have discovered that abominations are combinations! If we can just figure out *which* combination each abomination is, we’ll have a formula for $A(m, l)$.

Comparing the two tables shows that TABLE 3.8.63 is obtained from Pascal’s triangle by rotating it 45 degrees counterclockwise. We can use this observation to match up abominations and combinations.

TABLE 3.8.64 is a version of TABLE 3.8.63 in which is have replaced each value by the corresponding binomial coefficient $\binom{\mathcal{M}}{\mathcal{L}}$ from TABLE 3.4.18. The script letters have been chosen to avoid confusing the abomination and combination indices.

ℓ m	1	2	3	4	5	6	7	8	9	10
0	$\binom{0}{0}$	$\binom{1}{1}$	$\binom{2}{2}$	$\binom{3}{3}$	$\binom{4}{4}$	$\binom{5}{5}$	$\binom{6}{6}$	$\binom{7}{7}$	$\binom{8}{8}$	$\binom{9}{9}$
1	$\binom{1}{0}$	$\binom{2}{1}$	$\binom{3}{2}$	$\binom{4}{3}$	$\binom{5}{4}$	$\binom{6}{5}$	$\binom{7}{6}$	$\binom{8}{7}$	$\binom{9}{8}$	$\binom{10}{9}$
2	$\binom{2}{0}$	$\binom{3}{1}$	$\binom{4}{2}$	$\binom{5}{3}$	$\binom{6}{4}$	$\binom{7}{5}$	$\binom{8}{6}$	$\binom{9}{7}$	$\binom{10}{8}$	$\binom{11}{9}$
3	$\binom{3}{0}$	$\binom{4}{1}$	$\binom{5}{2}$	$\binom{6}{3}$	$\binom{7}{4}$	$\binom{8}{5}$	$\binom{9}{6}$	$\binom{10}{7}$	$\binom{11}{8}$	$\binom{12}{9}$
4	$\binom{4}{0}$	$\binom{5}{1}$	$\binom{6}{2}$	$\binom{7}{3}$	$\binom{8}{4}$	$\binom{9}{5}$	$\binom{10}{6}$	$\binom{11}{7}$	$\binom{12}{8}$	$\binom{13}{9}$
5	$\binom{5}{0}$	$\binom{6}{1}$	$\binom{7}{2}$	$\binom{8}{3}$	$\binom{9}{4}$	$\binom{10}{5}$	$\binom{11}{6}$	$\binom{12}{7}$	$\binom{13}{8}$	$\binom{14}{9}$
6	$\binom{6}{0}$	$\binom{7}{1}$	$\binom{8}{2}$	$\binom{9}{3}$	$\binom{10}{4}$	$\binom{11}{5}$	$\binom{12}{6}$	$\binom{13}{7}$	$\binom{14}{8}$	$\binom{15}{9}$
7	$\binom{7}{0}$	$\binom{8}{1}$	$\binom{9}{2}$	$\binom{10}{3}$	$\binom{11}{4}$	$\binom{12}{5}$	$\binom{13}{6}$	$\binom{14}{7}$	$\binom{15}{8}$	$\binom{16}{9}$

TABLE 3.8.64: WHAT COMBINATION $\binom{\mathcal{M}}{\mathcal{L}}$ GIVES $A(m, \ell)$?

Looking at TABLE 3.8.64, it’s immediately clear that in each column where m is fixed the value of \mathcal{L} is also fixed and that the two are related by $\mathcal{L} = m - 1$. As for the value of \mathcal{M} , it increases by 1 both when we go down a column (that is, when ℓ increases by 1) and when

3.8 Counting by the “divide and conquer” method

we go across a row (that is, when m increases by 1). In other words, \mathcal{M} and $\ell + m$ vary identically, However they do not quite match, comparing them on a few cells we see that $\mathcal{M} = \ell + m - 1$.

When a mathematician has a guess that they’re pretty sure of, they call it a **conjecture**. Our comparison makes us pretty sure that

ABOMINATION CONJECTURE 3.8.65: $A(m, \ell) = \binom{\ell+m-1}{m-1} = C(\ell + m - 1, m - 1)$.

PROBLEM 3.8.66: Use TABLE 3.8.63 and the COMBINATION FORMULA 3.4.15 to check ABOMINATION CONJECTURE 3.8.65 for the following choices of (m, ℓ) :

- i) $(m, \ell) = (3, 2)$.
- ii) $(m, \ell) = (3, 5)$.
- iii) $(m, \ell) = (3, 9)$.
- iv) $(m, \ell) = (4, 3)$.
- v) $(m, \ell) = (4, 6)$.
- vi) $(m, \ell) = (4, 10)$.
- vii) $(m, \ell) = (7, 7)$.
- viii) $(m, \ell) = (9, 8)$.

Looks pretty good. If our guess is right, we can immediately answer the THE m&m’s PROBLEM 3.8.60: the number of to color 56 m&m’s using 6 colors is $A(6, 56) = C(56 + 6 - 1, 6 - 1) = C(61, 5) = 5,949,147$.

Are we done? Well, yes and no. We have a lot of computational evidence that abominations are combinations, we think we can see in the ABOMINATION CONJECTURE 3.8.65 what the pattern is, and we’ve got a lot of computational evidence that this formula does indeed tell us which combination gives which abomination. This kind of evidence is called **inductive** because, like the evidence used in testing our answer to PROBLEM 2.1.21, it’s based on many observations. As in that problem, we’re pretty convinced by it. That’s the “Yes” part of the answer. If I offered to bet you my \$1 against your \$10 that the

3.8 Counting by the “divide and conquer” method

answer to the **THE m&m's PROBLEM 3.8.60** is *not* 5,949,147, you'd say yes in a New York minute, and you'd feel like you found a dollar in the street without even having to bend over to pick it up. Before we discuss the “No” part of the answer, let's look at a couple of problems.

PROBLEM 3.8.67: Recall that a whole number bigger than 1 is called **prime** if it has no positive whole number divisors except 1 and itself. For example, 2, 3, 5, 7 and 11 are prime, but $4 = 2 \cdot 2$, $6 = 2 \cdot 3$, $8 = 2 \cdot 2 \cdot 2$, $9 = 3 \cdot 3$ and $10 = 2 \cdot 5$ are not. To test if a number n is prime is easy, if a bit tedious when the number is large. You just find $b = \frac{n}{a}$ for every whole number a between 2 and $(n - 1)$. If b is *ever* a whole number, then n is not prime and if b *always* has a fractional part, then n is prime. You can speed things in many ways. For example, note that, if $n = a \cdot b$ and $a < b$, then $a \leq \sqrt{n} \leq b$ —if both a and b were bigger than \sqrt{n} then $a \cdot b > \sqrt{n}\sqrt{n} = n$ and likewise if both are smaller, then $a \cdot b < n$. Therefore, if you haven't found a divisor by the time a reaches \sqrt{n} , then you can stop: n is prime. You can also skip any divisors a that are *not* prime themselves; for example, there no point in dividing by 6 because 6 divides n evenly, then so do its factors 2 and 3 and we'd already know that n was not prime.

Consider the conjecture:

PRIME CONJECTURE 3.8.68: For any positive whole number x , $P(x) = x^2 + x + 41$ is prime.

i) Verify the **PRIME CONJECTURE 3.8.68** for x from 1 to 10 by dividing $P(x)$ by all prime numbers less than its square-root.

Partial solution

I'll do $x = 10$ when $P(10) = 10^2 + 10 + 41 = 151$. I check that $\frac{151}{2} = 75\frac{1}{2}$, $\frac{151}{3} = 50\frac{1}{3}$, $\frac{151}{5} = 30\frac{1}{5}$, $\frac{151}{7} = 21\frac{4}{7}$ and $\frac{151}{11} = 13\frac{8}{11}$. Since $13^2 = 169 > 151$, this shows that $P(10) = 151$ is prime as conjectured

ii) Verify the **PRIME CONJECTURE 3.8.68** for x from 11 to 30 by looking up the value $P(x)$ in this [table](#) of prime numbers.



3.8 Counting by the “divide and conquer” method

Partial solution

I'll do $x = 11$. Here $P(11) = 11^2 + 11 + 41 = 173$ and since 173 is the last entry in the 4th row of the [table](#), it's prime.

Once again, we have a *lot* of inductive evidence for the [PRIME CONJECTURE 3.8.68](#). Can we stop?

iii) Check the [PRIME CONJECTURE 3.8.68](#) for $x = 40$. Hint: The square root of 1681 is 41.

So $P(40)$ is *not* prime even though every one of $P(1)$ to $P(39)$ is!

CONVICTION IS NOT KNOWLEDGE 3.8.69: *Inductive evidence may convince us that a statement is true, but we can never know with absolute “impossible that it could be false” certainly that a statement is true unless we have inductively checked every last case.*

Here's another example that makes the same point. You may have heard of [Fermat's Last Theorem](#) which says that you can't write the n^{th} -power of a positive whole number as the sum of two other n^{th} -powers for any $n > 2$ —that is, $x^n + y^n = z^n$ has no positive whole number solutions with $n > 2$. Pierre Fermat claimed, but did not prove, this in 1637. His conjecture was only proved by the English mathematician Andrew Wiles in 1995—357 years later.

But many mathematicians had struggled with the problem throughout those centuries. Leonhard Euler, who we met at the end of [SECTION 1.4](#) as the man who named [THE NUMBER \$e\$ 1.4.47](#) conjectured, in 1769, that no n^{th} -power could be written as a sum of fewer than n other n^{th} -powers. For $n = 3$, this is what Fermat claimed but for $n \geq 4$, Euler's conjecture is an even stronger claim. Euler's conjecture was also unresolved for almost 300 years, and an immense amount of inductive evidence that it was true was amassed over this period. Here's your chance to disprove a conjecture by the great Euler.

PROBLEM 3.8.70: Show that $27^5 + 84^5 + 110^5 + 133^5 = 144^5 = 61,917,364,224$.



3.8 Counting by the “divide and conquer” method

This example was found in 1966 by Leon J. Lander and Thomas R. Parkin and it is known to be the smallest. **Counterexamples**, as mathematicians like to call an example which shows that a statement is at least sometimes false, with 4 4th-powers were found by Noam Elkies in 1988 but the smallest of these (found by Roger Frye) is

$$95800^4 + 217519^4 + 414560^4 = 422481^4 = 31,858,749,840,007,945,920,321.$$

Again the moral is that statements can be true almost all the time and still be fail to be universally true. You may have verified the first million cases of the statement, but there’s no guarantee that it won’t be wrong in the million and first. Or you may have to wait until the billion and first case to find the counterexample.






If we want to really know, in the absolute certainty sense, that a statement is true, we must find arguments that show *why* it is true in *every* case. In down to earth terms, we need an explanation not just a calculation. In fancier language, we need a **deductive** argument or **proof** that applies to every case.

In the **PRIME CONJECTURE 3.8.68**, the ways we checked that values of the polynomial $P(x)$ were prime had nothing to do with $P(x)$. We just took the value and either started doing trial divisions or looked it up in the online table. The fact that any value happened to be prime gave us no clue as to why it was prime, and made it no easier to check the next value.

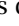

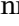
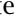
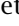
Our evidence for the **ABOMINATION CONJECTURE 3.8.65** is of the same nature. We use the **m&m’s RELATION 3.8.61** to generate a table of values, but nothing about any of those values suggested the kind of **R?** “No” and **O?** “No” choice that we associate with a combination. Looking at the table did make it clear that we were indeed computing combinations, and it was easy enough to identify which ones they were, but again, we have gained no insight into *why* these counts match up. Why, for example, does the m in the abomination turn (after being decremented) into the \mathcal{L} in the combination? We

3.8 Counting by the “divide and conquer” method

have no idea. So all our calculation provides no better guarantee that $A(6, 56) = 5,949,147$ than all our calculations with $P(x)$ offered that $P(40)$ would be prime.


What we need is a way to associate to any of the $C(\ell + m - 1, m - 1)$ combinations arising from making $(m - 1)$ choices from a set of $\ell + m - 1$ possibilities, a bag of ℓ m&m's in m colors. Can we make a picture of the combinations that at least suggests m&m's? To avoid too many complications, let's set $\ell = 10$ and $m = 5$ and fix the colors to be , , ,  and . And let's try using white (i.e. uncolored) m&m's for the set of possibilities and black ones for our actual choices. Here is a typical choice.



Looking at this picture, we see that the placement (i.e. choice) of the 4 black m&m's has left 10 white m&m's that are divided into 5 groups. That's exactly what coloring 10 m&m's using 5 colors would do. This suggests that we use one color for each group. To keep things definite lets color the groups from left to right in the order , , ,  and . What we'll then see is:



PROBLEM 3.8.71: Color the white m&m's corresponding to the combinations below in the same way.



So far so good. Unfortunately, we won't always see 5 groups of white m&m's. How, for example, should we color the following combination?



Take a moment and see if you can convince yourself that there's only one logical way.


3.8 Counting by the “divide and conquer” method

I hope you saw that instead of worrying about groups, we just need to view the black m&m's as *instructions* telling when to stop using one color and start using the next. We start coloring white m&m's red until we reach the first black m&m. It tells us to switch to blue and we color white m&m's blue until we reach the second black m&m which switches us to yellow.

Likewise the third and fourth black m&m's switch us from yellow to green and from green to orange respectively. In coloring the combination above, we're using yellow when we hit those two consecutive black m&m's, the third and fourth. As we cross over these, we switch from yellow to green and from green to orange without coloring any white m&m's. In other words, what we have above is bag that happens to contain 0 *green* m&m's. We really still have 5 groups of white m&m's, it's just that one of them happens to be the empty group. We get:




PROBLEM 3.8.72: Color the white m&m's corresponding to the combinations below which one or more of the 5 groups is empty.



Was there anything special about $\ell = 10$ and $m = 5$?

PROBLEM 3.8.73:

i) Color the white m&m's corresponding to the combinations below with $\ell = 14$ and $m = 6$ using tan as the 6th color.



ii) Color the white m&m's corresponding to the combinations below with $\ell = 16$ and $m = 4$ dropping the color orange.




3.8 Counting by the “divide and conquer” method



OK, we now understand how to go from the combinations with count $C(\ell + m - 1, m - 1)$ to the abominations—bags of **m&m’s**—with count $A(m, l)$. We’re still not done. How do we know we get all the abominations in this way? We need to see that our procedure is reversible, or to use the preferred term, to find its **inverse**. In essence, we’ve got a “logarithm” and we need an “exponential”.

That’s easy. Given a bag of m **m&m’s** in ℓ colors, we just lay out the m candies from left to right, grouping those of each color and putting the colors in the same order used above. Then we insert a black **m&m** after each group and white wash all the colored **m&m’s**. The only point that calls for any care is handling missing colors. If the bag contains no green **m&m’s**, we still have to put down an (empty) green group delimited by a black **m&m**. Here are a couple of examples, showing the ‘raw’ bag of **m&m’s**, the intermediate colored row with black **m&m’s** added, and the final black-and-white row that gives a combination.



Of course, these were the first two examples I did, chosen to make it visible that this inverse procedure really takes an abomination back to the combination it came from.

PROBLEM 3.8.74: Diagram the three stages of the abomination-to-combination procedure for the bags of **m&m’s** constructed in **PROBLEM 3.8.71** and **PROBLEM 3.8.72**.

3.8 Counting by the “divide and conquer” method

What have we achieved? We have paired off each of the $C(\ell + m - 1, m - 1)$ combinations with one (and only one) of the $A(m, l)$ abominations and vice-versa. That's only possible if the number of combination and the number of abominations are the same. So we have deduced **ABOMINATION CONJECTURE 3.8.65**.

And that means that we now *know*—not just are pretty sure but know with absolute certainty—that the answer to the **THE m&m's PROBLEM 3.8.60** is 5,949,147.

Yes, it was a lot of work to get from conviction to knowledge. But knowledge like this, that can be relied on unquestioningly, is never easy to come by—in most situations, impossible. Perhaps the greatest key to the power of mathematics is that it enables us, if we think long enough and work hard enough, to command such knowledge in ways that almost no other discipline can. And it's just as powerful in handling situations that really matter as it is in counting **m&m's**.

Having done all this work, we're ready to record what we've accomplished and I need to come clean and tell you that the standard name for an **abomination** is a **multiset**. The “multi” indicates that “Are Repetitions allowed?” is “Yes” and the set indicates that, as with sets, “Does Order matter?” is “No”. There's also a more standard notation for the count $A(m, l)$, namely $\binom{m}{l}$, which is intended to remind you that this count is just a “rotated” **combination** or **binomial coefficient**. So I'll state our formula as:

MULTISET OR ABOMINATION FORMULA 3.8.75: *The **multiset** or **abomination** count $\binom{m}{l} = A(m, l)$ equals the combination $\binom{\ell + m - 1}{m - 1} = C(\ell + m - 1, m - 1)$.*

I said earlier that *you* won't need to count abominations in this course and I'll stand by that. However, problems which involve counts where repetition is allowed but order does *not* matter do arise fairly often enough in mathematics. The **MULTISET OR ABOMINATION FORMULA 3.8.75** has been known for at least 100 years, and

3.8 Counting by the “divide and conquer” method

probably much longer, but it was only in the late 1960s that the name **multiset** was coined by the Dutch mathematician Nicolas de Bruijn in correspondence with the great American computer scientist Donald Knuth. The notion of a multiset has had many other names (among them bunch, bag and heap but *not* abomination) and while today these are becoming less common they have still not completely disappeared.

Chapter 4

Fat chance

In [CHAPTER 2](#), we took an informal look at some of the many ways that intuition goes wrong when thinking about events with an uncertain or random element. Take a moment to recall the experiments you carried out in playing “[\\$7’LL GET YOU \\$12](#)” and their crazy results, or the paradox of the playoffs in “[WE WUZ ROBBED](#)”. These examples show that we need a careful mathematical way to describe questions involving random events if we hope to draw correct conclusions about and make intelligent decisions involving them. Probability provides this guide to discovering what the right expectations about uncertain phenomena are.

We also saw, for example in looking at [CHUCK-A-LUCK](#), that, even when we think we understand the theory of a random process well and have precise expectations, it can still be very difficult tell if experimental observations support our predictions. If our expectations are correct and we can make a *large enough* number of experiments, our observations will more and more clearly match our expectations. But this is almost never practical, so if we want practical tests of predictions we need to understand how reliably they are confirmed by the observation we are actually able to perform. Statistics allows us



4.1 Experiments, outcomes and sample spaces

to quantify how much confidence about the correctness of our expectations we can derive from any collection of observations.

In [CHAPTER 3](#), we built up the language of sets and strategies for counting that will enable us to address these two needs. We’re going to compute probabilities by counting. Indeed, in many of the exercises in that chapter, we’ve already used these tools to verify claims and predictions from [CHAPTER 2](#). Now we’re ready apply these tools.

Very few new definitions or techniques are needed. Coming fully to grips with probability will not be easy, but the difficulties are in learning how to ask the right questions. Once these are asked, finding the answers usually involves only straightforward counting of the type we’ve already mastered. So let’s cut to the chase.

4.1 Experiments, outcomes and sample spaces

The good news in this section is that we are already very familiar with the basic setup for studying probability. Almost all problems reduce to questions about the number of elements in subsets of a universal set. The bad news is that we don’t use the familiar terms “**universal set**”, “**element**” or “**subset**”. You need to start off by learning the tiny dictionary in [TABLE 4.1.1](#).

SET TERM	PROBABILITY TERM
universal set	sample space
element	outcome
subset	event
disjoint	mutually exclusive

TABLE 4.1.1: SET-PROBABILITY DICTIONARY

4.1 Experiments, outcomes and sample spaces

Why do we introduce 4 new terms for perfectly familiar concepts? It's not because there's *any* tricky difference in meaning between the terms in the left and right columns. They really mean exactly the same things. In one way, that's good: there's nothing new to learn and no worry about keeping subtle differences between the two sets of terms straight. In another way, it's maddening. There's really no need for the right column at all, so why bother with it.

Reason one is a feeble one. The terms **sample space**, **outcome**, **event** and **mutually exclusive** are traditional in probability. So you simply need to know them to read about or discuss probability, in the same way that, however illogical the spellings or pronunciations of many English words may be, you need to know how them to read and converse.

Reasons two and three are better. Elements and subsets come up in a great many contexts, so seeing these terms doesn't give you any context. When you see any of **sample space**, **outcome**, **event** or **mutually exclusive**, you immediately know what's up because these terms are *only* used to discuss probability. This is one big advantage.

The other advantage is that the right column of [TABLE 4.1.1](#) conveys the active spirit in which we think about probability. The key metaphor is that of an **experiment**. Don't think Erlenmeyer flasks and white labcoats here. Doing almost anything and observing what happens counts as a probability experiment. the only requirements are the doing and the observing.

EXAMPLE 4.1.2: Here are some experiments we've already discussed (without referring to them as experiments), and that we'll look at in the sequel. I have colored the **doing** and the **observing** to distinguish them.

- i) **Roll 2 dice** and **observe the numbers that come up on each die**.
- ii) **Toss a coin 3 times** and **record the sequence of Hs and Ts that appear**.

4.1 Experiments, outcomes and sample spaces

- iii) Deal 5 cards from a standard deck and note what poker hand you get.
- iv) Choose a committee of 6 Senators and list its members.
- v) Elect a President Pro Tempore, a Secretary and a Sergeant at Arms from the Senate and record who holds each office

It's common to specify only the experiment that is performed and leave it up to the reader to infer what observation to make. Usually, but not always, there's only one sensible set of observations. For example, the experiment in i) above might come up in study a dice game like craps in we are only interested in the total of the numbers on the two dice. Could we then just observe this total and not the two numbers in the individual dice? Or in ii), we might only care how many heads appear, and not what tosses they occur on. Could we then just observe this number and not the individual tosses? The answer, in both cases, is “No”, but I need to show you a few more ideas before I can explain why. The full story can be found in [OBSERVE SEQUENCES NOT SUMMARIES 4.3.14](#). Until then, I'll always make explicit what observation should be made.

How do we tie in the probability terms with such experiments? Two are very easy. An **outcome** is simply any *one* of the observations that might result from our experiment. The **sample space** is simply the set of *all* such observations or outcomes. We usually denote outcomes by lower case letters like x and y or a , b and c , and denote the sample space by a capital letter, most commonly by S .

EXAMPLE 4.1.3: Here are the sample spaces S for the experiments in [EXAMPLE 4.1.2](#), described—remember a sample space is just a set—by giving the admission test for each, or what's the same, by saying what are the possible outcomes (translation: elements of S).

- i) Ordered pairs of numbers from 1 to 6 (or two letter sequences in the alphabet $\{1, 2, 3, 4, 5, 6\}$). This S is just the set D^2 of [PROBLEM 3.2.8](#).

4.1 Experiments, outcomes and sample spaces

- ii) Three letter sequences in the alphabet $\{H, T\}$. We counted this S —it has $2^3 = 8$ elements—in [PROBLEM 3.6.14](#).
- iii) 5-element subsets of the 52 cards in a standard deck. Of course, we think of this S as the set of poker hands and we counted it in [PROBLEM 3.6.17.ii](#)).
- iv) 6-element subsets of the set of 100 Senators. We counted this S in [EXAMPLE 3.8.23](#).
- v) 3-element lists from the set of 100 Senators. We counted this S in [PROBLEM 3.5.26](#).

An **event** is just a subset of *some* of the outcomes in the sample space. As with subsets we use upper case letters to describe events, but the most common letters used are E (for event), F and so on. It's in dealing with events that the active spirit of probability comes to the fore. Since an event is, like every subset, a set in its own right, you might think we'd specify events by giving their admission tests. In a sense, we do. But that's definitely not how we *think* of events. We almost always describe an event by saying what *happened* when we performed our experiment.

PROBLEM 4.1.4: Here are some events we might be interested in when we consider some of the experiments in [EXAMPLE 4.1.2](#). Use each “what happened” description to tell if each outcome listed is or is not an element of the event (which is, remember, just a subset of the sample space).

- i) When we roll 2 dice:
 - a. the total on the dice is 8: $(5, 4)$, $(6, 2)$.
 - b. the first die comes up odd: $(3, 4)$, $(4, 4)$.
 - c. both dice come up even: $(3, 4)$, $(2, 2)$.
- ii) When we toss 3 coins:
 - a. the first toss is a head: TTH, HHT.
 - b. there are exactly 2 tails: THT, HTH
 - c. there is at least one head: HHH, TTT
- iii) When we pick a poker hand, we get:

4.1 Experiments, outcomes and sample spaces

- a. a full house: $J♥ J♠ J♣ 9♥ 9♦$, $J♥ 9♥ 6♥ 4♥ 3♥$.
b. a straight: $J♥ 10♠ 9♣ 8♥ 7♦$, $J♥ J♠ J♣ 9♥ 9♦$.
c. a flush: $J♥ 10♠ 9♣ 8♥ 7♦$, $J♥ 9♥ 6♥ 4♥ 3♥$.

Once we have described an event, the next thing we'll want to do in almost every case is count it. Doing this involves exactly the kind of counting we spent the last chapter mastering.

PROBLEM 4.1.5: Here are some events we might be interested in when we consider experiments in [EXAMPLE 4.1.2](#). Use each “what happened” description to count the number of outcomes in the subset associated to each event, using the [DIVIDE AND CONQUER COUNTING STRATEGY 3.8.1](#). I'll work a few parts for you because they provide a quick review of most of the main ideas involved.

- i) When we roll two dice:
a. the total on the dice is 5.
b. the first die comes up a 3.
c. the first die comes up odd.

Solution

There are 3 odd possibilities for the number on the first die and then 6 for the number on the second die, so by [AMOUNTS 3.8.3](#) we get $3 \cdot 6 = 18$.

- d. both dice come up odd.

Solution

There are 3 odd possibilities for the number of each of the two dice. Since we want to count rolls with the first odd and then the second odd we multiply to get $3 \cdot 3 = 9$.

- e. exactly 1 of the dice comes up odd.

Solution

Here we must have either the first die odd and then the second die even or else the first die even and then the second die odd. (Why is this or else and not either or both?) Since there are 3 odd and 3 even numbers of each die, [AMOUNTS 3.8.3](#) tells us the count is $3 \cdot 3 + 3 \cdot 3 = 18$.

f. at least 1 of the dice comes up odd.

Solution

We can do this in three ways by giving three different “divide-and-conquer” descriptions of this event.

We can apply the principle [DIVIDE INEQUALITIES USING](#) [orelse 3.8.22](#) to reexpress this event as exactly 1 of the dice comes up odd [orelse](#) exactly 2 (that is, both) come up odd and apply [AMOANS 3.8.3](#) and parts [i\)d](#) and [i\)e](#) to get $9 + 18 = 27$.

We can also say that we want all of the 36 rolls [butnot](#) those with both dice even (I’ll leave you to check that there are 9 of these) and again apply [AMOANS 3.8.3](#) to get $36 - 9 = 27$.

Finally, we can say that we want the event “first die odd” or “second die odd”—there are 18 outcomes in each by applying [i\)c](#). Here we have an [eitherorboth](#) or because the two events intersect in the event “both dice odd” of which there are 9 by [i\)d](#). But now we have all the counts needed to apply [AND-OR FORMULA FOR ORDERS 3.7.15](#) to get $18 + 18 - 9 = 27$.

In the next two parts, assume that the Senate contains 60 Democrats and 40 Republicans.

ii) When we pick Senate officers:

- All 3 are Republican.
- The President Pro Tempore and Secretary are Democrats and the Sergeant at Arms is a Republican.
- The President Pro Tempore is a Democrat.

iii) When we pick the Senate committee:

- all 6 Senators are Democrats.
- exactly 4 Senators are Republican.
- at least 2 Senators are Democrats.

4.2 Probability measures

Our study of probability will be based on what we've learned about sets and counting, but more is involved. Now that we've learned to use the [TABLE 4.1.1](#) to translate from set terminology to probability terminology, we're ready to look at the key extra ingredient. This is the notion of a **probability measure** (also often called a probability distribution) on a sample space S .

Probabilities of Outcomes and Events

Remember the key probability metaphor: we perform an experiment and make an observation. Each possible observation is an **outcome** x of our experiment and the set of all possible outcomes is the **sample space** S . We'll often want to carry out the same experiment many times and record the outcome each time. When we do, we'll call each repetition of the experiment a **trial** (and we'll speak of performing or carrying out a trial, just as we do an experiment).

What a probability measure \Pr amounts to is a numerical *prediction* or expectation, for each outcome x in the sample space S , of *how often we expect to observe* x . More precisely, a probability measure gives us a number $\Pr(x)$ for each outcome x and predicts that, if we perform a *large* number of trials, then the fraction of those trials in which the observed outcome will be x will be *close to* $\Pr(x)$.

The basic idea is thus very simple, as we can see from an example. Consider the experiment: toss a coin and observe which side lands up. Our sample space $S = \{H, T\}$. If the coin is **fair**—by which we mean simply that each face is equally likely to come up—then we expect to see each of the outcomes H and T half the time. The probability measure with values $\Pr(H) = \frac{1}{2}$ and $\Pr(T) = \frac{1}{2}$ encodes this expectation.



4.2 Probability measures

This apparent simplicity is deceptive. The difficulty is hidden in the vagueness of the italicized words *large* and *close to*. Both terms weaken the force of the prediction made by the number $\Pr(x)$, but both are essential. First, the number $\Pr(x)$ makes no sensible prediction about any single trial. If we toss a coin once, we'll see either 1 head or 0, never $\frac{1}{2}$ a head.

Even if we toss the coin many times—that is, make a large number N of trials—we don't expect the fraction of trials in which we see a head to be exactly $\frac{1}{2}$. The number of trials NH in which we'll see a head is a whole number, so if the number of trials N is odd the fraction $\frac{NH}{N}$ can't possibly equal $\frac{1}{2}$. But even if N is even and it's possible for $\frac{NH}{N}$ to equal $\frac{1}{2}$, this possibility won't be very likely. For example, as I mentioned in [He's ON Fire!](#), if $N = 100$, we'll only see exactly $\frac{1}{2} \cdot 100$ or 50 Hs about 8% of the time.

So the most we can say is that in a *large* number of trials, the fraction where we observe a head will be *close to* $\frac{1}{2}$. Even this mealy mouthed prediction is not very clear. Just *how many* trials is large number? Just *how close* to $\frac{1}{2}$ do we expect the fraction of trials in which we observe a head to be? The plain fact is that we get no guidance on either of these questions from our probability measure. So the predictions that a probability measure makes are maddeningly difficult to stick a fork into; that fraction $\frac{1}{2}$ seems so precise until we try to say what it tells us about an actual series of coin tosses.

There is a solution. Statistics tells us how to use probability measures to make precise quantitative predictions that we can test by performing trials. But statistics stands on the shoulders of probability. Before we can profitably study it, we first need learn probability itself much better. While we're doing this, we'll just have to be satisfied to view the predictions made by probability measures as an abstract idealization of what happens when we really perform experiments.

Fortunately, there are a few consequences that we can draw from



4.2 Probability measures

these predictions. Let's make a virtue out of a vice and make the idealistic assumption that $\Pr(x)$ is the fraction of trials in which we observe the outcome x . What can we conclude?

First, if we *never* observe x , that fraction is 0; if we *always* observe x , that fraction is 1. So any $\Pr(x)$ must lie between 0 and 1. Second, in any single trial we will observe exactly 1 of the outcomes in the sample space S . This means that, when we make any number N of trials and sum up the number of times N_x that we observe x for *all* outcomes x in S , we'll get exactly N . Dividing by N , this says that if we sum the fractions $\frac{N_x}{N}$ for all x , we'll get exactly 1. But if, as we idealize, $\Pr(x)$ equals $\frac{N_x}{N}$, then this means that the sum of $\Pr(x)$ and for all x must equal exactly 1. This may seem bit complicated but actually, you knew this all along. When I set out the example of tossing a fair coin, I said what fair meant: the probabilities of heads and tails should be *equal*. But how did we know that the common value P had to be $\frac{1}{2}$? Because we know that the chance of seeing either a head or a tail is 1, so $P + P = 1$ and that forces P to be $\frac{1}{2}$.

It turns out that these are the only restrictions that every probability measure satisfies, so it's worth recording them in a formal definition.

PROBABILITY MEASURE 4.2.1: *By a probability measure on a sample space S , we mean the choice of numbers $\Pr(x)$ for each outcome x in S that satisfy:*

- i) *For every $x \in S$, $0 \leq \Pr(X) \leq 1$.*
- ii) *The sum of $\Pr(X)$ for all the outcomes $x \in S$ is exactly 1:*

$$\sum_{x \in S} \Pr(x) = 1.$$

When we think of $\Pr(x)$ as an idealized fraction of trials in which we expect to observe the outcome x , we can restate these properties.

- i) *The fraction of trials on which we expect to observe x is between 0 (never) and 1 (always).*
- ii) *In every trial, we observe one and only one outcome x in S .*



4.2 Probability measures

Next, we can assign probabilities not just to outcomes (elements of S) but also to events (subsets of S).

PROBABILITY OF AN EVENT 4.2.2: *A probability measure on the outcomes x of a sample space S defines a probability $\Pr(E)$ for any event $E \subset S$ by the rule*

$$\Pr(E) = \sum_{x \in E} \Pr(x).$$

In words, the probability of an event E is just the sum of the probabilities of all the outcomes x that belong to E .

Like the probability of an outcome, we can think of the probability $\Pr(E)$ of an event E as the (idealized) fraction of trials in which we expect to observe an outcome x that belongs to E . Indeed, the fraction of trials when the observed outcome *belongs* to E will just be the sum, over the outcomes (elements) x of E , of the fraction of trials in which the outcome *equals* x , and this is just a restatement of the definition of $\Pr(E)$ above.

Basic Probability Formulae

Next observe that the each term in the sum $\sum_{x \in E} \Pr(x)$ that defines $\Pr(E)$ is a term in the sum $\sum_{x \in S} \Pr(x)$ in [PROBABILITY MEASURE 4.2.1iii](#). What terms in the sum for S are missing from the sum for E ? Exactly those for which the outcome x does *not* lie in E , what's the same, exactly those for which x belongs to the complement E^{cs} . Since we'll almost always have in mind subsets of a fixed sample space S , we'll generally use the simpler notation E^c in which the S is omitted. In other words, because every outcome x in S lies in E or else in E^{cs} , we can split up

$$\sum_{x \in S} \Pr(x) = \sum_{x \in E} \Pr(x) + \sum_{x \in E^c} \Pr(x).$$

We can think of the sum $\sum_{x \in S} \Pr(x)$ on the left hand side of this equation in two different ways. First, we can simply say this sum



equals 1 by [PROBABILITY MEASURE 4.2.1.ii](#)). Since the two terms on the right hand side equal $\Pr(E)$ and $\Pr(E^c)$ respectively, we find that:

COMPLEMENT FORMULA FOR PROBABILITIES 4.2.3: *For any event $E \subset S$, $\Pr(E) + \Pr(E^c) = 1$. Equivalently, $\Pr(E^c) = 1 - \Pr(E)$,*

In words, the probability that E does not happen is 1 minus the probability that E does happen. This way of putting the formula makes it seem obvious that it ought to be true. We can view the fact that it follows from the definition of [PROBABILITY MEASURE 4.2.1](#) as evidence that this definition agrees with our intuition about probability.

The second way of viewing $\sum_{x \in S} \Pr(x)$ is as the probability $\Pr(S)$ of the event S itself—after all, an event is just a subset of S and S is always a subset of itself. This means that we rewrite the [COMPLEMENT FORMULA FOR PROBABILITIES 4.2.3](#) as $\Pr(E^c) = \Pr(S) - \Pr(E)$. In this form, it looks just like the [COMPLEMENT FORMULA FOR ORDERS 3.7.24](#) except that we relate the probabilities of E , E^c and S instead of their orders.

This makes it natural to wonder whether the other formulae involving orders in [SECTION 3.7](#) have probability versions. The answer is yes, and, as above, the formulae are identical except that every order ($\#$) becomes a probability (\Pr). First we have an analogue of the [AND-OR FORMULA FOR ORDERS 3.7.15](#).

AND-OR FORMULA FOR PROBABILITIES 4.2.4: *For any events E and F of S ,*

$$\Pr(E \cap F) + \Pr(E \cup F) = \Pr(E) + \Pr(F) .$$

Exactly the same argument that shows why the [AND-OR FORMULA FOR ORDERS 3.7.15](#) holds also shows why this formula is true. Each of the four probabilities is a sum of terms $\Pr(x)$ for certain outcomes x , namely those in the corresponding event. I claim that for any outcome x the number of terms $\Pr(x)$ that appear on the left and right side are the same. [FIGURE 4.2.5](#) below makes this easy to see.






FIGURE 4.2.5: Picturing the And-Or Formula

Elements x in the top left quarter of the diagram lie in E and in $E \cup F$ but not in F or in $E \cap F$ so they contribute a *single* $\Pr(x)$ to each side. Elements x in the bottom left lie in F and in $E \cup F$ but not in E or in $E \cap F$ so they also contribute a *single* $\Pr(x)$ to each side. Elements in the top right lie in both E and F and hence in both $E \cup F$ and $E \cap F$ so they contribute *two* $\Pr(x)$'s to each side. Elements in the bottom left are in none of the four sets and contribute *no* $\Pr(x)$ to either side. In all cases, the contribution of x to each side is the same as claimed.

There's also a probability form of the [ORELSE FORMULA FOR ORDERS 3.7.19](#) but with one small wrinkle. You may have noticed that we have never mentioned or used the last entry in Set-Probability Dictionary in [TABLE 4.1.1](#). In set theory, when two sets E and F have empty intersection—that is, have no *elements* in common—we say that E and F **disjoint**. In probability, when two events E and F have empty intersection—that is, have no *outcomes* in common—we say that E and F are **mutually exclusive**.

In terms of our active spirit in which we think about probability experiments, the notion “mutually exclusive” is commonly rephrased in several ways. saying that E and F are mutually exclusive means that they cannot both happen (or be observed) in a single trial. If the outcome lies in E , it can't lie in F and vice-versa. If we observed E , the possibility that we might have observed F is *excluded*. In terms of operations, mutually exclusive sets are those for which an `orelse` or disjoint union $E \dot{\cup} F$ is defined. We then have:

4.2 Probability measures

OR ELSE FORMULA FOR PROBABILITIES 4.2.6: *If E and F are mutually exclusive events of S (that is, the subsets E and F of S are disjoint so that $E \cup F$ is defined), then*

$$\Pr(E \cup F) = \Pr(E) + \Pr(F).$$

As with the version for orders, this is just the special case of the **AND-OR FORMULA FOR PROBABILITIES 4.2.4** in which (since $E \cap F = \emptyset$), the term $\Pr(E \cap F) = 0$.

Two warnings are in order here. First, just as with subsets, a very common cause of dope-slapping errors is to treat either or both unions of events that are *not* mutually exclusive as if they were or else or disjoint unions. The **OR ELSE FORMULA FOR PROBABILITIES 4.2.6** is simpler than the **AND-OR FORMULA FOR PROBABILITIES 4.2.4** and, under pressure, it's easy to forget that the simpler form only applies *when* the events are *known* to be mutually exclusive. If you do not *check* whether they are mutually exclusive before applying the **OR ELSE FORMULA FOR PROBABILITIES 4.2.6**, Murphy's law tells us that two events will *not* be mutually exclusive and you'll get the wrong answer.

The second warning is about one of the great unsolved mysteries of contemporary mathematics. In **INDEPENDENCE**, we'll look at the extremely important notion of independent events. Like "mutually exclusive", "independent" is a property that any *pair* of events E and F may or may not have. Other than this the two concepts have *nothing* in common. Yet, students confound these two concepts on tests again and again. So please, repeat after me three times:

"The checks for 'mutually exclusive' and for 'independent' are *completely* different." "The checks for 'mutually exclusive' and for 'independent' are *completely* different." "The checks for 'mutually exclusive' and for 'independent' are *completely* different".

OK. I know some of you will *still* get these mixed up, but at least I've tried.

4.2 Probability measures

Here are a couple of easy problems to let you work with the ideas and formulae. I have worked a few parts of each to get you started.

PROBLEM 4.2.7: Consider the experiment of drawing a card from a standard deck and recording the card you picked. The sample space S for this experiment has 52 outcomes which we'll think of as the cards shown in [FIGURE 3.3.11](#) and a probability measure for this experiment assigns each card x a probability of $\frac{1}{52}$.

i) Verify that the given values of $\Pr(x)$ do meet the conditions for a [PROBABILITY MEASURE 4.2.1](#).

Solution

Since $0 \leq \frac{1}{52} \leq 1$, condition 1).i) holds for each card x . The sum $\sum_{x \in S} \Pr(x)$ in 1).ii) consists of 52 terms each equal to $\frac{1}{52}$ so it equals 1 as required.

ii) Find the probabilities of observing the following events.

- a Jack.
- a heart.
- a black card.
- a black jack
- a red card.
- a black card or a Jack
- a face card (K, Q or J)
- a spot card (one that is not a face card).
- a red face card
- a red card or a face card
- neither a red card nor a face card.

Solution to parts ii)a to ii)f

The basic idea here is that, since all the outcomes have probability $\frac{1}{52}$, when we add up $\Pr(x)$ for the outcomes x in any event E , we'll get $\frac{1}{52}$ exactly $\#E$ times for a total of $\frac{\#E}{52}$. So the probability of getting one of the 4 Jacks is $\frac{4}{52}$, of getting one of the 13 hearts is $\frac{13}{52}$ of getting one of the 26 black cards is $\frac{26}{52}$, and of getting one of the 2 black jacks is $\frac{2}{52}$.

4.2 Probability measures

For events E in ii)e and ii)f, we can either use counting formulae to express $\#E$ in terms of orders we know, or use the corresponding probability formulae to express $\Pr(E)$ in terms of probabilities we know. In the next problem, only the latter approach will work, so to get warmed up, that's what I'll use here. Try to use these ideas in working the other parts.

The events 'red card' is the complement of the event 'black card' so by the [COMPLEMENT FORMULA FOR PROBABILITIES 4.2.3](#), $\Pr(\text{'red card'}) = 1 - \Pr(\text{'black card'}) = 1 - \frac{26}{52} = \frac{26}{52}$.

The events 'a black card' and 'a Jack' intersect as we saw in ii)d: that is they are not mutually exclusive. So we need to use the [AND-OR FORMULA FOR PROBABILITIES 4.2.4](#) to find the probability of the union event 'a black card or a Jack'. We find $\Pr(\text{'a black card'} \cup \text{'a Jack'}) = \Pr(\text{'a black card'}) + \Pr(\text{'a Jack'}) - \Pr(\text{'a black card'} \cap \text{'a Jack'}) = \frac{26}{52} + \frac{4}{52} - \frac{2}{52} = \frac{28}{52}$.

iii) The solution above used the fact that the vents 'black card' and 'red card' were complements of each other. Which other pairs of events in ii) are complements of each other?

iv) Complementary subsets (like 'black card' and 'red card') are always disjoint so complementary events are always mutually exclusive. Which other pairs of events in ii) are mutually exclusive?

In working the previous problem, we started by observing that the probability of any event E was the fraction $\frac{\#E}{52}$. Notice that the denominator of all this fractions is the same—52. This means that we can add or subtract two such probabilities by simply adding or subtracting their numerators—which are *whole* numbers! So even though all the probabilities in the problem are fractions, the only arithmetic you need to perform with them is whole number arithmetic.

Unless, of course, you converted these fractions to decimals, say using a calculator. As soon as you do this, the probability of each event becomes a messy decimal with no apparent relation to any other probability and now you need your calculator add them.

4.2 Probability measures

Take part ii).ii)f. Which calculation is easier? The decimal version $0.500000000 + 0.076923076 - 0.038461538 = 0.538461538$ or the fraction version $\frac{26}{52} + \frac{4}{52} - \frac{2}{52} = \frac{28}{52}$. For the same reason, you should avoid reducing probability fractions to lowest terms because doing so throws away the common denominator.

LEAVE PROBABILITIES AS FRACTIONS 4.2.8: Whole number arithmetic is usually all you need when probabilities are kept as fractions so converting to decimal form makes calculations messier and more error prone. We'll also see later that the numerator and denominator of a probability fraction often contain information that's lost when you convert to decimal form.

PROBLEM 4.2.9: Consider the experiment of drawing a tile from a bag containing a set of 100 English Scrabble™ tiles and recording the letter on the tile.

The sample space S for this experiment consists of the 26 letters of the alphabet plus the blank which I'll indicate by a '_' (but refer to as a letter). To avoid confusion with the letter 'x', I'll denote a typical outcome of this experiment by an italic ℓ . The table below gives a probability measure for this experiment.

ℓ	a	b	c	d	e	f	g	h	i	j	k	l	m	n	o	p	q	r	s	t	u	v	w	x	y	z	_
$\Pr(\ell)$.09	.02	.02	.04	.12	.02	.03	.02	.09	.01	.01	.04	.02	.06	.08	.02	.01	.06	.04	.06	.04	.02	.02	.01	.02	.01	.02

- i) Verify that the given values of $\Pr(\ell)$ do meet the conditions for a PROBABILITY MEASURE 4.2.1.
- ii) Find the probabilities of observing the following events.
 - a. a vowel (consider 'y' to be a consonant here).

Solution

We proceed as in PROBLEM 4.2.7 except that since the outcomes now all have different probabilities we actually have to add up the probability of each vowel. We get $\Pr(a) + \Pr(e) + \Pr(i) + \Pr(o) + \Pr(u) = .09 + .12 + .09 + .08 + .04 = .42$

- b. a vowel or a '_'.

Solution

4.3 Equally Likely Outcomes

The events ‘a vowel’ and ‘_’ have probabilities .42 and .02 and are mutually exclusive so we can apply the **ORELSE FORMULA FOR PROBABILITIES 4.2.6** to get $.42 + .02 = .44$.

c. a consonant.

Solution

We could add up the probabilities of the 21 consonants but it’s easier to observe that this event is the complement of a vowel or a ‘_’ and apply the **COMPLEMENT FORMULA FOR PROBABILITIES 4.2.3** to get $1 - .44 = .56$.

d. a letter contained in the word “mathematics”.

Solution

Once again, we have no alternative to adding up. The only point to note is that we only want to add up the probabilities of the repeated letters ‘m’, ‘a’ and ‘t’ once each. We get $\Pr(m) + \Pr(a) + \Pr(t) + \Pr(h) + \Pr(e) + \Pr(i) + \Pr(c) + \Pr(s) = .02 + .09 + .06 + .02 + .12 + .09 + .02 + .04 = .46$

e. a letter not contained in the word “mathematics”.

f. a vowel or a letter contained in the word “mathematics”.

g. a letter contained in the word “life”.

h. a consonant or a letter contained in the word “life”.

i. a letter neither a consonant nor contained in the word “life”.

j. a letter contained in the name “Barack_Obama”.

k. a letter contained in both “mathematics” and in “Barack_Obama”.

iii) Which pairs of events in ii) are complements of each other?

iv) Which pairs of events in ii) are mutually exclusive?

v) What property should we *not* confuse with “mutually exclusive”?

4.3 Equally Likely Outcomes

Now that we know what probability measures are and their most important properties, we can ask “Where do they come from?” Suppose



4.3 Equally Likely Outcomes

we have an experiment in mind and we know the sample space S and all the outcomes x . How do we come up with the numbers $\Pr(x)$?

We can at least say what won't work. We can't find $\Pr(x)$ by performing a bunch of trials and recording what fraction of the time the observed outcome was x . Sure, such a procedure will give us a fraction that we expect to be close to $\Pr(x)$. But just as $\Pr(x)$ is almost always different, often very different, from this observed fraction, so this fraction will very rarely equal the ideal probability $\Pr(x)$. At best, we can get a rough guess in this way. Worse still, if we performed a second series of trials we'd almost certainly get rather different guesses than those we saw the first time. Finally, to get even crude estimates in this way, you have to perform many trials, or collect a lot of data from trials that others have performed. This is often just not possible, and even when it is, it's a lot of work.

Despite all this, people do attempt to determine probability measures in this way. Because the measures are deduced from experimental trials, they are often called **experimental probability** measures. Other common terms are **empirical probability** or relative frequencies. Often, finding empirical probabilities is the only way to get your hands on a probability measure you want to study. But, in this course, we won't use experimental probability measures very much because we just don't have time for the onerous process of data collection that's required.

Basic Equally Likely Outcomes Formulae

Instead, we will usually restrict ourselves to studying experiments where it's possible to pin down the right probability measure mathematically. To do so, we need to make one key assumption that generalizes the coin toss example above. We assume that all the outcomes in the sample space are equally likely to occur—that is, the probability $\Pr(x)$ is the same number for every outcome x in S . It's



4.3 Equally Likely Outcomes

easy to see that this **equally likely outcomes** assumption determines a unique probability measure. Let's check this and see what consequences it has. Then I'll address the more tricky question of when we can expect this assumption to hold.

The key lies in the equation in **ii)** of the definition of a **PROBABILITY MEASURE 4.2.1**. We are assuming that all the terms $\Pr(x)$ in the sum $\sum_{x \in S} \Pr(x)$ are equal. There are $\#S$ such terms, one for each x in S . So we can rewrite the sum as

$$\underbrace{\Pr(x) + \Pr(x) + \cdots + \Pr(x)}_{\#S \text{ terms}} = \Pr(x) \cdot \#S.$$

But **ii)** says that this sum equals 1, so $\Pr(x) \cdot \#S = 1$ and hence $\Pr(x) = \frac{1}{\#S}$.

As a bonus, the assuming **equally likely outcomes** let's us eliminate the sum from the formula for the **PROBABILITY OF AN EVENT 4.2.2**. We just substitute $\Pr(x) = \frac{1}{\#S}$ and find that $\Pr(E) = \sum_{x \in E} \Pr(x) = \sum_{x \in E} \frac{1}{\#S}$. Now we are just adding up $\#E$ terms (one for each $x \in E$) all of which equal $\frac{1}{\#S}$ so we get the product $\#E \cdot \frac{1}{\#S} = \frac{\#E}{\#S}$. We can also see why this ought to hold by just considering the interpretation of $\Pr(E)$ as the fraction of trials when we expect the observed outcome x to belong to E . Since we expect all outcomes x equally often, the fraction $\Pr(E)$ should just be the fraction of all the outcomes in the sample space S that belong to E . That's just another way of saying $\frac{\#E}{\#S}$.

EQUALLY LIKELY OUTCOMES PROBABILITY MEASURE 4.3.1: *For any sample space S , there is a unique probability measure \Pr with the **equally likely outcomes** property that every outcome x in S has the same probability—that is, is equally likely to be observed in any trial. This probability measure is given by the formula that, for every outcome x ,*

$$\Pr(x) = \frac{1}{\#S}.$$

4.3 Equally Likely Outcomes

EQUALLY LIKELY OUTCOMES FORMULA 4.3.2: *If S is an equally likely outcomes probability space and E is any event in S , then*

$$\Pr(E) = \frac{\#E}{\#S}.$$

Since we'll use these formulae constantly in working with probabilities, I want to pause to make a few comments about doing so. First, this formula is a sort of warp drive for calculating probabilities that makes it practical to deal with experiments with very large sample spaces and events. Imagine for a moment trying to study an experiment with a few thousand outcomes whose probabilities vary from outcome to outcome. Just to describe the probability measure for such an experiment, we'd have to provide a table of values of $\Pr(x)$ like the table of probabilities in [PROBLEM 4.2.9](#); except that instead of having 27 columns, this table would have to have *thousands* of columns. Worse still, most interesting events E would also contain thousand of outcomes x and, to calculate the probability of such an event E , we'd have to *add up* thousands of values $\Pr(x)$ from our table. Neither the table nor the sums are impossible, but I'd need a bodyguard if I assigned such problems as homework.

By contrast, an **equally likely outcomes** probability measure is completely determined by the single value $\frac{1}{\#S}$. More important still, there's no need to do any outcome-by-outcome summing to calculate the probability of an event E . All we need to be able to do is to *count* E to find $\#E$. And that's something we can do very efficiently even if our experiment has a sample space (and events) that contain not thousands, but billions of outcomes.

Moreover, we can often have our cake and eat it too, to some extent. Quite often, when we're collecting data, we have in mind an equally likely outcomes sample space but we have no way to perform experiments that utilize it. A typical example arises when trying to understand opinions or voting patterns. One man (or woman), one vote is just a way of saying that counting votes is modelled by an equally

4.3 Equally Likely Outcomes

likely outcomes sample space S . We can think of opinions as a less formal type of vote. But, to use S we'd need to be able to observe every vote (or opinion) which is usually a practical (or legal) impossibility. What we can do is survey a sample of the population and get estimates for how likely people are to vote, or think, in certain ways. Such surveys can be viewed as estimates of the probabilities of various events of S ("registered Republicans voting for Obama") and we can often use the S we have in mind to analyze these estimates.

We also see, in the next section, that it's often possible to construct sample spaces in which outcomes have differing probabilities out of equally likely outcomes spaces (as in [EXAMPLE 4.4.15](#) and the following examples). Although the arithmetic becomes a bit more complicated, we can understand probabilities in such spaces by using the component equally likely outcomes spaces.

But for now we'll stick to equally likely outcomes sample spaces themselves. In these cases, the upshot is that calculating probabilities comes down to counting, not adding. We can—and should—view the formula $\Pr(E) = \frac{\#E}{\#S}$ as expressing the probability of E as a fraction whose numerator and denominator are both counts. Thus, each of the numerator and denominator *separately* are answers to "how many?" questions and carry useful information. The numerator answers "How many outcomes are in E ?" and the denominator "How many outcomes are in S ?"

This provides one more reason to [LEAVE PROBABILITIES AS FRACTIONS 4.2.8](#). When you convert an equally likely outcomes probability from fraction to decimal form, you throw away the information—the counts—that comprise the numerator and denominator. And, as I already noted, you make *more*, not less work for yourself. Every single probability arising from an equally likely outcomes sample space S has the *same denominator* $\#S$. So you can add or subtract such probabilities (as you'll often need to so) simply by adding or subtracting the *whole number* numerators. For the same reasons, it's



4.3 Equally Likely Outcomes

not smart to cancel common factors from the numerator and denominator. You find the numerator as $\#E$ and the denominator as $\#S$: leave them both just as you found them.

We can apply these ideas to give equally likely outcomes versions of the **BASIC PROBABILITY FORMULAE** of the previous section with the common denominator $\#S$ built in for convenience. You'll use a single formulae to solve for several of its terms, so I have given version for each as a convenience while you are novices. Rather than memorize all these variants, I recommend that you practice using the key formula $\Pr(E) = \frac{\#E}{\#S}$ to get each from the master version in **BASIC PROBABILITY FORMULAE** as you work the problems in the rest of this chapter.

EQUALLY LIKELY OUTCOMES FORMULAE FOR PROBABILITIES 4.3.3:

Let \Pr be the equally likely outcomes probability measure on a sample space S . Then:

i) *If E and F are any events of S ,*

a. $\Pr(E \cup F) = \frac{\#E + \#F - \#(E \cap F)}{\#S}.$

b. $\Pr(E \cap F) = \frac{\#E + \#F - \#(E \cup F)}{\#S}.$

c. $\Pr(E) = \frac{\#(E \cup F) + \#(E \cap F) - \#F}{\#S}.$

ii) *If—and only if— E and F are mutually disjoint events of S ,*

a. $\Pr(E \cup F) = \frac{\#E + \#F}{\#S}.$

b. $\Pr(E) = \frac{\#(E \cup F) - \#F}{\#S}.$

iii) *If E is any event of S , $\Pr(E^c) = \frac{\#S - \#E}{\#S}.$*

Here are a few problems for practice.

PROBLEM 4.3.4: Consider the experiment of rolling a blue die and a red die and observing the number that comes up on each and assume that every possible pair of numbers is equally likely to be observed.



4.3 Equally Likely Outcomes

- i) What is the order $\#S$ of the sample space S ? Hint: Refer to [PROBLEM 3.6.12.ii](#)).
- ii) In [PROBLEM 3.8.41](#), you should have found the following table for the number of rolls that give each total from 2 to 12.

TOTAL	2	3	4	5	6	7	8	9	10	11	12
NUMBER OF ROLLS	1	2	3	4	5	6	5	4	3	2	1

TABLE 4.3.5: COUNTS FOR TOTALS ON 2 DICE

Use this table to find the probabilities of each of the events below. Leave each probability in fraction form.

- A total of 11.
- A total of 11 or 12.
- A total less than or equal to 10. Hint: It's easiest to use the preceding answer.
- A total less than or equal to 4.
- A total of 6 or 12
- A total divisible by 6. Hint: You just computed this.
- A total divisible by 3.
- A total divisible by 2.
- An odd total. Hint: A total is odd exactly when it's not divisible by 2.
- A total divisible by 2 *or* divisible by 3. Hint: A number is divisible by both 2 *and* 3 exactly when it is divisible by 6.

In the next problem, it's possible to count the given events either by giving a shorthand count for the subset or by listing the elements of the subset. The same shorthands will come up in the next problem, but the events will be too big to list, so you'll need to use the shorthands there. To warm up, Try to find the shorthands first in this problem, and then list elements to check your count.

PROBLEM 4.3.6: Consider the experiment of tossing a fair coin 3 times and recording which side comes up on each toss. In [EXAMPLE 4.1.3.ii](#)), we saw that the sample space for this experiment is the set



4.3 Equally Likely Outcomes

of 3-letter sequences in H and T,

$$S = \{ \text{HHH, HHT, HTH, HTT, THH, THT, TTH, TTT} \}$$

and $\#S = 8$. Assume that each 3-letter sequence of Hs and Ts appears equally often.

i) Show that the probability of seeing any given number of heads is given by TABLE 4.3.7. Hint: You may want to refer to EXAMPLE 3.6.15

HEADS	0	1	2	3
PROBABILITY	$\frac{1}{8}$	$\frac{3}{8}$	$\frac{3}{8}$	$\frac{1}{8}$

TABLE 4.3.7: PROBABILITIES FOR HEADS IN 3 COIN TOSSES

ii) What is the probability of each of the following events?

- The first toss is a head.
- At least 1 toss is a head.
- The first toss is a head and at least 2 tosses are heads.
- The first toss is a head and the last toss is a head.
- The first toss is a head or the last toss is a head.

PROBLEM 4.3.8: Consider the experiment of tossing a fair coin 10 times and recording which side comes up on each toss, and assume that each 10-letter sequence of Hs and Ts appears equally often.

- What is the order $\#S$ of the sample space for this experiment?
- Fill in the probabilities of seeing each given number of heads in the table below.

HEADS	0	1	2	3	4	5	6	7	8	9	10
PROBABILITY											

TABLE 4.3.9: PROBABILITIES FOR HEADS IN 10 COIN TOSSES

iii) What is the probability of each of the following events?

- At least 8 tosses are heads.
- At least 2 tosses are heads.
- The number of heads is even.
- The first 3 tosses are heads.
- The last 2 tosses are heads.



4.3 Equally Likely Outcomes

- f. The first 3 tosses are heads and the last 2 tosses are heads.
- g. The first 3 tosses are heads or the last 2 tosses are heads.

Here's another dice problem based on the game chuck-a-luck. The necessary counting formulae are in [PROBLEM 3.8.37](#) and [PROBLEM 3.8.38](#) and some applications are given in [EXAMPLE 3.8.36](#).

PROBLEM 4.3.10: Consider the experiment of rolling 5 dice and observing the number that comes up on each and assume that every possible quintuple of numbers from 1 to 6 is equally likely to be observed.

- i) What is the order $\#S$ of the sample space S ?
- ii) Find the probabilities of each of the events below. Leave each probability in fraction form.
 - a. Exactly 2 dice comes up showing 6.
 - b. No die comes up showing 6.
 - c. At least 1 die comes up showing 6
 - d. Exactly 3 dice come up showing 4
 - e. Exactly 3 dice come up showing 4 and exactly 2 dice comes up showing 6. Hint: If you knew which 2 dice showed a 6, you'd also know which 3 showed a 4. How many choices are there for this subset of 2 dice?
 - f. Exactly 3 dice come up showing 4 or exactly 2 dice come up showing 6.
 - g. Exactly 3 dice come up showing 4 or exactly 3 dice come up showing 6.

Finally, another basically very easy question where we've already done all the counting but where the terminology makes things tricky: I have provided hints to guide you.

PROBLEM 4.3.11: Consider the experiment of drawing 5 cards from a standard deck and recording the resulting poker hand and assume that every poker hand is equally likely to be selected.

- i) What is the order $\#S$ of the sample space S ? Hint: Refer to [PROBLEM 3.6.17.ii](#)).



4.3 Equally Likely Outcomes

ii) **POKER RANKINGS 3.8.58** gives counts of many subsets of S —that is, events E . Use these counts to find the probability of each of the following events. Leave each probability in fraction form.

- a. A straight flush.
- b. A flush.
- c. A hand containing only cards from a single suit. Hint: This includes flushes *and* straight flushes. Remember, in poker, calling a hand a flush implies that it's not a straight flush.
- d. A straight.
- e. A straight or a flush. Hint: This is a dirty question because poker is an exception to the rule that “or includes both”. In poker if I say I have a straight (or flush), I am implying that I do not have a straight flush and hence that I do not have a flush (or straight). In other words, ‘straight’ and ‘flush’ are mutually exclusive events.
- f. A straight or a flush or a straight flush.
- g. A full house
- h. Three-of-a-kind
- i. A pair
- j. A pair or three-of-a-kind Hint: The idea in the previous hint applies. Why?
- k. A pair or three-of-a-kind or a full house.

Setting up Equally Likely Outcomes Sample Spaces

We've seen that *any* sample space S carries a unique **EQUALLY LIKELY OUTCOMES PROBABILITY MEASURE 4.3.1**. Moreover, this measure reduces calculating probabilities of events to counting outcomes and so greatly speeds up and simplifies such calculations.

But, we should also remember that a probability measure is more than just a set of numbers. Each $\Pr(x)$ is prediction for how often we should expect to see the outcome x if we run a large number of trials of our experiment and the motivation for calculating a probability



4.3 Equally Likely Outcomes

$\Pr(E)$ is to understand how often to expect to see the event E in such trials. So garbage in, garbage out applies: if the outcomes x are *not* all equally likely, then the predictions $\Pr(x) = \frac{1}{\#S}$ —and, along with them, the predictions $\Pr(E)$ —will generally be *wrong*.

So we need to understand when the equally likely outcomes assumption really holds. The most honest answer is *never*. Even in the toy example of tossing a single coin, the outcomes heads and tails will never be equally likely. Remember that we called a coin that had this property **fair**. The simple fact is that there is no such thing as an exactly fair coin because marking the two sides differently perturbs the coin's symmetry. For example, if you stand a bunch of pennies on edge on a table, then bang the table, most of the pennies will land heads up because the “head” is incised more deeply.

Even if we could make a fair coin (and we could get pretty close if we were willing to work hard enough), people are incapable of flipping a coin “fairly”. High-speed photography shows that often the coin only wobbles and the same side remains up through an entire toss. Even when people toss more carefully, the coin lands noticeably more often than not with the same side up as was up before it was tossed. Persi Diaconis, a very distinguished Stanford statistician **Diaconis, Persi** actually had a machine built that could toss a coin so that it would come up heads (or, if you like tails) every time and his website contains a (very advanced) paper on the subject.

Let me emphasize that the issue here is not that in a large but finite series of trials we can only hope that each comes up close to half the time. Rather it's that, as we do more and more trials, it will become clearer and clearer that one side comes up more often than the other. We can make this bias small—say, less than 1% or even 0.01%—if we are careful enough about our coin and our tosses, but we can't make it go away completely.

How about a dishonest answer instead? So what if there's no *exactly* fair coin. The equally likely outcomes assumption captures what

4.3 Equally Likely Outcomes

we'd like to think a coin does, it's a reasonably accurate estimate of what coins actually do, and it makes analyzing the resulting events mathematically easy. So we're just going to ignore the slight deviations and pretend that real coins satisfy the ideal of fairness. This mild self-deception is harmless (especially if we acknowledge that we're engaging in it) and is the only practical way to come to grips with the probabilities that arise in tossing coins.

For similar reasons, there's no perfectly fair die on which each number comes up exactly $\frac{1}{6}$ th of the time and no perfectly shuffled deck from which we are equally likely to deal any poker hand but we'll pretend that such ideal dice and decks exist. We'll take our lie even further.

AT RANDOM 4.3.12: By choosing an element at random from a set S , we mean performing an “thought” experiment with sample space S and equally likely outcomes probability measure. In other words, each element of S is equally likely to be chosen in our experiment.

I called the experiment above a thought experiment because we will usually not even *attempt* to specify how to make the choice so that all outcomes are equally likely. Not that this omission really matters, because, as with the coins and dice and decks, such choices are at best only approximately possible in practice.

At this point, we've made life pretty easy for ourselves. All we have to do is take our favorite sample space, utter the magic words “equally likely” or “at random” and we're off to the races. Well, not quite. First, an observation that's useful in counting such choices.

REPETITION AND AT RANDOM CHOICES 4.3.13: When we make several choices at random from the *same* set A of possibilities, the answer to the question “Are Repetitions allowed?” is assumed to be “Yes” by default—that is, unless it is explicitly indicated that repetitions are *not* allowed.



4.3 Equally Likely Outcomes

Thus if we make ℓ choices “at random” from a set A of m possibilities, then, by default, outcomes will be sequences of length ℓ from A , so we’ll find $S = A^\ell$ and $\#S = m^\ell$.

So, experiments where we make several choices at random from a set A are easy to set up. There’s just one pitfall to avoid, alluded to in the discussion following our first [EXAMPLE 4.1.2](#). To answer the questions we’re interested in about such an experiment, we often need only some summary totals derived from the sequence of choices and not the full sequence itself. It’s tempting to observe only these totals and use them as our outcomes, rather than the full sequence, because we get a much smaller sample space. We must avoid this, however, because the equally likely outcomes probability measure on these smaller spaces predict *incorrectly* the frequencies of these summary totals. I’ll first state the rule to follow, then give some examples to explain how to apply it and why it’s needed.

OBSERVE SEQUENCES NOT SUMMARIES 4.3.14: *Suppose an experiment involves picking a random sequence s of length ℓ from a set A . Recall that, by the words **sequence** and **random**, we mean that:*

- i) *We make a sequence of ℓ choices from a fixed set A with repeated choices allowed (and order mattering).*
- ii) *In making every successive choice, each element of A is **equally likely** to be selected, regardless of any preceding choices.*

Then the only sample space on which the equally likely outcomes probability measure correctly predicts how often we’ll see each outcome is the space $S = A^\ell$ in which the outcomes are just the sequences s . In other words, in such experiments, our observation must be the full sequence, and not any summary of it that yields a smaller sample space S' , even if we only want to study events in S'

To explain what sort of alternate S ’s we need to avoid, and why they don’t work, let’s look at the examples in i) and ii) of [EXAMPLE 4.1.2](#).

In i), we roll a red die and a blue die and observe the numbers on each getting a the sample space $S = D^2$ of sequences of length 2



4.3 Equally Likely Outcomes

in the numbers from 1 to 6 which has $\#S = 6^2 = 36$. The equally likely outcomes probability on this space thus says that each roll has probability $\frac{1}{36}$. Let's see why this is the right expectation, say for the roll $(3, 2)$. By [OBSERVE SEQUENCES NOT SUMMARIES 4.3.14.ii](#)), the chance of seeing a 3 on the red die is $\frac{1}{6}$ (all 6 numbers are equally likely) and then the chance of seeing a 2 on the blue die is also $\frac{1}{6}$ so we expect $Pr((3, 2)) = \frac{1}{6} \cdot \frac{1}{6} = \frac{1}{36}$. Of course, the same argument applies to any roll.

This is exactly the sample space we used in [PROBLEM 4.3.4](#). Notice that all the questions we answered there involved only the *total* on the two dice which is a number from 2 to 12. Why not just observe this total—call it t —and use the sample space $T = \{2, 3, \dots, 11, 12\}$? Well, first notice that these totals are *not* equally likely: in fact, none of the 11 totals has probability $\frac{1}{11}$. [TABLE 4.3.5](#) shows their actual probabilities. Fine, but so what if they're not *equally* likely? Once we have [TABLE 4.3.5](#) we know *how* likely each is, and that's all we needed to answer the questions in [PROBLEM 4.3.4](#).

True enough, but there's one fatal flaw. The only way to *find* the probabilities in the table is to view each total t , not as a single outcome in T , but as an *event* or *subset* of outcomes in S , count the rolls s that total to t and apply the [EQUALLY LIKELY OUTCOMES FORMULA 4.3.2](#). To sum up, the right probability measure on T is *not* the equally likely outcomes one, and the only way to find the *right* one is to use the equally likely outcomes probability on the “right” sample space S .

[PROBLEM 4.3.6](#) and [PROBLEM 4.3.8](#) are very similar. Especially in the latter, it's very tempting to hope that we could get by with a sample space containing the 11 total numbers of heads in [TABLE 4.3.9](#) as outcomes, instead of the sample space S of 1024 sequences of length 10 in H and T. But once again, the only way to find the probability in [TABLE 4.3.9](#) is by counting how many sequences contain each number of Hs. There's no way to get started except by using



4.3 Equally Likely Outcomes

the sample space S of full sequences.

And, once again, the equally likely outcomes probability is the right one on S . We expect, for example, to see an H on the first roll $\frac{1}{2}$ the time, a T on the second roll $\frac{1}{2}$ the time, and a T on the third roll $\frac{1}{2}$ the time. So we expect to see the sequence HTT with probability $\frac{1}{2} \cdot \frac{1}{2} \cdot \frac{1}{2} = \frac{1}{8}$ (and likewise for any other sequence of tosses).

You may be wondering why the rule—OBSERVE SEQUENCES NOT SUMMARIES 4.3.14—against observing only partial, or summary information applies only to experiments that consist of choosing sequences. Well, it's true that *whatever* flavor the outcomes to an experiment have, you shouldn't try to condense or summarize your observation of them. However, summarizing is just not very tempting in the other cases: you'll naturally tend to record all the relevant data when an experiment involves choosing a list or subset, so there's no need for any injunction in these cases.

Let me close this section by admitting that I have glossed over one very important issue. It's hidden in OBSERVE SEQUENCES NOT SUMMARIES 4.3.14.II), in the assumption that each choice is an equally likely one, uninfluenced by any choices that may have preceded it. In the examples above, this assumption seems natural (and *does* hold). But, how should we check this less familiar cases? To answer this question, we need to introduce the notions of conditional probability and independence. This is the subject of the next section.

There's one more observation that's helpful in setting up and counting sample spaces and events in them.

AND IN PROBABILITY 4.3.15: *When “and” is used in describing an experiment and its sample space, it's usually andthen@andthen. When “and” is used in describing events, it's usually andalso@anda lso.*

The reason has to do with the active spirit in which we think of probability. When we're describing an experiment and setting its sample space, we talking about what we *did*. Then “and” is usually used extensively to connect a sequence of actions: we do this first andthen

4.3 Equally Likely Outcomes

we do that second. By the [MULTIPLICATION PRINCIPLE 3.7.1](#), this kind of “and” results in a sample space S that is **product** of the sample spaces for the first action and for the second.

When we’re describing an event, we talking about what *happened*. Then “and” is usually used to restrict what *happened* or was observed: we observed this and also we observed that. That is, when we use “and” to describe an event, we are usually expressing it’s set of outcomes as the **intersection** of two larger events. You can take advantage of this using [EQUALLY LIKELY OUTCOMES FORMULAE FOR PROBABILITIES 4.3.3.i](#)), if you already know the order of the union of the two larger events. But more commonly, you need to find the right way of “dividing” the experiment into pieces with so that each of the two events describes a subset of a single piece. You’ll see examples of this in [PROBLEM 4.3.16](#) and [PROBLEM 4.3.17](#) (each time “both” is used there it means “the first and also the second”)

Now for a little more practice in setting up sample spaces. Remember to leave your answers in fraction form.

PROBLEM 4.3.16: Consider the experiment of drawing a card at random from a standard deck, replacing the card and shuffling the deck, and then drawing a second card at random from the deck.

- i) What should the sample space S be for this event and what is $\#S$? Hint: Remember [REPETITION AND AT RANDOM CHOICES 4.3.13](#).
- ii) Find the probability that:
 - a. The 7 of ♥ is picked both times.
 - b. The two cards picked are the same.
 - c. Both cards have the same value (i.e. are a pair).
 - d. Both cards have the same suit.
 - e. Both cards have the same color.
 - f. Both cards are face cards.
 - g. The first card is a face card or the second card is a face card.
 - h. Exactly one card is a face card

4.3 Equally Likely Outcomes

PROBLEM 4.3.17: Consider the experiment of drawing a card at random from a standard deck and then drawing a second card at random from the deck *without* replacing the first card.

- i) What should the sample space S be for this event and what is $\#S$? Hint: In this variant the answer to **R?** is “No”; the usual rule that **R?** is “Yes” given by **REPETITION AND AT RANDOM CHOICES 4.3.13** is overruled by that “without replacing”.
- ii) Find the probability that:
 - a. The 7 of ♥ is picked both times.
 - b. The two cards picked are the same.
 - c. Both cards have the same value (i.e. are a pair).
 - d. Both cards have the same suit.
 - e. Both cards have the same color.
 - f. Both cards are face cards.
 - g. The first card is a face card or the second card is a face card.
 - h. Exactly one card is a face card

Now that we’re getting quite familiar with understanding how to set up the sample spaces for an experiment, we can start using more informal language to ask questions. Indeed, it’s very common to simply indicate what happens in the experiment and then start right in asking about the probability of seeing various events. It’s then up to you to determine what observation is right, describe the corresponding sample space is, and find $\#S$. You’re doomed if you jump right and try to start computing probabilities, because until you know S you don’t know what outcomes to count to get your numerators and until you know $\#S$ you don’t have the right denominator.

FIND S AND $\#S$ FIRST 4.3.18: *In any probability problem, before you start to calculate probabilities, make sure to:*

- i) *Describe the sample space S : that is, determine the appropriate observations and outcomes for the experiment and describe the set of all outcomes.*
- ii) *Find $\#S$: that is, use your description of S to count the outcomes.*

4.3 Equally Likely Outcomes

First, here's a problem that doesn't ask you to write down the sample space, but does start by asking you to count it.

EXAMPLE 4.3.19: Each week during a 10 week season a coach chooses a player at random from a basketball team of 12 players to record the results of the weekly team free throw shooting practice.

- i) How many ways can he choose the 10 recorders?

Solution

We are making 10 choices (one per week) from a set of 12 possibilities (the players). There's nothing to make us overrule the default of [REPETITION AND AT RANDOM CHOICES 4.3.13](#) so the sample space S consists of sequences of 10 players and $\#S = 12^{10} = 61,917,364,224$. Now we are ready to compute probabilities.

- ii) Find the probability that:

- None of the 3 centers is ever picked.
- A center is picked the first week of the season.
- A center is picked the first week of the season but not the last week of the season.
- A center is picked the first 3 weeks of the season.
- No player gets picked as a recorder twice.
- At least one player gets picked twice.

Solution

- Here the outcomes are sequences of 10 choices from the 9 non-centers. There are $9^{10} = 3,486,784,401$ of these so the probability is $\frac{3,486,784,401}{61,917,364,224}$.
- Here there 3 are choices the first week, then 12 for the other 9 weeks. [AMOANS 3.8.3](#) tells us to multiply these so there are $3 \cdot 12^9 = 15,479,341,056$ and a probability of $\frac{15,479,341,056}{61,917,364,224}$.
- Here there 3 are choices the first week, 9 for the last (the 9 non-centers) and 12 for the other 8 weeks. [AMOANS 3.8.3](#) tells us to multiply these so there are $3 \cdot 9 \cdot 12^8 = 11,609,505,792$ and a probability of $\frac{11,609,505,792}{61,917,364,224}$.



4.3 Equally Likely Outcomes

- d. Now there 3 choices the first 3 weeks and 12 for the other 7 weeks, giving $3^3 \cdot 12^7 = 967,458,816$ and a probability of $\frac{967,458,816}{61,917,364,224}$.
- e. Now the answer to **R?** is “No” and we need to count lists of length 10 from a set of 12 elements getting the permutation $P(12, 10) = 239,500,800$ and a probability of $\frac{239,500,800}{61,917,364,224}$.
- f. This event is the complement of the previous one so we can just apply **COMPLEMENT FORMULA FOR PROBABILITIES 4.2.3** or **EQUALLY LIKELY OUTCOMES FORMULAE FOR PROBABILITIES 4.3.3.iii)** to get the probability is $1 - \frac{61,677,863,424}{61,917,364,224}$.

PROBLEM 4.3.20: An English professor assigns each of the 10 students in her seminar to write on one of 5 authors at random.

- How many ways can the authors be assigned?
- Find the probability that:
 - None of the students is assigned to write on Virginia Woolf.
 - Jane and Tom both get assigned to write on Virginia Woolf.
 - Exactly 2 students get assigned to write on Virginia Woolf. Hint: How many choices are there for the pair of students? You may find it useful to refer to **PROBLEM 4.3.6** or **PROBLEM 4.3.8**.
 - At most 2 students get assigned to write on Virginia Woolf.

Here are typical problems that will trip you up if you don't pay attention to **FIND S AND #S FIRST 4.3.18** because they start right out by asking you for some probabilities. Remember, always start by determining what outcomes you'll observe and counting the sample space they form. Only then are you ready to start finding any probabilities you are asked for.

PROBLEM 4.3.21: A computer science quiz contains 15 True/False questions. If you select your answer to each question at random, what is the chance that:

- You'll get no questions right.
- You'll get the first question right and the last question right.
- You'll get exactly 8 questions right.

4.3 Equally Likely Outcomes

iv) You'll get at least 13 questions right.

PROBLEM 4.3.22: In Manhattan Avenues run north-south and number up from east to west and streets run east-west and number up from south to north. You walk from 2nd Avenue and 4th Street to 9th Avenue and 8th Street, going only west (up the avenues) and north (up the streets). If you select your route at random, what is the probability that you'll,

- i) pass through the intersection of 9th Avenue and 4th Street?
- ii) pass through the intersection of 6th Avenue and 6th Street?
- iii) pass through the intersection of 9th Avenue and 4th Street and through the intersection of 6th Avenue and 6th Street? Hint: Don't count, draw a picture.
- iv) reach 9th Avenue before you reach 8th Street?
- v) walk at least one block along 5th Street?

Hint: All routes are 11 blocks long and go west 7 blocks and north 4 blocks. Knowing which 4 of the 11 blocks you were walking north tells you which 7 you were walking east, and hence what your route was.

First steps in the minefield

To show how far we have already come in understanding probabilities, I'll close this section by setting up carefully the sample spaces that arise in two of the examples from [SECTION 2.1](#), and then verifying the informal probabilities that were calculated there. In fact, you'll be able to do all the necessary calculations: I'll just ask the questions.

PROBLEM 4.3.23: First, we'll deal with [LIGHTNING STRIKES TWICE](#). We'll start with a warmup, then move to the real thing.

- i) Suppose we roll a red die and a blue die and observe the number that comes up on each.



4.3 Equally Likely Outcomes

- a. What is the chance that both numbers that come up will equal 4?
 - b. What is the chance that the 2 numbers that come up will be equal to each other?
- ii) Suppose we observe the winning number in the New Hampshire pick4 (remember this is just any 4-digit number between 0000 and 9999) and, on the same day, observe the winning number in the Massachusetts pick4.
- a. What is the chance that both winning numbers will equal 8092?
 - b. What is the chance that the 2 winning numbers will be equal to each other?
- iii) What is the chance of that the number 8056 will come up tomorrow in both drawing of the New York win4?
- iv) What is the chance of the “phenomenal 100 million to one” event that the same number (from 0000 to 9999) will come up tomorrow in both drawing of the New York win4?
- v) What is the chance of the “one in a million” event that the same number (from 000 to 999) comes up two days in a row in the Nebraska pick3?

Next, we'll turn to the fictitious example [AIG GIVES BACK: A FAIRY TALE WITH A MORAL](#). We did all the calculating in [EXAMPLE 3.8.42](#). You can refer to that problem for the numbers you'll need here.

PROBLEM 4.3.24: Consider the basic trial of picking at random one of 227,719,424 Americans citizens and recording his or her unique social security number. For this trial we can think of the sample space as the set A of all 227,719,424 social security numbers. Now consider the experiment of repeating this basic trial 18,000 times, each time picking one of the 227,719,424 Americans at random, and observing all 18,000 social security numbers.

- i) What is the sample space S for this experiment and what is $\#S$?
- ii) What is the probability that all 18,000 social security numbers observed are different?



4.3 Equally Likely Outcomes

iii) What is the probability that least one American is selected more than once?

Here's a final problem that ties up these two examples. About a year and a half after the Massachusetts pick4 (originally called the State Lottery) began operating in April, 1976, the Lottery's public relations director was asked by a reporter whether any winning number had come up more than once. His response, off the top of his head, was, "Of course, not" because there were 10,000 possible number and only 500 numbers had been picked.

PROBLEM 4.3.25: What is the probability that the public relations director was right, if each day's number were truly drawn at random? If he *was* right, what conclusion would you come to about the drawings?

Some graduate students at M.I.T. reached the same conclusion: either there *were* repeated numbers, or the Lottery was being fixed. There was enough of an outcry that the State Lottery Commissioner was forced to write a letter of correction to the Boston Globe, pointing out that had been several repeated winning numbers.

To sum up these examples. if we make a sequence of random choices from a large set of possibilities, the chance of seeing a previously specified number (say, 8092) on any choice is small. Seeing a given previously specified twice in a row is *extremely* unlikely. But the chance of seeing the same (but unspecified) number twice in a row is *not* tiny: it's exactly the same as the chance of seeing a specified number once. And, after a fairly small number of choices have been made, it's almost certain that you'll have seen some number come up more than once.

Or, more graphically. If there are a lot of "places", it's unlikely for lightning to strike in any previously chosen place at all. For lightning to strike twice *in a row* in any previously chosen spot is *extremely* unlikely. But for lightning to strike twice in a row in some unspecified

4.4 Conditional Probability

place is again only unlikely. And if we watch a lot of lightning strikes, we're almost sure to see lightning strike twice somewhere.

4.4 Conditional Probability

We now come the spot where many of the bodies are buried, the topic of **conditional probability**. Why does this notion cause so much difficulty for so many students? Not because it involves any complicated formulas or tricky counting. Like ordinary probabilities, conditional probabilities turn out to be fractions whose numerator and denominator each count some set of outcomes.

No, the difficulties are in understanding what *question* is being asked. Mastering conditional probabilities comes down to learning to read questions carefully and detect the clues that tell you what kind of probability is being discussed in each sentence. I'll give you a very cut-and-dried method for finding these clues. Once you've located them, finding the answer is usually a piece of cake.

Basic Conditional Probabilities

What we are trying to understand about an experiment with conditional probabilities is very simple. The ordinary probability of an event E tells us how often to expect to observe an outcome in E if we know only the measure \Pr and *no* other information. Conditional probabilities for E tell us how these expectations should be adjusted in the light of *partial* information about the outcome. We provide this partial information by specifying a second event F in which we assume that the observed outcome *does* lie.

In other words, a conditional probability predicts the fraction of outcomes that will lie in E , *assuming* that it's known that the outcome



4.4 Conditional Probability

already lies in F . As we'll soon see this probability may or may not equal, or be close to the probability of E itself. What we're generally interested in understanding, is just how the assumption that F happened affects (or doesn't) the chance that E will happen.

A typical question where we'd like such information is in understanding **prevalence** and risk factors for diseases. I'll stick to some examples dealing with cancer rates and based on a recent [SEER report](#) of the National Cancer Institute. In 2006, there were a total of 11,384,892 cases on cancer in the U.S. of which 6,216,003 were in men and 5,168,889 in women. So the probability that a randomly chosen cancer victim is male is about 55% and female about 45%. Is this true of specific types of cancers? To answer this question, we can assume that the cancer is of a particular tissue, and then ask for the conditional probability that the victim is male or female.

At one extreme, we have cancers in gender specific tissues like the prostate and ovaries, where all cases are, by definition, in one sex. All the 2,177,975 cases of prostate cancer are in men and all the 176,007 cases of ovarian cancer are in women. So if we assume the cancer is of the prostate, then the chance that the victim is female is 0 (or 0%) and, if we assume that the cancer is of the ovary, then the chance that the victim is female is 1 (or 100%). Other cancers like breast cancer occur almost, but not quite, always in one sex: of 2,546,325 cases of breast cancer only 13,132 (or about 0.5%) were in men. Here again, we attribute the difference to gender specific biology, especially hormonal differences.

Other cancers appear to strike men and women with roughly equal frequency. For example, there were 536,944 cases of colorectal cancer, 367,925 of melanoma and 171,522 of lung cancer in men versus 567,158, 390,763 and 193,474 in women. So, for example, the chance that a victim of melanoma is male is $\frac{367,925}{858,688}$, or just over 48%. Finally, there are cancers that are substantially more frequent in one sex than the other, but where it's not clear whether this is due to



4.4 Conditional Probability

biological or behavioral differences. Lung cancer used to be in this category 50 years ago, attacking men far more often than women, because smoking was much more prevalent amongst men than women in the first half of the 20th century. What we see today is the result of the increase in smoking amongst women from 1935-1965. In fact, because male smoking rates peaked (at 67% !) at the end of World War II, lung cancer rates in men have been declining for 15 years, while those for women are only peaking now.

In these examples, the SEER report handed me the counts that formed the numerator and denominator of my conditional probability. How do we find them for ourselves? [FIGURE 4.4.1](#) provides a geometric picture of what's involved that leads to a simple formula. Here $\Pr(E)$ measures the fraction of trials when the outcome lies




FIGURE 4.4.1: Picturing the Conditional Probability $\Pr(E|F)$

in the upper E rectangle over the total fraction of 1 in the large S square. On the right, we assume that the observed outcome was in the F rectangle, so we can ignore outcomes outside this rectangle. Now to say that the outcome lies in E means that it's in the upper square, where the E and F rectangles overlap or intersect, which is just $E \cap F$. So, on the right, the conditional probability we are after is the fraction of trials when the outcome lies in the upper $E \cap F$ square over fraction when it lies in the F -rectangle.

In both cases, we are taking the ratio of the upper half over the whole. But in the conditional probability we forget everything outside the right F rectangle because we *assume* the outcome lies there.

4.4 Conditional Probability

All we need is to introduce the notation for conditional probability and we can translate this picture into a formula. Fix a sample space S with a probability measure \Pr and two events E and F of S .

CONDITIONAL PROBABILITY FORMULA 4.4.2: *We write $\Pr(E|F)$ —read, the conditional probability of E given F —for the probability of observing an outcome in E assuming or given that the observed outcome does lie in F . Then,*

- i) *For any probability measure \Pr , $\Pr(E|F) = \frac{\Pr(E \cap F)}{\Pr(F)}$.*
- ii) *If \Pr is the equally likely outcomes probability measure on S , then*
$$\Pr(E|F) = \frac{\#(E \cap F)}{\#F}.$$

The first formula is just the translation of the picture “blue square” over “red rectangle”. It’s the really only really key formula. The formula in ii) is an easy consequence. Using the **EQUALLY LIKELY OUTCOMES FORMULA 4.3.2**, $\Pr(E \cap F) = \frac{\#(E \cap F)}{\#S}$ and $\Pr(F) = \frac{\#F}{\#S}$. Multiplying both sides by $\#S$, these tell us that

$$\frac{\#(E \cap F)}{\#F} = \frac{\#S \cdot \Pr(E \cap F)}{\#S \cdot \Pr(F)} = \frac{\Pr(E \cap F)}{\Pr(F)} \text{ after cancelling.}$$

This version is most convenient for calculating in the usual equally likely outcomes case, but we’ll use the first quite a bit because it’s the one that reminds us of **FIGURE 4.4.1** and the idea behind conditional probability.

If we clear the denominator in the **CONDITIONAL PROBABILITY FORMULA 4.4.2** we get a restatement so useful it deserves its own name.

INTERSECTION PROBABILITY FORMULA 4.4.3: *For any probability measure \Pr , $\Pr(E \cap F) = \Pr(F) \cdot \Pr(E|F)$.*

This formula is intuitively clear when re-expressed in the active spirit of probability. What has to happen for the an outcome in the event $E \cap F$ to be observed? By definition, both E and F have to happen. The factor $\Pr(F)$ is the chance of the second of these two events occurring. Of outcomes where F happens, what fraction also lie in E ? This is exactly what we mean by the conditional probability $\Pr(E|F)$,

4.4 Conditional Probability

depicted on the right of [FIGURE 4.4.1](#). That is the chance that E and F both happen is the product of the chance that first F happens and that then that, given this, E happens too.

Of course, we could also think of first asking for E to happen and then for F to happen, given that E is known to occur: this gives $\Pr(E \cap F) = \Pr(E) \cdot \Pr(F|E)$. Warning: you need to keep your E s and F s straight here. Neither $\Pr(E) \cdot \Pr(E|F)$ nor $\Pr(F) \cdot \Pr(F|E)$ bears any relation to $\Pr(E \cap F)$. See, for example, [ORDER MATTERS FOR CONDITIONAL PROBABILITY 4.4.4](#).

This formula is called the *Multiplication Formula* in many books but we won't use this term. A special name is not needed and this one gets confused with the [MULTIPLICATION PRINCIPLE 3.7.1](#). Not surprisingly, this variant is used to compute probabilities of intersection events. We'll look at some examples at the end of this subsection.

Up to this point, when we have considered probabilities that involve more than one event like $\Pr(E \cup F)$ or $\Pr(E \cap F)$, the order in which we took E and F did not affect the probability—for the simple reason that union and intersection are symmetric operations to the events do not depend on the order: $E \cup F = F \cup E$ and $E \cap F = F \cap E$. Not so for conditional probability.

ORDER MATTERS FOR CONDITIONAL PROBABILITY 4.4.4: Always pay careful attention to the order of the events E and F in a conditional probability because the two probabilities $\Pr(E|F)$ and $\Pr(F|E)$ are almost never equal.

In finding $\Pr(E|F)$, we assume F is observed and ask what is the chance of observing E . In finding $\Pr(F|E)$, it's E that we assume observed and F that we ask about. And ask what is the chance of observing E . The fractions that give these probabilities always have the same numerator— $\Pr(E \cap F)$ which does not depend on the order of E and F —but the denominators—the probability of the assumed event—are different because F is assumed in the first case and E in the second.



4.4 Conditional Probability

Let's try a few examples to see that using the [CONDITIONAL PROBABILITY FORMULA 4.4.2](#) is easy and confirm that [ORDER MATTERS FOR CONDITIONAL PROBABILITY 4.4.4](#). I have solve a few parts to get you started but the plan is very simple. Note that I always start by describing the sample space S and counting it, in keeping with our rule [FIND \$S\$ AND \$\#S\$ FIRST 4.3.18](#). Strictly speaking, we'll never need to know $\#S$, as it does not appear in the the [CONDITIONAL PROBABILITY FORMULA 4.4.2](#), but we do need to understand what outcomes to count the events that do appear in this formula and the best way to make sure of this is to start with S itself. Next, we count each of the events E , F and $E \cap F$. Then, we simply plug into [CONDITIONAL PROBABILITY FORMULA 4.4.2.ii\)](#) to find either $\Pr(E|F)$ or $\Pr(F|E)$.

PROBLEM 4.4.5: We roll blue and red dice and record the number on each as usual. For each pair E and F of events, find the two conditional probabilities $\Pr(E|F)$ and $\Pr(F|E)$.

- i) E = “total of 9”, F = “blue die comes up 3”.

Solution

As usual, the sample space S consists of the 36 sequences of length 2 in the numbers 1 to 6. Here $\#E = 4$ (see [PROBLEM 4.3.4](#) for a table of these counts) and $\#F = 6$ because we can have any number on the red die. If $E \cap F$ occurs then we have a total of 9 with a 3 on the the blue die, so the red die must come up 6. Hence $\#(E \cap F) = 1$. So $\Pr(E|F) = \frac{\#(E \cap F)}{\#F} = \frac{1}{6}$ and $\Pr(F|E) = \frac{\#(E \cap F)}{\#E} = \frac{1}{4}$. Note that the two are not equal.

- ii) E = “total of 11”, F = “blue die comes up 4”. (Here the two conditional probabilities do turn out to be equal)

- iii) E = “total of 6”, F = “blue die comes up 5”.

PROBLEM 4.4.6: We toss 4 coins and record the sequence of heads and tails as usual. For each pair E and F of events, find the two conditional probabilities $\Pr(E|F)$ and $\Pr(F|E)$.

- i) E = “three heads”, F = “first toss is a head”.

Solution



4.4 Conditional Probability

The sample space S consists of the $2^4 = 16$ sequences of length 4 in the letters H and T. As in [PROBLEM 4.3.6](#) and [PROBLEM 4.3.8](#), the number of outcomes with exactly ℓ heads is the combination $C(4, \ell)$, so $\#E = C(4, 3) = 4$. Sequences in F consist of an H followed by an sequence of length 3 so $\#F = 2^3 = 8$. Sequences in $E \cap F$ consist of an H followed by a sequence of length 3 with exactly 2 more H's so $\#(E \cap F) = C(3, 2) = 3$. So $\Pr(E|F) = \frac{\#(E \cap F)}{\#F} = \frac{3}{8}$ and $\Pr(F|E) = \frac{\#(E \cap F)}{\#E} = \frac{3}{4}$. Note again that the two are not equal.

ii) E = “last toss is a head”, F = “first toss is a head”.

iii) E = “at least 3 heads”, F = “first toss is a head”.

PROBLEM 4.4.7: We select a poker hand at random. For each pair E and F of events, find the two conditional probabilities $\Pr(E|F)$ and $\Pr(F|E)$.

i) E = “straight”, F = “hand contains 5♥, 6♦, 7♥ and 8♣”.

Solution

The sample space S consists of the 5 card subsets of a 52 card deck with order $C(52, 5) = 2,598,960$, as in [PROBLEM 3.6.17](#). Likewise, we can read off $\#E = 10,200$ from the count of straights in [POKER RANKINGS 3.8.58](#). The hands in F consist of the 4 given cards, plus one of the other 48 so $\#F = 48$. When will this card give a straight? When it's a 4 or a 9 so $\#(E \cap F) = 8$. So $\Pr(E|F) = \frac{\#(E \cap F)}{\#F} = \frac{8}{10,200} = \frac{1}{1275}$ and $\Pr(F|E) = \frac{\#(E \cap F)}{\#E} = \frac{8}{48} = \frac{1}{6}$.

Here the two conditional probabilities are not just different, they are very different. We can see that this is to be expected if we go back and think about what the conditional probability means. The value $\Pr(E|F) = \frac{1}{1275}$ is the answer to the question, “How often will a straight contain the cards 5♥, 6♦, 7♥ and 8♣?": not often, because of the more than ten thousand of straights, only 8 have these cards. The value $\Pr(F|E) = \frac{1}{6}$ is the answer to the question, “If I'm given the cards 5♥, 6♦, 7♥ and 8♣ and draw a fifth, how often will I make a straight?": quite often, because 8 those 48 fifth cards do the job.

4.4 Conditional Probability

ii) E = “straight or straight flush”, F = “hand has a 5, a 6, a 7 and an 8”.

Solution

Once again, we can read off $\#E = 10,200 + 40 = 10,240$ from [POKER RANKINGS 3.8.58](#). To determine a hand in F , we now need to choose 4 times (once for each card) to determine the 4 suits, and then pick one of the other 48 cards. Since in picking the suits **R?** is “Yes” we have $4^4 = 256$ possibilities so $\#F = 256 \cdot 48 = 12,288$.

Some of the hands in F are now flushes, which wasn’t possible in i), and to make it easy to count $E \cap F$, I added the straight flushes to E . This means that, once again, the hands in $E \cap F$ are those in F that have values in sequence. This happens whatever the suits of the first 4 cards as long as the fifth card is a 4 or a 9 so $\#(E \cap F) = 256 \cdot 8 = 2048$. So $\Pr(E|F) = \frac{\#(E \cap F)}{\#F} = \frac{2048}{12,288} = \frac{1}{6}$ and $\Pr(F|E) = \frac{\#(E \cap F)}{\#E} = \frac{2048}{10240} = \frac{1}{5}$.

The value $\Pr(E|F)$ is unchanged because given a hand in F what we need to get one in E is the same: a 4 or 9 are our fifth card. But the value $\Pr(F|E)$ is much larger because now F contains all straights whose low card is a 4 or 5—since these are 2 of the 10 possible low values for a straight, we expect, and get, probability $\frac{2}{10} = \frac{1}{5}$.

iii) E = “straight or straight flush”, F = “hand has a 5, a 6, a 8 and an 9”.

iv) E = “full house”, F = “hand has at least three 5s”.

v) E = “full house”, F = “hand has at least two 7s”.

Let’s take stock. A few of the problems above called for some thought, but the thought was needed to count events and had nothing to do with the fact we were going to use these counts to find a conditional probability. Using the [CONDITIONAL PROBABILITY FORMULA 4.4.2](#) itself was always just plug and chug.

Distinguishing Conditional Probabilities

OK. Now we are ready to tackle the hard part of working with conditional probabilities: understanding what question is being asked. Let's start with some examples, more than a few, to get a feel for the difficulty.

EXAMPLE 4.4.8: Consider an experiment that involves picking a student in your class at random and recording his or her age, sex and major. We'll think of the sample space S for this experiment as the set of students in your class, but we know how to ask questions about these three attributes, so we can consider the events E = "21 or older", F = "male", G = "psychology major" and H = "biology major".

What probability—expressed in terms of the events E , F , G and H —should we calculate to answer each of the following questions about students chosen at random in your class? The answers involve not only these events, but their unions, intersections, complements and conditional probabilities. See how many you can write down on your own before taking a look at the answer key.

- i) What's the chance the student is 21 or older?
- ii) What's the chance the student is female?
- iii) What's the probability the student is 21 or older or is male?
- iv) What fraction of the male students are 21 or older?
- v) What's the probability of picking a male aged 21 or older?
- vi) How many students 21 or older are male?
- vii) What's the chance the student is majoring in biology or psychology?
- viii) Given that the student is a psychology major, what's the chance she's a female?
- ix) How many students are male psychology majors?
- x) If a student is female, what's the chance she's a psychology major?



4.4 Conditional Probability

- xi) Assuming he's male, what's the chance he's a psychology major?
- xii) What fraction are psychology majors?
- xiii) What percentage of biology majors are 21 or older?
- xiv) How many students are 21 or older and biology majors?
- xv) How many biology majors are males 21 or older?
- xvi) How many male biology majors are 21 or older?

Answer key

- i) $\Pr(E)$
- ii) F^c
- iii) $\Pr(E \cup F)$
- iv) $\Pr(E|F)$
- v) $\Pr(F \cap E)$
- vi) $\Pr(F|E)$
- vii) $\Pr(G \cup H)$
- viii) $\Pr(F^c|G)$
- ix) $\Pr(F \cap G)$
- x) $\Pr(G|F^c)$
- xi) $\Pr(G|F)$
- xii) $\Pr(G)$
- xiii) $\Pr(E|H)$
- xiv) $\Pr(E \cap H)$
- xv) $\Pr((E \cap F)|H)$
- xvi) $\Pr(E|(F \cap H))$

If you didn't get a lot of these right, don't worry. Most students feel that they're swimming in event stew when first faced with such a range of questions. But, in a moment, I'll outline an easy method for telling which chunk in the bowl each such question is pointing at.

First, a few comments about the nature of the task. Most of these questions involve more than one event. The few that don't, like [i\)](#) and [xii\)](#), are easy. Sometimes the individual events are complemented as in [ii\)](#): again, as long as we stay alert and remember that [ANTONYMS](#)



4.4 Conditional Probability

DESCRIBE COMPLEMENTS 3.7.29 and that “female” is a way to say “not male”, or than “under 21” means “not 21 or over” these are straightforward.

But most of these questions involve two of the events. When they do the difficulty is not in identifying *what* the two events are. For example, all of **iii)**, **iv)**, **v)** and **vi)** involve the age and sex of the student so we know we’re dealing with E and F . What a bit tricky about these is deciding *how* the two events are used to build the question. Ignoring for the moment complements, we now have 3 basic ways to construct a question from E and F . By taking a union, by taking an intersection, or by taking a conditional probability. This last really amounts to 2 possibilities itself. Because **ORDER MATTERS FOR CONDITIONAL PROBABILITY 4.4.4**, swapping the event that is assumed to have happened and the one whose probability is being asked about changes the question. The upshot is that we not only need to find ways to distinguish union, intersection and conditional probability, but we need to find a way to tell which is the given or assumed event, in a conditional probability question.

Union questions are almost always easy to detect. The word “or” is a giveaway, and there’s no way to ask a union question without using it. We just have to distinguish the connecting “or” that we see in **iii)** or **vii)** from the “or” in “21 or older” that is internal to the description of E . I chose this way of describing E to make this point: I could have described E as “over 20” to avoid such confusion.

Likewise, when we see the word “and” linking two events, we’re pretty sure we’re dealing with an andalso intersection. Remember though that this applies only *inside* a fixed sample space or universal set: when we’re describing or count a sample space S , it’s more likely to be an andthen “and” telling us to take the product of two component sets of choices of get S . Still, in questions like those in **EXAMPLE 4.4.8**, an “and” pretty much tells us we have an intersection.



4.4 Conditional Probability

The first difficulty we need to confront is that we can also ask about intersections *without* using “and”. We just have to concatenate the events we want to intersect as in [v](#)). We’re clearly asking for students who are both “male” and “21 or older”—we could rephrase the question as, “What’s the probability of picking a student who is male *and* 21 or older?”—but, as it stands, there’s no “and” to tip us off.

There are giveaway words for conditional probabilities too. Remember the idea is to assume that one event—it’s usual to call this the **given event**—happened. So if we see something like “assuming that” or “given that”, as in [xi](#)) or [viii](#)), we know we’re dealing with a conditional probability in which the event that is assumed or given is the **given**.

Another giveaway word is the conjunction “**if**” that introduces a conditional subordinate clause, and this is where conditional probability got its name. The **if** precedes a property that we are to mentally assume in the question. So when we see an **if** clause in a question, as in [viii](#)), we know that it describes the given event in a conditional probability.

The second difficulty that faces us is that we can also ask for conditional probabilities without using any of the giveaway words “if” “given” or “assume” (or any of their synonyms like “suppose that”). Let’s look again at [xiii](#)): “What percentage of biology majors are 21 or older?”. This is not an intersection question. We could ask about $E \cap H$ without using an “and” but we’d say, “What percentage are biology majors 21 or older?” What’s more, we *know* something about these students—they’re biology majors. So this is a conditional probability question in which the given event is H (“biology majors”) and we want to know how many of this given are in E (“21 or over”).

Let’s look now at the trickiest parts of [EXAMPLE 4.4.8](#) to identify. Questions [iv](#)), [v](#)) and [vi](#)) *all* involve both “male” and “21 or older”—that is both E and F . None of these questions uses any of the intersection or conditional probability giveaway words, and all three are

4.4 Conditional Probability

asking for different probabilities. You can see from the Answer Key that **v)** asks about the intersection $\Pr(E \cap F)$, while **iv)** asks about the conditional probability $\Pr(E|F)$ in which F is given, and **vi)** asks about $\Pr(F|E)$ in which E is given. Remember that because **ORDER MATTERS FOR CONDITIONAL PROBABILITY 4.4.4**, these are two *different* conditional probabilities. How do we tell what each of these questions is asking?

Here then are the problems that confront us. How do we distinguish the intersection from the two conditional probabilities? And how do we tell which conditional is which? In other words, how do we tell which is the given event and which the event whose probability we are asking about? There two good tests. One is based on meaning and the other on structure. Remember that a conditional probability question is one in which we assume one event occurred and restrict our attention to it, then ask about the probability that a second event occurred. The next rule simply tells you to use this characteristic “division of labor” between the two events to recognize conditional probabilities.

GIVENS ARE KNOWN 4.4.9: *If a probability question involves two events and does not contain any giveaway words, ask: “Does the question restrict our attention to outcomes in one event and ask about the other, or does it ask about both events?”*

- i) If the question is asking about both, then it’s seeking an intersection probability.*
- ii) If it assumes one and asks about the other, then the question is asking about a conditional probability and the assumed event is the given one.*

Let’s see how to apply this rule to **iv)**, **v)** and **vi)**. The question in **iv)**—“What fraction of the male students are 21 or older?”—considers male students and asks about their age: our attention is restricted to the event F and we’re asked about E . So it’s seeking the conditional probability $\Pr(E|F)$. The question in **v)**—“What’s the probabil-



4.4 Conditional Probability

ity of picking a male aged 21 or older?”—asks about *both* gender and age but assumes neither. So it’s seeking the intersection probability $\Pr(E \cap F)$. The question in vi)—“How many students 21 or older are male?”—restricts our attention to students in E (who are 21 or older) and asks about F (what fraction are male). So it’s seeking a conditional probability in which E is the given, $\Pr(F|E)$.

PROBLEM 4.4.10: GIVENS ARE KNOWN 4.4.9 applies even if there are giveaway words, though it’s not needed in such cases. Apply this rule to verify the answer to parts of vii)–xi) of EXAMPLE 4.4.8.

On occasion, you may find yourself confused about the intended meaning of a question and unable to decide for sure how to apply GIVENS ARE KNOWN 4.4.9. In such situations, there’s a fallback rule—GIVENS ARE IN THE SUBJECT 4.4.11—that let’s you use the grammatical structure of the question to look for the same conditional probability “division of labor”.

GIVENS ARE IN THE SUBJECT 4.4.11: *If a probability question involves two events and does not contain any giveaway words, ask: “Is the reference to one of the events contained in the subject of the question?” You can usually recognize these without a lot of parsing because you can delete the reference from the question and still have a grammatically complete sentence.*

- i) *If neither event is in the subject, the question is asking about an intersection probability.*
- ii) *If one event is in the subject, then the question is asking about a conditional probability and the event in the subject is the given one.*

This is a lot more complicated to say than it is to do. To warm up, let’s try it first on xiii), “What percentage of biology majors are 21 or older?”, which refers to E = “21 or older” and H = “biology majors”. The reference to E is the object of the main clause so we learn nothing. But the reference “of biology majors” to H is in the subject. So GIVENS ARE IN THE SUBJECT 4.4.11 tells us that we are dealing with the conditional probability $\Pr(E|H)$ in which H is the given event.



4.4 Conditional Probability

Next, let's try it on [iv](#)), [v](#)) and [vi](#)) which refer to $E = \text{"21 or older"}$ and $F = \text{"male"}$. In [iv](#)), the phrase "of the male students" referring to F is in the subject so the rule tells us that we are dealing with the conditional probability $\Pr(E|F)$ in which F is the given event. In [v](#)), both events occur as parts of the object, so here we are being asked about the intersection probability $\Pr(E \cap F)$. Finally, in [vi](#)), the reference to E occurs as a modifier to students, the subject, ("students 21 or older") so here we have the conditional probability $\Pr(F|E)$ in which E is the given event. Note also that in deleting "of the male student" in [iv](#)) or "21 or older" in [vi](#)) leaves a complete sentence, while deleting "male" or "21 or older" in [v](#)) does not.

PROBLEM 4.4.12: Like [GIVENS ARE KNOWN 4.4.9](#), the rule [GIVENS ARE IN THE SUBJECT 4.4.11](#) applies even if there are giveaway words, though it's not often needed in such cases. Apply this rule to verify the answer to parts of [vi](#))-[xi](#)) of [EXAMPLE 4.4.8](#).

Do these rules make every possible question completely cut and dried? No, because we're often interested in questions that involve more than 2 events, like [xv](#))-[xvi](#)) above. But we can use giveaway words, and the ideas that [GIVENS ARE KNOWN 4.4.9](#) and [GIVENS ARE IN THE SUBJECT 4.4.11](#) to break down these more complicated events.

PROBLEM 4.4.13: Here we'll take apart [xv](#))-[xvi](#)) in steps. To begin with, we'll focus on finding a given, if there is one, and not worry about connecting to the events E - H .

- i) First use [GIVENS ARE KNOWN 4.4.9](#) to determine what part of the question describes what is known or assumed about the outcomes (if anything) and what part describes what we are asking about them.
- ii) Now use [GIVENS ARE IN THE SUBJECT 4.4.11](#) to check your answers by finding which events (if any) occur in the subject and which in the object of the question.
- iii) Use your answer to say whether the question involves a conditional probability or not. Then state the probability being sought in the words used in the question.

4.4 Conditional Probability

Solution to xvi)

- i) Here the questions tells to restrict our attention to “male biology majors” and then asks us about those who are “21 or older”.
- ii) The events that occur in the subject are “male and biology major” and those that occur in the the object are “21 or older”.
- iii) This asks for a conditional probability with “male and biology major” as the given event and “21 or older” as the queried event. So we want $\Pr(\text{“21 or older”} | \text{“male and biology major”})$.

The next phase is to express the given subject event and queried object event in terms of $E-H$ using the basic operations of union, intersection and complement. For each of [xv](#))-[xvi](#)):

- i) Express the given/subject event event in terms of $E-H$ using the basic operations of union, intersection and complement.
- ii) Do the same for queried/object event.
- iii) Combine these expressions with your answers above to determine what probability is being asked about in each part in terms of $E-H$.

Solution to xvi)

- i) The given event is “male and biology major” and this is $F \cap H$.
- ii) The object event is “21 or older” which is just E .
- iii) We were looking for $\Pr(\text{“21 or older”} | \text{“male and biology major”})$ or $\Pr(E | (F \cap H))$

Probability Measures for Compound Experiments

We next want to look at some problems that use the variant in part [INTERSECTION PROBABILITY FORMULA 4.4.3](#): $\Pr(E \cap F) = \Pr(F) \cdot \Pr(E|F)$. This formula provides far the easiest way to find intersection probabilities, when, as is often the case, we know, the necessary conditional probabilities. Using the formula in this way is straightforward, as you’ll see.



4.4 Conditional Probability

What's a bit surprising, and calls for a bit more thought, are some of the applications of this formula. It tells us how to find the appropriate probability measure for experiments in which outcomes are *not* equally likely. These arise—quite commonly—when we consider **compound experiments** that:

- i) The compound experiment involves a succession of 2 (or more) component experiments.
- ii) Each component experiment has an equally likely outcomes probability measure.
- iii) The outcome of the first experiment affects *what second experiment we perform*.

In such compound experiments, the equally likely outcomes probability is seldom the right one, and **INTERSECTION PROBABILITY FORMULA 4.4.3** tells us what probability measure is appropriate.

I'll start with the most classic sources of examples, urns and balls. The setup was inspired by Athenian legal system of the 5th century. Each juror had a “guilty” and an “innocent” bronze disc: his vote went into a bronze urn and the other was discarded into the wooden one (to maintain voting secrecy). Versions of the system have been widely used since. In the 18th century and 19th centuries, membership in gentlemen's clubs was often decided by dropping small balls colored either white or black into a box. A white ball was a vote to elect, a black one to reject. Such election were usually by veto rather than plurality: a single black ball caused the rejection of the applicant, and, to this day, when a veto is exercised we speak of black-balling.

The balls were named **ballota** (Italian for small ball) and the box was known as a **ballot box**. In the example above, used by the **Association of the Oldest Inhabitants of the District of Columbia**, members picked up either a black or white ball from the tray at the rear and deposited it into the small chute from which it fell into the draw-



FIGURE 4.4.14: A 19th century ballot box

ers below to be counted. The names have survived to the present although the original procedure is now only used by Freemasons.

For our purposes, an urn is any container holding a number of objects that we'll call balls with a narrow opening that lets us choose balls without seeing them. We imagine that—for example, by shaking the urn—it is possible to randomize the contents so that every ball is equally likely to be chosen. We allow the ball to be visually different—we will usually suppose that they are labeled and colored—so that we partially or totally identify what ball has been selected. This “urn and ball” model has been studied by many of the great creators of probability or 300 years. Jakob Bernoulli (whom we met in [BERNOULLI'S LIMIT FOR \$\ln\$ 1.4.42](#)) considered such problems in 1713, and in 1795 the great Pierre-Simon Laplace uses them dozens of times in his expository [Essai philosophique sur les probabilités](#).

EXAMPLE 4.4.15: Suppose we are given a red urn containing 8 white balls labelled $RW1$ through $RW8$ and 8 black balls labelled $RB1$ through $RB8$, and a green urn containing 2 white balls labeled $GW1$ and $GW2$ and 2 black balls labelled $GB1$ and $GB2$. Our experiment consists of picking an urn at random and then picking a ball from the chosen urn at random and recording the label of the ball we selected. Remembering [FIND \$S\$ AND \$\#S\$ FIRST 4.3.18](#), we ask:

- What's the sample space S for this experiment?
- What's the appropriate probability measure on S ?

4.4 Conditional Probability

The first question is easy: S consists of the 20 labels on the balls in the two urns.

But there's a twist. Although we picked an urn at random and a ball at random from that urn, the balls are *not* equally likely to be picked. We could arrange this, but we'd need to perform a *different* experiment: pour *all* the balls from both urns into a yellow urn and choose a ball at random from that urn. What difference does introducing the extra step of picking one of the urns make? Well, it introduces a conditional probability. Once we understand how, we can use the [EQUALLY LIKELY OUTCOMES FORMULA 4.3.2](#) to unwind the *unequally* likely outcomes in our experiment. The outcome will be that the balls in the red urn each have probability $\frac{1}{32}$ of being picked while those in the green each have probability $\frac{1}{8}$ of being picked!

Let's focus on ball $RW4$ and let's use R and G for the events of picking the red urn and green urn. Since we picked an urn at random, each of R and G has probability $\frac{1}{2}$ of being chosen. Suppose that we chose the red urn. When we pick a ball at random from this urn, each of the 16 balls is equally likely to be picked. So we know that, if we picked the red urn, the chance of seeing ball $RW4$ is $\frac{1}{16}$. This tells us exactly that the conditional probability $\Pr(RW4|R) = \frac{1}{16}$. But now [INTERSECTION PROBABILITY FORMULA 4.4.3](#) tells us that

$$\Pr(RW4 \cap R) = \Pr(R) \cdot \Pr(RW4|R) = \frac{1}{2} \cdot \frac{1}{16} = \frac{1}{32}.$$

But the outcome $RW4$ is contained in the event R (the ball $RW4$ is in the red urn), so $\Pr(RW4) = \Pr(RW4 \cap R) = \frac{1}{32}$. The same argument applies to any ball in the red urn by simply changing the label.

PROBLEM 4.4.16: Argue exactly as above that:

- i) $\Pr(G) = \frac{1}{2}$
- ii) $\Pr(GB2|G) = \frac{1}{4}$
- iii) $\Pr(GB2 \cap G) = \frac{1}{8}$ Hint: Apply [INTERSECTION PROBABILITY FORMULA 4.4.3](#).

Conclude that the probability of observing any ball in the green urn is $\frac{1}{8}$.



4.4 Conditional Probability

We can check these probabilities by noting that the probabilities of all the balls do indeed total $16 \cdot \frac{1}{32} + 4 \cdot \frac{1}{8} = 1$. In hindsight, it's also easy to see why we should expect the balls in the green urn to have 4 times as great a chance of being selected. In total, the balls in each of the red and green urns have probability $\frac{1}{2}$ of being selected. But there are 4 times as many balls in the red urn so each is only $\frac{1}{4}$ th as likely to be chosen.

A depressingly good parallel to this undemocratic example is provided by the **United States Senate**: if we think of Senate seats as urns and voters as balls, then the probabilities in the example represent the relative importance or weight of votes in different states. If you live in a urn that contains a lot of balls like California or Texas, your vote means a lot less than if you live in an urn with very few balls like Vermont or Wyoming. Voters in Wyoming count more than 69 times as much as those in California!

EXAMPLE 4.4.17: In the experiment of **EXAMPLE 4.4.15**, let W denote the event of choosing a white ball and B that of choosing a black ball. Find each of the probabilities below:

- i) $\Pr(R \cap W)$
- ii) $\Pr(R \cap B)$
- iii) $\Pr(G \cap W)$
- iv) $\Pr(G \cap B)$

Solution

Now that we understand the probability measure on S , we can do this by summing probabilities of outcomes. For example, $R \cap W$ consists of the 8 outcomes $RW1$ to $RW8$, each of which has probability $\frac{1}{32}$ so $\Pr(R \cap W) = 8 \cdot \frac{1}{32} = \frac{1}{4}$. Each of the other probabilities also turns out to equal $\frac{1}{4}$ by the same argument. This approach, however, forces us to work with outcomes in the compound experiment where outcomes are not equally likely.

However, there's a second approach, again using **INTERSECTION PROBABILITY FORMULA 4.4.3**, that let's us think in terms of

4.4 Conditional Probability

events in component experiments which remember each *do* have equally likely outcomes. For example, the formula says that $\Pr(R \cap W) = \Pr(R) \cdot \Pr(W|R)$. We know the chance of picking the red urn at random is $\frac{1}{2}$. Likewise the chance of picking a white ball from the red urn at random is $\frac{8}{16}$ —8 of the 16 balls are white, and we just apply the **EQUALLY LIKELY OUTCOMES FORMULA 4.3.2**. So $\Pr(R \cap W) = \frac{1}{2} \cdot \frac{8}{16} = \frac{1}{4}$. Similarly, we get $\Pr(G \cap B) = \Pr(G) \cdot \Pr(B|G) = \frac{1}{2} \cdot \frac{2}{4} = \frac{1}{4}$. I'll leave the other two to you.

We'd now like to go one stage further and repeat the procedure we've just carried out without referring to labels on the balls. We'll simply pretend that we can distinguish the balls in some way without worrying about how. We'll work out the right probability measure, just for practice, but then we'll see that we can compute the probabilities of the most salient events without referring to this measure explicitly.

EXAMPLE 4.4.18: Suppose now the red urn contains 10 white and 2 black balls and the green urn contains 2 white and 10 black balls.

- What's the sample space S for this new experiment?
- What's the appropriate probability measure on S ?
- What's the probability of choosing the green urn and a black ball?
- What's the probability of choosing the red urn and a black ball?
- What's the probability of choosing a black ball?
- If we chose a black ball, what's the chance we chose the green urn?

Solution

- Once again, our sample space is just the set of 24 balls in the two urns.
- Suppose x is any of the 12 balls in the red urn. Then (replacing RW4 above by x ,

$$\Pr(x) = \Pr(x \cap R) = \Pr(R) \cdot \Pr(x|R) = \frac{1}{2} \cdot \frac{1}{12} = \frac{1}{24}.$$



4.4 Conditional Probability

Likewise, each ball in the green urn has a chance of $\frac{1}{24}$ of being picked. This checks with the total of 24 balls in the two urns.

iii) $\Pr(G \cap B) = \Pr(G) \cdot \Pr(B|G) = \frac{1}{2} \cdot \frac{10}{12} = \frac{10}{24}.$

iv) $\Pr(R \cap B) = \Pr(R) \cdot \Pr(B|R) = \frac{1}{2} \cdot \frac{2}{12} = \frac{2}{24}.$

v) Here we need to recognize that a black ball either comes from the green urn or the red but never both so $B = (B \cap G) \dot{\cup} (B \cap R).$

Applying **EQUALLY LIKELY OUTCOMES FORMULAE FOR PROBABILITIES**

4.3.3.ii), $\Pr(B) = \Pr(B \cap G) + \Pr(B \cap R) = \frac{10}{24} + \frac{2}{24} = \frac{10}{24}.$

vi) This question is easy as long as you don't let it make you queasy. Yes, we are assuming something about the second stage of our experiment (choosing a ball) and asking something about the first (choosing an urn), and the question seems to travel back in time. But it makes perfect sense, and better still, we only have to plug in values we already have to find this conditional probability. We want to know $\Pr(G|B)$ —we know the ball is black, and wonder what color was the urn—so we just plug in

$$\Pr(G|B) = \frac{\Pr(G \cap B)}{\Pr(B)} = \frac{\frac{10}{24}}{\frac{12}{24}} = \frac{10}{12}.$$

A few comments about this example. I set this one up intentionally so that all the outcomes were equally likely—each of the 24 balls has a chance of $\frac{1}{24}$ of being picked. So we can check all the answers we got above by counting. For example to check **iii)–v)**, just note that 10 of the 24 balls were black balls in the green urn, 2 were black balls in the red urn, and 12 were black. We can even check **vi)** by noting that 10 of the 12 black balls were in the green urn. Warning: in the next problem, and in most problems of this type, only the method of of the solution will work, because the outcomes will not be equally likely. For example, although there are 2 black balls in each urn, the answer to the question, “If we chose a black ball, what’s the chance we chose the green urn?” will *not* be $\frac{2}{4}$!

PROBLEM 4.4.19: Suppose now the red urn contains 10 white and 2 black balls and the green urn contains 4 white and 2 black balls.



4.5 Organizing Related Probabilities

- i) What's the sample space S for this new experiment?
- ii) What's the appropriate probability measure on S ?
- iii) What's the probability of choosing the green urn and a black ball?
- iv) What's the probability of choosing the red urn and a black ball?
- v) What's the probability of choosing a black ball?
- vi) If we chose a black ball, what's the chance we chose the green urn?

4.5 Organizing Related Probabilities

As the examples above demonstrate, we'll often be interested in relating a lot of different probabilities determined by a few basic events, by taking intersections, complements and conditionals. In this section, I want to discuss the two most convenient ways to organize such collections of related probabilities and display them together efficiently. These are tables and tree diagrams. In each case, we mainly need to understand how to read and write down such displays—which mainly involves keeping straight where to look for or set down probabilities of different types. At the end, we'll see that tables and trees are just different ways to display the same kinds of collections of information, but I'll indicate what kinds of problems fit better with each and we'll work some examples.

Working with Probabilities in Tables

Let's start with the basic case, in which we have just two events, E and F . Our display, in this case, will contain the following kinds of probabilities:

PROBABILITIES FROM TWO EVENTS 4.5.1:



4.5 Organizing Related Probabilities

i) **Simple probabilities** that involve just a single event. These are the probabilities of E and F and of their complements E^c and F^c . As a shorthand, I'll speak of an E -event to mean either E or E^c and an F -event, either F or F^c below.

When we combine a pair of events, we'll always choose one E -event and one F -event. The reason is simple: E and E^c are mutually exclusive, so intersection and conditional probabilities involving this pair all automatically 0.

ii) **Intersection probabilities** that involve the intersection of an E -event and an F -event. These are the probabilities $\Pr(E \cap F)$, $\Pr(E \cap F^c)$, $\Pr(E^c \cap F)$, and $\Pr(E^c \cap F^c)$.

iii) **Conditional probabilities** that involve the probability of an E -event given an F -event, or vice-versa. These are now eight of these, because **ORDER MATTERS FOR CONDITIONAL PROBABILITY 4.4.4**: first $\Pr(E|F)$, $\Pr(E|F^c)$, $\Pr(E^c|F)$, and $\Pr(E^c|F^c)$, and then, $\Pr(F|E)$, $\Pr(F|E^c)$, $\Pr(F^c|E)$, and $\Pr(F^c|E^c)$.

That's a mess'o'probabilities—16 in all. The key observation is that all these numbers are determined by the 4 intersection probabilities, or counts. Let's see how.

i) First, the 4 simple probabilities. This comes down to the picture




FIGURE 4.5.2: A probability Q-diagram

in Figure 4.5.2. If we know the probabilities in the four quadrants, we know the probabilities in top and bottom, or left and right rectangles. Algebraically, if we intersect the events in E , E^c and S with F , we get

4.5 Organizing Related Probabilities

$E \cap F$, $E^c \cap F$ and F . The relation $E \dot{\cup} E^c = S$ of [COMPLEMENT RELATIONS 3.7.23](#) yields $(E \cap F) \dot{\cup} (E^c \cap F) = F$. Applying the [OR ELSE FORMULA FOR PROBABILITIES 4.2.6](#), we find that $\Pr(E \cap F) + \Pr(E^c \cap F) = \Pr(F)$. Likewise, if we intersect with, F^c we find that $(E \cap F^c) \dot{\cup} (E^c \cap F^c) = F^c$ and $\Pr(E \cap F^c) + \Pr(E^c \cap F^c) = \Pr(F^c)$.

PROBLEM 4.5.3: Counts work just the same way. Show how to find $\#E$, $\#F$, $\#E^c$ and $\#F^c$ from $\#(E \cap F)$, $\#(E \cap F^c)$, $\#(E^c \cap F)$, and $\#(E^c \cap F^c)$.

ii) Once we have the all the simple and intersection probabilities, we have every conditional probability since the [CONDITIONAL PROBABILITY FORMULA 4.4.2](#) $\Pr(E|F) = \frac{\Pr(E \cap F)}{\Pr(F)}$ expresses each conditional at the ration of an intersection and a simple probability.

PROBLEM 4.5.4: Show that any 3 of the 4 intersection probabilities determine the fourth (and hence all 16 [PROBABILITIES FROM TWO EVENTS 4.5.1](#)). Hint: What is the union of the 4 quadrants in [FIGURE 4.5.2](#)?

Table displays take advantage of this observation by laying out the 4 intersection probabilities in a way that makes it easy to read off all the others. It’s easiest to see how from an example.

EXAMPLE 4.5.5: Suppose we are just given the intersection counts for a pair of events as in

	Male	Female	Totals
Honors Core	142	162	
Regular Core	563	645	
Totals			

Here our sample space S is the set of entering freshmen at a university, and the events are E = “student is Female” (with complement the set of Male students) and F = “student is taking the Honors core” (with complement the students taking the Regular core). As we’ll see, there often little to be gained by introducing letters to denote the events: we just need to keep the meaning of each straight. The columns of the table already tell us about the gender (E) and the rows tell us about the core (F).

4.5 Organizing Related Probabilities

The **interior** or upper left portion of the table we have been provided with contains intersection counts which are located where the corresponding column and row meet. For example, the number $\#(E^c \cap F)$ of Male Honors students is the 142 located where the Male column intersects the Honors row. The number 645 counts Female students taking the Regular core, or $\#(E \cap F^c)$.

Next, by simply totaling along the rows and columns, we complete the table.

	Male	Female	Totals
Honors Core	142	162	304
Regular Core	563	645	1208
Totals	705	807	1512

Thus, the right and bottom **borders** of the table giving totals tell us about simple events. For example, the total number 807 of Female students gives us $\#E$ and the bottom right entry with its grand total 1512 tells us the order of S .

Now, to get simple and intersection probabilities, we just divide each entry in the table by $\#S = 1512$. For example, the chance that a student is a Male in the Regular core, $\Pr(E^c \cap F^c)$, is $\frac{563}{1512}$. The chance that a student is in the Honors core, $\Pr(F)$, is $\frac{304}{1512}$.

Conditional probabilities are always an intersection probability divided by the given (simple) probability. So to find these we take the intersection entry for the pair of events and divide by total entry for the given event. For example, the chance that a Female student is in the Honors core, $\Pr(F|E)$, is $\frac{162}{807}$ —the number of Female Honors students divided by the number of Female students. The chance that a Regular student is male, or $\Pr(E^c|F^c)$, is $\frac{563}{1208}$.

PROBLEM 4.5.6: Use the table in [EXAMPLE 4.5.5](#), to find the probability that:

- A student is Male.
- A student is Male and in the Honors core.

4.5 Organizing Related Probabilities

- iii) An Honors student is Male.
- iv) An Male student is in Honors core.

LOCATING PROBABILITIES IN TABLES 4.5.7: In tables, *interior* cells give intersection counts or probabilities, and *border* cells give simple counts or probabilities. To find each kind of probability below, use the ratio indicated:

- i) (Simple) Border total for the event row or column over the bottom right, grand total.
- ii) (Intersection) Interior value where the row and column of the two events meet over the bottom right, grand total.
- iii) (Conditional) Interior value where the row and column of the two events meet over the border total for the given event row or column.

PROBLEM 4.5.8: Complete the following table of percentages of people with Anomalous and Normal *trichromacy* (red-green-blue vision) by Gender. Hint: As a check, what is the grand total, if we are dealing with percentages?

	Men	Women	Totals
Anomalous	3.2	0.2	
Normal	47.2	49.4	
Totals			

Use your table to find the probabilities that:

- i) A randomly chosen person is a Woman.
- ii) A randomly chosen Woman has Normal trichromacy.
- iii) A randomly chosen person is a Woman with Normal trichromacy.
- iv) A person with Normal trichromacy is a Woman
- v) A person with Anomalous trichromacy is a Woman

We often want to split up a sample space not into just 2 disjoint pieces but into several. In other words, we want to consider a **PARTITION OF S** 3.4.10 of S into pairwise disjoint sets whose union is all

4.5 Organizing Related Probabilities

of S . Tables work just as well to handle probabilities with such partitions. The *only* difference is that there will be as many columns (or rows) as there are subsets in the partition. Let's try an example, again involving **trichromacy**, but where we distinguish various subtypes.

PROBLEM 4.5.9: Complete the following table of percentages of people with Normal and three Anomalous types of **trichromacy** by Gender.

	Men	Women	Totals
Protanomaly	0.66	0.01	
Deuteranomaly	2.50	0.18	
Tritanomaly	0.01	0.01	
Normal	47.21	49.42	
Totals			

Use your table to find the probabilities that:

- i) A randomly chosen person is a Man.
- ii) A randomly chosen Man has Deuteranomaly.
- iii) A randomly chosen person is a Man with Deuteranomaly.
- iv) A person with Protanomaly is a Man
- v) A person with Deuteranomaly is a Man
- vi) A person with Tritanomaly is a man

A table can also be specified by data other than the intersection probabilities. The most common way to do this is to give the simple probability for one event and conditional probabilities in which this event is the given. In this case, we can use each conditional probability to find a corresponding intersection probability by multiplying using the **INTERSECTION PROBABILITY FORMULA 4.4.3**. A classic example is the analysis of the British National Security Strategy that we looked at in **TROLLING FOR TERRORISTS**. In the next problem, we'll verify the counts given there.

PROBLEM 4.5.10: Suppose that Britain contains 50,000,000 Innocent Citizens and 4,500 Terrorists. A computer system attempts

4.5 Organizing Related Probabilities

to identify Britons as either innocent or terrorist by monitoring private information about them. Assume that the system classifies 99% of Innocent Citizens as innocent and 99% of Terrorists as terrorist. Use this information to complete the table of counts below. You should find that the total number of innocent Britons is 49,500,055.

	Innocent Citizens	Terrorists	Totals
innocent			
terrorist			
Totals			

Use your table to find the probabilities that:

- i) A randomly chosen Briton is an Innocent Citizen.
- ii) A randomly chosen Briton is an Innocent Citizen who is classified as innocent.
- iii) A randomly chosen Briton is a Terrorist.
- iv) A randomly chosen Briton is an Terrorist who is classified as terrorist.
- v) A randomly chosen Briton who is classified as innocent is a Terrorist.
- vi) A randomly chosen Briton who is classified as terrorist is a Terrorist.

The same difficulties that arise in trying to identify terrorists also come up in many other identification problems. The most famous class occur in medicine when trying to diagnose illnesses. Let’s record some useful terms before discussing the issues.

DIAGNOSTIC TESTING 4.5.11: A character of a population is just a subset of it. In other words, if we view the members of the population as outcomes in a sample space S then members with the character are a subset or event of S . It’s convenient to write C^+ for this event (thinking of it as with the members who *have* the character) and C^- for its complement (members who do *not* have the character).



4.5 Organizing Related Probabilities

We prevalence of the character is the **equally likely outcomes** probability $\Pr(C^+) = \frac{\#C^+}{\#S}$ that a randomly selected member has the character.

A test for a character C^+ is just another event T^+ of S , with complement T^- . We say that outcomes (i.e. members of S) in T^+ **test positive** for the characteristic C^+ and that outcomes in T^- **test negative**.

Members in the four quadrants determined by C^\pm and T^\pm are named as shown in **FIGURE 4.5.12**.




FIGURE 4.5.12: True and false positives and negatives

The **specificity** of a test T for C^+ is the probability that a person who actually *has* the character *tests positive* for it: this is just the conditional probability $\Pr(T^+|C^+)$.

The **sensitivity** of a test T for C^+ is the probability that a person who does *not have* the character *tests negative* for it: this is just the conditional probability $\Pr(T^-|C^-)$

The **positive predictive value** of a test T for C^+ is the probability that a person who *tests positive* for the character does *has* it: this is just the conditional probability $\Pr(C^+|T^+)$.

The **negative predictive value** of a test T for C^+ is the probability that a person who *tests negative* for the character does *not have* it: this is just the conditional probability $\Pr(C^-|T^-)$.

4.5 Organizing Related Probabilities

The terms “false positive” and “false negative” come from medicine where they stand for the two different ways that a diagnostic test can fail. First, the test may classify ill patients—that is, those who have the condition or character being tested for—as healthy. Such an error is called a **false negative**: negative because the test reports *no* illness and false because the patient really *is* ill. But the test may also classify healthy patients—that is, those who don’t have the condition or character being tested for—as ill. Such an error is called a **false positive**: positive because the test reports an ill patient and false because the patient is actually healthy. In the cases “true positive” and “true negative”, the test is correct, correctly classifying the ill and the healthy, respectively.

In **PROBLEM 4.5.10**, we are trying to diagnose terrorism. So the Terrorists who are classified as innocent are the false negatives, and the Innocent Citizens who are classified as terrorist are the false positives. In our terrorist example, the test correctly classifies 99% of Innocent Citizens as innocent so it has a **specificity** of 99%. It also has a **sensitivity** of 99% because it classifies 99% of Terrorists as terrorist, but, for most tests, the specificity and sensitivity are not equal.

In **PROBLEM 4.5.10**, you should have found that 4,455 Terrorists and 500,000 Innocent Citizens were classified terrorist, so the **positive predictive value** of the test—the chance a terrorist is not really an Innocent Citizen—is $\Pr(\text{Terrorist}|\text{terrorist}) = \frac{4,455}{504,455} \approx 0.0088$. Likewise, 45 Terrorists and 49,500,000 Innocent Citizens were classified as innocent, so the **negative predictive value** of the test—the chance an innocent really is an Innocent Citizen—is $\Pr(\text{Innocent}|\text{innocent}) = \frac{49,500,000}{49,500,045} \approx 0.99999$.

You might be impressed by that negative predictive value of 0.99999.

PROBLEM 4.5.13: Show that a test (easily administered) that declares *everybody* innocent has a negative predictive value of $\frac{50,000,000}{50,004,500} \approx 0.99991$.



4.5 Organizing Related Probabilities

In other words, getting a high *negative* predictive value is easy: just declare few people **terrorist**. Getting a high *positive* predictive value is hard. Can we see why?

As we noted in [TROLLING FOR TERRORISTS](#), it's natural to expect that the errors we should worry about are the false negatives. After all, these are Terrorists on whom we are not going to keep an eye because they have been declared **innocent**. Similarly, in a medical situation, the important problem seems to be ill patients who have been declared healthy because these patients won't be looking for treatment for the condition.

But in both cases, it's really the false positives that concern those administering the test. The reason is that we're trying to diagnose a rare condition that applies only to a tiny group (the terrorists or the patients with the illness). This means that the false negatives are a small percentage of a tiny group, so there just aren't very many of them. In our terrorist example, there is just 1 false negative in a million or 50 in all of Britain.

The true positives are most of this tiny group (the 4950 Terrorists classified as **terrorist**) but they're still a tiny group. The group of false positives is a small percentage of the main population (Innocent Citizens or healthy patients) and this group can be much larger than the tiny group were testing for. Of the 504,950 Britons classified **terrorist**, fully 500,000 are false positives and only 4,950 are true positives. That's less than 1%— it's the positive predictive value—so a **terrorist** classification provides essentially no evidence of being a Terrorist.

Graphically put, we're looking for a few needles in a big haystack with a test that reports **straw** or **needle**. The problem is that even if our test seems quite accurate, the pile of **needles** turns out to be a smaller haystack containing a few needles.

The problem doctors have to face is that its neither practical nor ethical to treat all the positives when most of them are really healthy.

4.5 Organizing Related Probabilities

Impractical because you’re spending most of your resources treating the healthy and unethical because most treatments have unpleasant or even dangerous side effects. The next problems illustrate some of these issues.

The **Center for Disease Control** (CDC) estimated that in 2006 the **prevalence of HIV in the United States** was about 0.0037 (roughly $\frac{1}{270}$ Americans are infected with the virus). The Food and Drug Administration (FDA) requires that a rapid HIV test must have a **sensitivity of 98% and a specificity of 98%**¹

EXAMPLE 4.5.14: Let’s analyze what to expect from employing such a test on a randomly selected sample of 1,000,000 Americans. First we set up a table.

	HIV ⁻	HIV ⁺	Totals
T ⁻	0.98 · 996300 = 976374	0.02 · 3700 = 74	976374 + 74 = 976448
T ⁺	0.02 · 996300 = 19926	0.98 · 3700 = 3626	19926 + 3626 = 23552
Totals	0.9963 · 1000000 = 996300	0.0037 · 1000000 = 3700	1000000

The steps, indicated by the calculations, are:

- i) Fill in the bottom row (HIV⁻ and HIV⁺) using the prevalence 0.0037.
- ii) Fill in the left column using the specificity, $\Pr(T^+ | \text{HIV}^+)$, plus the **INTERSECTION PROBABILITY FORMULA 4.4.3** and the **COMPLEMENT RELATIONS 3.7.23**.
- iii) Fill in the middle column similarly but using the sensitivity, $\Pr(T^- | \text{HIV}^-)$.
- iv) Fill in the right column by totaling the rows.

Now, we can read off the positive predictive value,

$$\Pr(\text{HIV}^+ | T^+) = \frac{\#(\text{HIV}^+ \cap T^+)}{\#T^+} = \frac{3626}{23552}$$

¹Actually, the requirement both stricter and looser. Since we have no equally likely outcomes model for these probabilities, they must be *estimated empirically*. The actual requirement is that the likelihood that the sensitivity and specificity are at least 98% is predicted by experimental trials to be at least 95%. This means both that it is consistent with the trials, though unlikely, that true levels are lower and also that it is likely that they are greater.

4.5 Organizing Related Probabilities

or about 15%, and the negative predictive value,

$$\Pr(\text{HIV}^- | T^-) = \frac{\#(\text{HIV}^- \cap T^-)}{\#T^-} = \frac{976374}{976448}$$

or better than 99.99%.

In other words, essentially everyone who tests negative really is HIV free but 85% of those who test positive are too. This is not hopeless, like [PROBLEM 4.5.10](#) where almost all the terrorists were Innocent Citizens. But it does pose a problem. We've now reduced our haystack from 1,000,000 to 23,552 and there are a lot of needles—3,626—in that haystack. How do we tell them from the remaining 19,926 pieces of straw? Perhaps you've noticed a related gap in my presentation. How did we ever find out that the specificity of the test was 98% when most of those who test positive are incorrectly diagnosed?

The answer to both questions is that we use a more **specific** test—that is one with fewer false positives—often called a **gold standard**. Warning: gold standard tests are seldom perfect; to qualify as a gold standard, a test just has to be better than any competitors. The **gold standard** in HIV testing is to first perform an HIV enzyme immunoassay (EIA) looking for HIV-1 and HIV-2 antibodies and, if this is positive, to confirm it with an HIV-1 Western blot or immunofluorescence assay (IFA). This combination has very high specificity.

Thus, the standard protocol in clinical situations is that a positive rapid HIV test result is an indication that further testing should be performed: the patient is given the EIA/IFA combination test. Likewise, when performing trials to assess the specificity and sensitivity of a rapid test, *every* patient is also given the gold standard tests and these are used to decide if trial test was right or wrong.

Why not just give everyone the gold standard test in a clinical situation too? Why even bother with the rapid tests? The first answer is that it's almost always much more expensive to a gold standard test and it takes much longer to get the result back. Although prices have



4.5 Organizing Related Probabilities

been falling on all tests, rapid tests cost as little as a tenth what gold standard tests do. In the U.S., the rapid test lets doctors save the difference on the 98% of the population that it reports as negative. In the developing world, rapid tests are often the only ones cheap enough. A second answer is that rapidity itself has advantages: up to a quarter of those who take a slower test never return to learn the result. The situation is similar for a range of diseases.

The next few problems deal with empirical data, so this is one time that it is easiest to convert all fractions to decimal form.

PROBLEM 4.5.15: OraQuick™ is one rapid HIV test that was used in testing programs in New York City in which the observed prevalence of HIV was 0.8%. In this problem, we'll compare claims about the test to observations. Since we are dealing with empirical data, this is one problem where it'll be easiest to convert all fractions to decimal form.

- i) OraQuick's [manufacturer claimed that its trials showed](#) that the test has a likely sensitivity of 99.8% (and was 95% certain to have a sensitivity higher than 99.6%) and that it had a specificity of 100%. Find the positive and negative predictive value of this test by creating a table like that in [EXAMPLE 4.5.14](#), but assuming the prevalence of 0.008 and a sample space of 31,122 patients.
- ii) From November 2007–April 2008, the New York clinics gave 31,122 patients the OraQuick test. They found 213 false positives and 231 true positives, both determined by follow-on gold standard tests. Use this data to create another table (again assuming prevalence of 0.008 and sensitivity of 100%).
- iii) What specificity and positive predictive value does *this* table give?

Now let's turn another important disease, breast cancer for which the rapid test is mammography (a radiological exam and reading) and the gold standard is biopsy (minor surgery to extract a sample of suspect tissue followed by microscopic examination to determine

4.5 Organizing Related Probabilities

malignancy). Because biopsies are very expensive and invasive, they are a last resort, even though mammography is a rather weak test.

PROBLEM 4.5.16: Here is some data from [a summary of a study](#) reported in the British Journal of Medicine.

- i) In a sample of 122,355 women aged 50 – 64 who underwent screening mammograms, the mammography was positive in 3,885, breast cancer was detected in 726 of whom 629 also had positive mammograms. Use this data to build a table and then use your table to find the specificity, sensitivity, positive and negative predictive value of screening mammograms.
- ii) Other studies report lower positive predictive values for screening mammograms than this one, and indicate that the value decreases quite rapidly for younger women. Apply the sensitivity and specificity you have found in i) to a sample of 10,000 women under 50 having a prevalence of 0.002 for breast cancer, and show that the positive predictive value for women in this age group is less than 0.05.

I close this subsection with a slightly more complicated example from genetics that we'll come back to use at the end of the section. We'll consider a single gene, and before we start we need to define a few basic genetic terms.

To simplify, we assume that our gene has just two **alleles** or variants, that we denote G and g . Each child gets a copy of one of its mother's two copies of this gene and a copy of one of its father's. We'll write, for example, Gg to indicate that the one copy was a G and the other a g without worrying about which came from which parent. Individuals have one of 3 possible **zygosities** or genetic types: those whose genes are GG are called G -homozygous, those whose genes are Gg are called heterozygous, and those whose genes are gg are called g -homozygous².

²The meanings here, "homo" for same, "hetero" for different, should be familiar in other contexts.

4.5 Organizing Related Probabilities

We also assume that these alleles are expressed in each individual in two observable or **phenotypic** traits G and g (like black hair and blond hair, or normal vision and color-blind vision). It is standard to use upper-case letters to denote **dominant** alleles and traits and lower-case letters to denote **recessive** alleles and traits. Here, **dominant** and **recessive** mean that individuals with at least one G allele (those who are GG , Gg or gG) exhibit the G character and that only those whose genes are gg show the g character.

PROBLEM 4.5.17: In this problem, we want to consider the genetic types of married couples. We assume that any two individuals of opposite sexes are equally likely to marry, and that the proportions of each of the three zygositys—the **zygotic frequencies**—is the same for men and women.

- i) First, suppose that we know that in a certain population 49% of people are GG , 42% are Gg and 9% are gg . These three numbers are the **zygotic frequencies**. Make a table with rows for the 3 zygositys of the wife and a total row, columns for the 3 zygositys of the husband and a total column. Use the percentages given to fill the **border** of the table, and then fill the 9 interior cells with the the probabilities of seeing a marriage of each type. (For example, of the 49% of women who are GG , 42% will marry men who are Gg so the chance of such a (GG, Gg) marriage will be $0.49 \cdot 0.42 = .2058$ or 20.58%. Likewise, 20.58% of marriages will be between a Gg woman and a GG man. Check that the 9 interior cells in your table sum to 1 (or 100%).
- ii) Now, suppose that we know that the **zygotic frequencies** in a certain population are that 20% of people are GG , 70% are Gg and 10% are gg . Make a table exactly like the one in part i) but using these percentages to fill the border of the table. Check that the 9 interior cells in your table sum to 1 (or 100%).
- iii) Now, we want to replace the percentages in part i) by variables, so we assume that the fraction of the population that is GG is x , the fraction that is Gg is y and the fraction that is gg is z . Of course, these fractions give the whole population so $x + y + z = 1$. Create a

4.5 Organizing Related Probabilities

table like that in i) but using these fractions to fill the **border**. Check that the total of all 9 interior cells in your table is $(x + y + z)^2$ (Hint: expand the square). Of course, this total is 1, as it should be, because $x + y + z = 1$.

Working with Probabilities in Tree Diagrams

Recall that the motivation for using tables was to have a way of displaying a few key probabilities involving E and F that makes it easy to find all the rest of the 16 probabilities in **PROBABILITIES FROM TWO EVENTS 4.5.1**. For tables, the key probabilities were the intersection probabilities, from which we got the simple probabilities by totaling and the conditionals by taking quotients.

There's a second common set of key probabilities that arise in the **compound experiments** of **PROBABILITY MEASURES FOR COMPOUND EXPERIMENTS**. To prepare for what's coming, I'll call the two events here G and W . Think of **EXAMPLE 4.4.17**, **EXAMPLE 4.4.18** and **PROBLEM 4.4.19** where G was "green urn" and W was "white ball" (and, of course $G^c = R$ or "red urn" and $W^c = B$ or "black ball"). The key probabilities are $\Pr(G)$ of the simple event G and the two conditional probabilities $\Pr(W|G)$ and $\Pr(W|G^c) = \Pr(W|R)$. Showing that these suffice is a good review of all the basic probability rules. The following problem takes you through this step by step. Work it carefully because we'll repeat these steps in most of the problems later in this section.

PROBLEM 4.5.18: State the formula that tells you how to do each of the following. As you do, find each new probability, assuming that $\Pr(G) = \frac{1}{2}$, $\Pr(W|G) = \frac{10}{12}$ and $\Pr(W|R) = \frac{4}{6}$.

- Find $\Pr(R)$ from $\Pr(G)$.
- Find $\Pr(B|G)$ and $\Pr(B|R)$ from $\Pr(W|G)$ and $\Pr(W|R)$.
- Find $\Pr(G \cap W)$ from $\Pr(G)$ and $\Pr(W|R)$.
- Find $\Pr(G \cap B)$ from $\Pr(G)$ and $\Pr(B|R)$.



4.5 Organizing Related Probabilities

- v) Find $\Pr(R \cap W)$ from $\Pr(R)$ and $\Pr(W|R)$.
- vi) Find $\Pr(R \cap B)$ from $\Pr(R)$ and $\Pr(B|R)$.
- vii) Find $\Pr(W)$ from $\Pr(G \cap W)$ and $\Pr(R \cap W)$.
- viii) Find $\Pr(B)$ from $\Pr(G \cap B)$ and $\Pr(R \cap B)$.
- ix) Find $\Pr(G|W)$, $\Pr(G|B)$, $\Pr(R|W)$, and $\Pr(R|B)$ from the values in iii)–viii).

Tree diagrams are a way of displaying this second set of key probabilities that guides us graphically in finding the rest. First some dendrological terms. Most of these are borrowed from the trees of the forest, but the meanings are a bit more precise and things are rotated clockwise 90° so that, instead of growing *up*, our trees grow from *left to right*³. A tree starts with a single **root** vertex at the far left. We'll draw this and other vertices to come as rectangles. Out of this root come line segments, heading right (but also up or down if they like), that we call first **branches**. At the end of each branch is another vertex (rectangle). It's standard to draw the first level branches so that all these vertices are vertically aligned. Next, second level branches come out of each of *these* vertices, again heading right and again ending in new vertices. At some point, usually after two sets of branches, we reach vertices that only have branches in (from the left) and none out (to the right). These vertices are called **leaf** vertices or usually, just **leaves**.

On the left of [FIGURE 4.5.19](#) is the most common tree, with two levels of **branches** and two branches coming out of each vertex (except the 4 **leaves**) and with its parts labeled. In the middle, is a more complicated tree with three levels of branches and more than two branches out in the second level. On the right is a less symmetric tree—we won't see many of these—where the branching is not uniform at each level, and not all leaves are on the last level.

Tree diagrams turn out to be perfect for recording the outcome of a compound experiment, like those we discussed in [PROBABILITY MEA-](#)

³It could be worse. In computer science, trees grow *down*




FIGURE 4.5.19: Sample tree diagrams


SURES FOR COMPOUND EXPERIMENTS. The basic idea is that vertices correspond to events (shown in red) and branches to probabilities (shown in blue). Let’s see how to display [PROBLEM 4.4.19](#) as a tree, then we’ll write down a general method for filling in tree diagrams. Along the way we’ll see all the probabilities in [PROBLEM 4.5.18](#) appear: these were also based on [PROBLEM 4.4.19](#) with G the event “green urn” (so $\Pr(G) = \frac{1}{2}$ and $R = R$) and W the event “white ball” (so $B = B$). (so $\Pr(W|G) = \frac{10}{12}$ because 10 of the 12 balls in the green urn are white and $\Pr(W|R) = \frac{4}{6}$ because 4 of the 6 balls in the red urn are white).

The first step is to draw the tree and label its vertices. Each branch is used to record what we observed—that is, what happened—in one of the component experiments and we describe this observation in the box for the vertex at the right end of the branch. To keep the number of branches to a minimum, we usually describe what happened

4.5 Organizing Related Probabilities

in terms of events rather than outcomes. Specifically, we'll usually focus on a **partition** of the sample space of outcomes of each component. Remember a partition is just a collection of events such that every outcome is in one and only one of the events.


In our example, we want to record the outcome—the urn picked, R or G —in the first component experiment. But when we pick a ball, all we care about is the *color* of the ball so we partition the sample space—set of all the balls—into the events B = “black” and W = “white”. The second level vertices get compound “two-letter” labels that tell us what pair of events leads is to them—for example, the top one is reached by picking the red urn and a white ball so we label it RW . It would be more accurate to use the label $R \cap W$ but dropping the “ \cap ” makes the diagram more legible. Thus we get the simple tree below.




The second step is to attach a probability to each branch of the tree, that records the chance of taking that branch from the starting vertex. Thus, on the first level branches, we want the chances of picking the 2 urns which are both $\frac{1}{2}$. On the second level branches we want the chance of picking a white or black ball *from the chosen urn*. These are *conditional probabilities* that depend on the number of balls of each color in that urn (10 white and 2 black for the red urn and 4 white and 2 black for the green urn). Thus the chance of going from vertex R to vertex W is $\Pr(W|R)$ and this is $\frac{10}{12}$ because 10 of the 12 balls in the red urn are white. The probabilities we need at this step

4.5 Organizing Related Probabilities

are the 3 key ones and those in parts i) and ii) of [PROBLEM 4.5.18](#). Now we've got the following picture.



The third step is to attach a probability to each of the **leaves**. These are the chances of reaching that leaf from the root at the far left. To do this we have to follow the right first branch and also follow the right second branch, so we need to use the [INTERSECTION PROBABILITY FORMULA 4.4.3](#). For example, next to the RW vertex to the right of the R vertex, we want the probability $\Pr(R \cap W) = \Pr(R) \cdot \Pr(W|R) = \frac{1}{2} \frac{10}{12} = \frac{10}{24}$. Here we've written down the probabilities in [iii\)-vi\)](#) of [PROBLEM 4.5.18](#). Once we've found all these probabilities, we also rewrite them, as in the picture below, over a common denominator. This makes these probabilities easy to add. We'll use this in the last step, but already we can take advantage of it to check all the calculations to this point. Just total the intersection probabilities as shown. If you do not get 1 something is wrong. Why?



4.5 Organizing Related Probabilities

The fourth and last step is to group the leaves corresponding to “like” outcomes in the second component experiment: here this just means the same colored ball. In the picture below, I indicate these groupings by red arcs (red because we’re tying together *events*) and next to each I’ve put its (blue) probability. Because the first level vertices represent mutually exclusive events (different urns), the leaves do too, so we just need to add the probabilities of the leaves being joined. For example, the chance of seeing a white ball is $\Pr(W) = \frac{10}{24} + \frac{8}{24} = \frac{18}{24}$. Here we’ve found the probabilities in parts vi) and viii) of [PROBLEM 4.5.18](#). Note how the common denominator from the previous step means that we just have to add numerators to find these probabilities. The common denominator also makes answering further questions about the diagram easier. Here’s the final tree diagram.

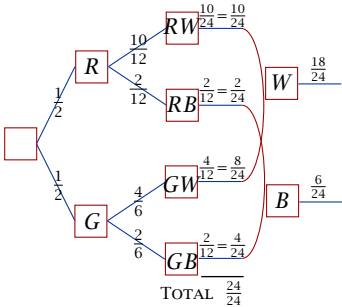


FIGURE 4.5.20: Tree diagram for [PROBLEM 4.4.19](#)

What about the probabilities in the last part, ix), of [PROBLEM 4.5.18](#)? These are usually what we want use the tree diagram to figure out! And they are now easily obtained by simply plugging in to the [CONDITIONAL PROBABILITY FORMULA 4.4.2](#). For example, we can read off both the numerator $\Pr(R \cap W)$ and denominator $\Pr(W)$ of $\Pr(R|W)$ to see confirm the value $\frac{10}{18}$ from [PROBLEM 4.4.19](#).

PROBLEM 4.5.21: Read off the other probabilities in ix) of [PROBLEM 4.5.18](#) from [FIGURE 4.5.20](#).

4.5 Organizing Related Probabilities

TREE DIAGRAM METHOD 4.5.22: *To draw a tree diagram that summarizes a compound experiment, follow the steps below:*

Step 1: FILL IN THE VERTICES: Draw the tree—vertices and branches—and fill in each vertex with the event you must observe to reach it. First level vertices correspond to events in the first stage of the compound experiment. Second level vertices correspond to events in the second stage, and so on.

Step 2: FILL IN THE EDGES: Label each edge with its probability. The first level edges are labeled with the simple probability of the event in the first stage to which they lead. Second and higher level edges are labeled with the conditional probability of following that edge given that you are already at its starting vertex.

Step 3: FILL IN THE LEAVES: Label each leaf with the intersection probability obtained by multiplying the probabilities on each edge from the root to the leaf.

Step 4: TOTAL LIKE LEAVES: Group leaves that correspond to the same event in a second (or later) stage of the experiment, and find the simple probabilities of these events by totalling the probabilities of the grouped leaves.

You are now ready to find any probabilities involving the events in your diagram.

To get a feel for this method, we'll first apply it to some examples we worked earlier.

PROBLEM 4.5.23: To get a feel for this method, we'll first apply it to some examples we worked earlier. Pay *careful* attention to *what probability* is being asked for in each part.

i) Draw a tree diagram for [EXAMPLE 4.4.17](#). Then use it to find the probabilities below.

a. You drew a white ball out of the green urn.

b. You drew a white ball given that you chose the green urn.



4.5 Organizing Related Probabilities

- c. You chose the green urn given that you drew a white ball.
- ii) Draw a tree diagram for [EXAMPLE 4.4.18](#). Then use it to find the probabilities below.
- If your urn was red, what's the chance your ball was black?
 - What's the chance you drew a black ball out of the red urn?
 - If your ball was black, what's the chance your urn was red?

Now let's do a couple of problems with a more complicated trees. What this means is that we **partition** the sample space at the first or second stage (or both) into more than 2 mutually exclusive events. Instead of dividing into just one event and its complement, we divide into several pieces. The footwork in filling out such a tree diagram is essentially the same in the example above. The only differences are that there is now one each for each event in the partition, so in [ix](#)) (FILL IN THE EDGES), we need to know the probabilities of each of these events. In example, these probabilities are almost always easy to read off. Here's an example.

PROBLEM 4.5.24: You are given 3 urns, a red urn that contains 8 white and 4 black balls, a green urn that contains 3 white and 15 black balls, and a yellow urn that contains 9 white and 3 black balls. An experiment consists of picking an urn at random and then picking a ball at random from that urn.

- i) Draw a tree diagram of this experiment. It will have three first level branches leading to vertices R , G and Y for the three urns, and two second level branches out of each of these leading to vertices RW and RB , GW and GB , and YW and YB .
- ii) Use your tree diagram to read off the probabilities below.
- If your urn was yellow, what's the chance your ball was black?
 - What's the chance you drew a black ball out of the yellow urn?
 - If your ball was black, what's the chance your urn was yellow?
 - If your ball was black, what's the chance the urn had more white balls than black ones? Hint: What color urns have more white balls than black?

4.5 Organizing Related Probabilities

PROBLEM 4.5.25: You are given 2 urns, a red urn that contains 6 white, 6 black and 5 pink balls, and a green urn that contains 3 white, 9 black and 3 pink balls. An experiment consists of picking an urn at random and then picking a ball at random from that urn.

- i) Draw a tree diagram of this experiment. It will have two first level branches leading to vertices R and G for the two urns, and three second level branches out of each of these leading to vertices for the color of the ball, RW , RB and RP , or, GW , GB and GP .
- ii) Use your tree diagram to read off the probabilities below.
 - a. If your urn was green, what's the chance your ball was pink?
 - b. What's the chance you drew a pink ball out of the green urn?
 - c. If your ball was not pink, what's the chance your urn was green?

PROBLEM 4.5.26: In this problem, we consider families with exactly 3 children and we assume that each child is equally likely to be a boy or a girl. Our experiment consists of first choosing a child at random and recording whether this child is the eldest (E), middle (M) or youngest (Y) child in its family. We then record whether the chosen child has an *older sister* (OS) or not (NOS).

- i) Draw a tree diagram of this experiment.
- ii) Use your tree diagram to read off the probabilities below.
 - a. The chance that the chosen child has an older sister.
 - b. The chance that the chosen child is the youngest if he or she has an older sister.
 - c. The chance that the chosen child is the eldest if he or she has an older sister.

CHALLENGE 4.5.27: Redo [PROBLEM 4.5.26](#), but considering families with exactly 4 children.

OK. Tree diagrams are not so bad to complete or work with. But they do seem to involve a lot of work to reach answers could also find without using them. Yes, once you have the tree diagram, it's easier to read off the answers, but are they really needed? I'll conclude with a deceptively simple problem that explains why we use them.



4.5 Organizing Related Probabilities

PROBLEM 4.5.28: A bag contains two fair coins that we'll call A and B . The coins are identical except that coin A has heads on both sides, coin B has a head and a tail. You first pick a coin at random from the bag, then you toss it and it comes up heads.

- Without worrying about formulas, what's the chance that the other side of the coin you picked has a tail on it?
- Now check your answer by drawing a tree diagram of this compound experiment (with events A and B at the first level and H and T at the second) and finding $P(B|H)$.

Why is the obvious answer $\frac{1}{2}$ wrong? The chance that we picked coin B was $\frac{1}{2}$ since we picked at random, and if we did pick B , then there's definitely a tail on the other side. Both those statements are true. The problem is that these statements describe an experiment different from the compound one we performed in which we didn't just pick a coin, we *first* picked a coin and *then* a side of that coin. What the tree diagram does for us is to give us a way of laying out such an experiment that forces us to respect its compound structure. It forces us to realize that there's a $\frac{1}{4}$ th chance of seeing either side of either coin and hence a $\frac{3}{4}$ th chance of seeing a side with a head on it. Since only 1 of those 3 sides is on coin B , the chance we picked B *given that we saw a head* is the ratio $\frac{1}{3}$ of these.

Finally, a question that many students ask at this point. Suppose you're given a bunch of data about some events and asked to lay it out and then answer some questions. When so you need to use a table and when do you need to use a tree-diagram? Never! at least in the sense that any problem you can answer using a table you can also answer using a tree-diagram, and vice versa. As we have seen, all the flavors of probability can be recorded in either. So if you have a strong personal preference for one format or the other, feel free to indulge it.

Analysing the game of Craps

In this section, we'll look at a final application of conditional probabilities that's quite a bit more complicated than any of the examples we have dealt with so far because it involves an *infinite* sample space.

We ask what are the chances of winning at the classic dice game of **craps**. Craps evolved from an older European dice game called **hazard** or, in French, *hasard*. A version was introduced to the United States in New Orleans in the early 19th century where it was known as *crapaud* or “toad” and the name craps is an English abbreviation of this name.

Let's start by explaining how the game is played. We begin by rolling 2 dice. This first roll is known as the **come-out**. You win if your come-out total is either 7 or 11 (a **natural**) and lose if your come-out total is 2, 3 or 12 (known as **craps**). If you roll another total, the game gets trickier. The total you rolled becomes your **point** and you begin to **shoot** the dice. Shooting involves rolling the dice until your roll totals either your point from the come-out roll, in which case you win, or a 7, in which case you “7 out” and lose.

How many times do you **shoot**? As many as it takes for you to roll either your point or a 7. That could be any number of rolls whatsoever. As the number of rolls increases, the chance that you won't have rolled either your point or a 7 gets smaller and smaller. But that chance is never 0 and we therefore need to consider shooting *every positive number of rolls*. That's what makes analyzing craps trickier than any of the examples we have dealt with so far. We'll need to put everything we have learned to this point together and we'll also learn a lesson that mathematics teaches again and again. It's a version of the old carpenter's rule about measuring and cutting: think twice, calculate once.

Let's start. The first point to observe is that craps is an multi-stage “experiment” with one stage for each roll. That means a tree diagram



4.5 Organizing Related Probabilities

is a perfect way to record the outcomes of every possible sequence of rolls. We won't be able to ever display the *entire* tree for a game of craps because it has an infinite number of nodes and branches but by studying portions of this tree, we will be able to determine the chance of winning the game.

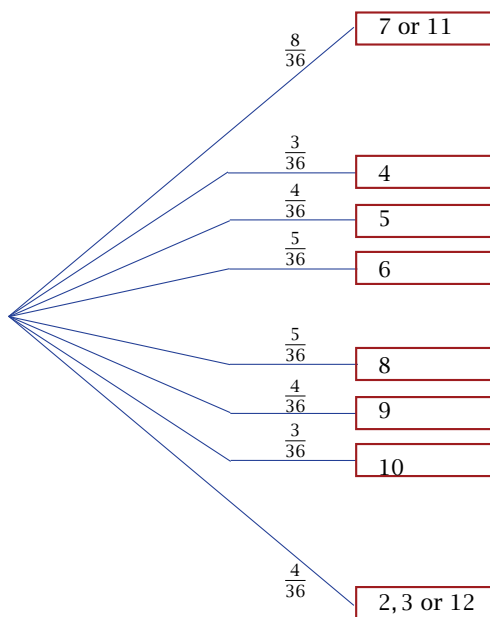


FIGURE 4.5.29: First level branches for Craps

Let's start with the first stage, the **come-out roll**, for which the tree is shown in [FIGURE 4.5.29](#). Although the picture looks complicated (it has 8 branches), the probabilities on all these branches can be read off from [TABLE 4.3.5](#). Moreover, the first level nodes for a natural and craps are already leaves (and go no further).

What about the 6 branches that lead to the points 4, 5, 6, 8, 9 and 10? They all continue indefinitely. That's exactly our difficulty. To come

4.5 Organizing Related Probabilities

to terms with this, let's focus on the branch that leads through the point 9.

The next key remark is that at the second stage, we only need to distinguish 3 possibilities. We might roll another 9 (shown going up) and win or we might roll a 7 (shown going down) and lose. These two branches again lead us to leaves. Any other total corresponds to following the middle, horizontal branch. But at the end of this branch, we are in exactly the same position that we were in at its start—shooting for a 9!.




FIGURE 4.5.30: Second level branches for a point of 9

The bad news is that we make no progress when we go down this branch. The good news is that it's the only branch that leads to third level branches. Better still, the third level branches look exactly like the second level ones: we can go up with a winning 9, down with

4.5 Organizing Related Probabilities

a losing 7 or, with any other total, go across and roll again. This pattern simply repeats indefinitely. It means that we can easily draw as many stages of the tree as we have the patience for. In [FIGURE 4.5.31](#) I have shown 4 stages with their branch probabilities (again, just transcribed from [TABLE 4.3.5](#)).




FIGURE 4.5.31: Multiple levels of the tree for a point of 9

At this point, we have—ignoring the infinite number of stages not shown—completed the second step in the [TREE DIAGRAM METHOD 4.5.22](#). The third stage in the method is to **FILL IN THE LEAVES** by multiplying along the branches. That’s just arithmetic which I have carried out below. I’ve written the product giving each leaf probability in two ways, first, in the order in which we encounter the factors as we move down the branches to the leaf in [FIGURE 4.5.32](#), and then with factors grouped into powers in [FIGURE 4.5.33](#).

What I did *not* do was to evaluate the products! Why not? Because there’s a pattern in the factors that get’s lost when we multiply them out. That pattern is the key to the seemingly impossible task completing the tree diagram. What’s left to do? According to [TREE DIAGRAM METHOD 4.5.22](#), we need to **TOTAL LIKE LEAVES**. In other words, we need to add up the entire *infinite* set of “up” probabilities that lead to a 9) to find the simple probability of a win (and likewise total

4.5 Organizing Related Probabilities




FIGURE 4.5.32: “Raw” leaf probabilities for a point of 9




FIGURE 4.5.33: “Grouped” leaf probabilities for a point of 9

the set of “down” probabilities that lead to a 7 and a loss). The total we’re after is:

$$\left(\frac{4}{36}\right)^2 + \left(\frac{4}{36}\right)^2 \left(\frac{26}{36}\right) + \left(\frac{4}{36}\right)^2 \left(\frac{26}{36}\right)^2 + \left(\frac{4}{36}\right)^2 \left(\frac{26}{36}\right)^3 + \cdots$$

or, after taking out the common factor,

$$\left(\frac{4}{36}\right)^2 \left(1 + \left(\frac{26}{36}\right) + \left(\frac{26}{36}\right)^2 + \left(\frac{26}{36}\right)^3 + \cdots\right)$$

Notice that we get each term in this new sum from the previous one by multiplying by $\frac{26}{36}$. What that means is that our total is the sum

4.5 Organizing Related Probabilities

of a **geometric series** with ratio $r = \frac{26}{36}$. It's to be able to recognize this geometric series that we needed to avoid multiplying out the leaf probabilities. Applying the **GEOMETRIC SERIES FORMULA 1.3.6**, we find that the total is

$$\frac{1}{1-r} = \frac{1}{1-\frac{26}{36}} = \frac{1}{\frac{10}{36}} = \frac{36}{10}.$$

That means that the total we're after is

$$\left(\frac{4}{36}\right)^2 \cdot \left(\frac{36}{10}\right) = \frac{4}{36} \cdot \frac{4}{10} = \frac{2}{45}.$$

To get the middle product, I just cancelled a 36 above and below. You should have reached this answer in the last part of **PROBLEM 1.3.8**.

PROBLEM 4.5.34: Imitate the argument above to show that the total of the “down” probabilities that lead to a 7 and a loss in **FIGURE 4.5.33** is the infinite sum

$$\left(\frac{4}{36} \cdot \frac{6}{36}\right) \left(1 + \left(\frac{26}{36}\right) + \left(\frac{26}{36}\right)^2 + \left(\frac{26}{36}\right)^3 + \cdots\right).$$

Then use the **GEOMETRIC SERIES FORMULA 1.3.6** to evaluate this sum as

$$\left(\frac{4}{36} \cdot \frac{6}{36}\right) \cdot \left(\frac{36}{10}\right) = \left(\frac{4}{36}\right) \cdot \left(\frac{6}{10}\right) = \frac{3}{45}.$$

We can check all these calculations quickly by adding the chance of winning after making a point of 9 to the chance of losing after making a point of 9. These should just total the chance that the come-out roll was a 9, which is $\frac{4}{36}$. And indeed, what we get is

$$\left(\frac{4}{36}\right) \cdot \left(\frac{4}{10}\right) + \left(\frac{4}{36}\right) \cdot \left(\frac{6}{10}\right) = \left(\frac{4}{36}\right) \cdot \left(\frac{4+6}{10}\right) = \frac{4}{36}.$$

What about the points other than 9? They can be handled in just the same way with only minor changes in arithmetic.

PROBLEM 4.5.35: In this problem, you'll calculate the chance of winning when you make a point of 8.

4.5 Organizing Related Probabilities

- i) Show that the tree that describes the shooting stages of the game of craps when the point is 8 has exactly the same branches and nodes as [FIGURE 4.5.31](#).
- ii) Use the fact that chance of rolling an 8 (either in the first stage to make the point or later on to win the game) is $\frac{5}{36}$ to fill in the branch probabilities on the 8 version of the tree. (You'll need to consider how the “across” probability changes too.)
- iii) Next, find the leaf probabilities on the 8 version of the tree in [FIGURE 4.5.32](#).
- iv) Finally use the [GEOMETRIC SERIES FORMULA 1.3.6](#) to show that the chances of winning and of losing at craps with a point of 8 are $\frac{5}{36} \cdot \frac{5}{11}$ and $\frac{5}{36} \cdot \frac{6}{11}$.
- v) Check your answers by showing that they total to the probability of rolling an 8 on the come-out roll.

At this point, it's pretty clear that with enough elbow grease we can compute the chance of winning after making any point. But before calculating once, let's stare at the values we have calculated for 9 and 8 and think twice. What probabilities have we been calculating anyway? For a point of 9, we were after the probability of coming out with a 9 and also winning (or losing). Symbolically, we wanted $\Pr(9 \cap \text{"win"})$. The [INTERSECTION PROBABILITY FORMULA 4.4.3](#) tells us we can write this as the product $\Pr(9 \cap \text{"win"}) = \Pr(9) \cdot \Pr(\text{"win"}|9)$. Our answer is also a product: $\Pr(9 \cap \text{"win"}) = \frac{4}{36} \cdot \frac{4}{10}$.

Can we match these up? For the first factor, the answer is certainly yes, as $\Pr(9) = \frac{4}{36}$. But that means that the two second factors must also match up. We must have $\Pr(\text{"win"}|9) = \frac{4}{10}$.

Could we have predicted this? Indeed we could. Once we've made a point of 9 we know the game will end when we shoot a 9 or a 7. In the former case, we win and in the latter we lose. In other words, the chance of winning by shooting a point of 9 is the chance of rolling a 9 given that we rolled either a 9 or a 7. In symbols, $\Pr(\text{"win"}|9) = \Pr(9|9 \cup 7)$.



4.5 Organizing Related Probabilities

PROBLEM 4.5.36: Use the **CONDITIONAL PROBABILITY FORMULA 4.4.2** to show that $\Pr(9|9 \cup 7) = \frac{4}{10}$ and deduce that $\Pr(\text{"win"}|9) = \frac{4}{10}$. Then show that $\Pr(\text{"lose"}|9) = \Pr(7|9 \cup 7) = \frac{6}{10}$.

Looking back at [FIGURE 4.5-31](#), we can give another interpretation. To get to a leaf of this tree, we have to go “up” or “down” at some stage, and of the 10 rolls that do one or the other 4 go up to a 9 and the other 6 down to a 7.

PROBLEM 4.5.37: Use the same argument to show that $\Pr(\text{"win"}|8) = \Pr(8|8 \cup 7) \frac{5}{11}$ and use this to check your answer to [PROBLEM 4.5.35](#).

PROBLEM 4.5.38: Show that $\Pr(\text{"win"}|10) = \frac{3}{9}$.

In hindsight, neither the tree diagram nor the geometric series formula is essential, though it's hard to imagine how we'd have understood what was going on without them. In any case, by imitating what we did for points of 8 and 9 we can now fill in [TABLE 4.5.39](#).

TOTAL T	2	3	4	5	6	7	8	9	10	11	12
$\Pr(T)$	$\frac{1}{36}$	$\frac{2}{36}$	$\frac{3}{36}$	$\frac{4}{36}$	$\frac{5}{36}$	$\frac{6}{36}$	$\frac{5}{36}$	$\frac{4}{36}$	$\frac{3}{36}$	$\frac{2}{36}$	$\frac{1}{36}$
$\Pr(\text{"WIN"} T)$	0	0	$\frac{3}{9}$	$\frac{4}{10}$	$\frac{5}{11}$	1	$\frac{5}{11}$	$\frac{4}{10}$	$\frac{3}{9}$	1	0
$\Pr(T \cap \text{"WIN"})$	0	0	$\frac{1}{36}$	$\frac{2}{45}$	$\frac{25}{296}$	$\frac{1}{6}$	$\frac{25}{396}$	$\frac{2}{45}$	$\frac{1}{36}$	$\frac{1}{18}$	0

TABLE 4.5.39: PROBABILITIES FOR WINNING AT CRAPS

Then we just need to total the bottom row of the table to find the probability of winning, getting

$$0+0+\frac{1}{36}+\frac{2}{45}+\frac{25}{396}+\frac{1}{6}+\frac{25}{396}+\frac{2}{45}+\frac{1}{36}+\frac{1}{18}+0=\frac{244}{495}\simeq 0.492929.$$

A few closing comments. On the one hand, notice how carefully tuned the game is. The chance of winning is just less than 50% and the house edge is only 1.41% ($0.507071 - 0.492929 = 0.141414$). On the other hand, the player winds on the come out roll exactly twice as often as he or she loses (there are 8 rolls that give a 7 or 11 but only



4.6 Applications of Bayesian probability

4 that give a 2, 3 or 12). This leads many players to feel that the game is actually tilted in their favor.

At craps, you can also bet with the house and against the shooter but with some proviso. Most typically, if the shooter rolls a 12 and loses then you do not get paid: this has the effect of giving the house an edge of 1.36%. This unusual two-sided betting was an invention of an American named John Winn, whose goal was to discourage the house from using loaded dice by letting you bet “with or against” the dice.

Both the “with or against” edges are much smaller than in any other casino game—and other casino rules, that I won’t go into, often allow the player to reduce this edge to under 1%. If you’re a sucker, don’t worry; the casino also has lots of so-called proposition bets on a craps table where it’s edge is as high as 16%.

4.6 Applications of Bayesian probability

The same ideas are applied, not just in toy problems like the last one, but also in many practical situations. In fact, these applications are so common they form a whole specialty, known as **Bayesian probability**. In most treatments of conditional probability, there’s a whole section of the text entitled **Bayes’ Theorem**. The theorem is named after a Scottish Presbyterian minister Thomas Bayes who stated a version of it in a paper that his friend, Richard Price had published after his death (if you like, you can read the [original article](#) as it appeared in the Philosophical Transactions of the Royal Society of London).

The “Theorem”, in its modern form, is really just a special case of the **CONDITIONAL PROBABILITY FORMULA 4.4.2**, but stated in a complicated way that makes it look both impressive and scary. As we’ll

4.6 Applications of Bayesian probability

see, it's neither—in fact, all it does is wrap up in a messy formula the tail end of the procedure we've been using in the [TREE DIAGRAM METHOD 4.5.22](#).

To keep things definite, let's see how [Bayes' Theorem](#) would handle a conditional probability like the chance of having picked the red urn, given that we chose a white ball. We start from $\Pr(R|W) = \frac{\Pr(R \cap W)}{\Pr(W)}$ and proceed in two step. One step re-expresses the denominator in this conditional probability and the other the numerator.

For the denominator, the idea is stupidly simple. Suppose that R , G and Y are three events that partition the sample space S (think of the three urns in [PROBLEM 4.5.25](#)) meaning that S is the disjoint union of these three subsets as shown in [FIGURE 4.6.1](#). Any event B (shown outlined in black) W (outlined in grey for visibility) is the disjoint union of its intersections with R , G and Y , again, as you can see.



FIGURE 4.6.1: Bayes' "Theorem"

In terms of disjoint unions, this means that $W = (W \cap R) \dot{\cup} (W \cap G) \dot{\cup} (W \cap Y)$, and in turn, applying the [OR ELSE FORMULA FOR PROBABILITIES 4.2.6](#), that $\Pr(W) = \Pr(W \cap R) + \Pr(W \cap G) + \Pr(W \cap Y)$. In the tree diagram method, we effectively do this, but without the formula, when we [GROUP LIKE LEAVES](#). That is, we add up the intersection probabilities for the RW , GW and YW leaves to find $\Pr(W)$.

[Bayes' "Theorem"](#) involves nothing more than replacing the denominator $\Pr(W)$ in a [CONDITIONAL PROBABILITY FORMULA 4.4.2](#) like $\Pr(R|W) = \frac{\Pr(W \cap R)}{\Pr(W)}$ with this sum to get

$$\Pr(R|W) = \frac{\Pr(W \cap R)}{\Pr(W \cap R) + \Pr(W \cap G) + \Pr(W \cap Y)} .$$

4.6 Applications of Bayesian probability

I've stated everything in terms of a partition into 3 pieces but the same idea clearly applies with any number of pieces E_1, E_2, \dots, E_n and leads to the really scary formula

$$\Pr(R|W) = \frac{\Pr(W \cap R)}{\sum_{i=1}^n \Pr(W \cap E_i)} = \frac{\Pr(W \cap R)}{\Pr(W \cap E_1) + \Pr(W \cap E_2) + \dots + \Pr(W \cap E_n)}.$$

that still just means GROUP LIKE LEAVES.

Relax. You don't need to learn this formula; you won't even have to use it.

OK. What about the numerator? Remember that our urn and ball problems have a temporal order: *first* we choose the urn, *then* we choose the ball. That's why, conditional probabilities in which the given is the color of the ball and we ask about the color of the urn seem like nonsense at first. How can use information that came *after* to we ask a question about what happened *before*? "Tree diagrams tell us how", is the answer.

Recall that we compute the numerator $\Pr(W \cap R)$ in a tree diagram using the [INTERSECTION PROBABILITY FORMULA 4.4.3](#), $\Pr(W \cap R) = \Pr(R) \cdot \Pr(W|R)$. Then we plug this into the [CONDITIONAL PROBABILITY FORMULA 4.4.2](#) to get the

TIME REVERSAL FORMULA 4.6.2: *If W has non-zero probability, then*

$$\Pr(R|W) = \frac{\Pr(R) \cdot \Pr(W|R)}{\Pr(W)}.$$

Note the way the forward conditional probability $\Pr(W|R)$ (ball given urn) is used to find the reversed one $\Pr(R|W)$ (urn given ball). Once again, there's no real need to memorize this formula, because, when you set up a tree diagram, you are automatically led to compute $\Pr(W \cap R)$ in this way. The [TREE DIAGRAM METHOD 4.5.22](#) is all you need to know.

This reversal of time is the essential element in most applications of [Bayes' theorem](#). For reasons that I do not understand, however, it's usually not even mentioned when the theorem is stated. What you

4.6 Applications of Bayesian probability

see is formula above, with its messy denominator and unchanged numerator.

Indeed, the term **Bayesian probability** is almost exclusively used to refer to probability questions that reverse the stages of our experiment in this way. There's no reason we couldn't use the same formulas that occur in **Bayes' "Theorem"** to rewrite $\Pr(R)$ as $\Pr(R \cap W) + \Pr(R \cap B)$ in the conditional probability formula for $\Pr(W|R)$. But no one ever does, because in this probability the given (red urn) happens *before* what we are asking about (white ball).

Put differently, Bayesian probability asks how we can use what we just saw to adjust our existing expectations about what we already observed. In fancy terms, "**Posterior knowledge** influences **prior expectation**". This is a useful notion that we *will* use and it's what we'll mean in the future by **Bayesian probability**. In this section, we'll look at a couple of the many applications of this sort of reversal of time.

Filtering spam

The easiest place to start, because it most clearly reflects the reversal of time is what's called Bayesian spam filtering. The problem here is to distinguish **spam** (emails trying to see you a con or a "product" that you don't want in your inbox) from **ham** (emails you want to receive), and filter our or remove the spam.

What's needed is a diagnostic test, like those discussed following **DIAGNOSTIC TESTING 4.5.11**. But there we were trying to spot rare conditions, like being a terrorist or having HIV, and were looking for a few needles in a haystack. Estimates are that about 70% or more of all emails are spam, so here most of the haystack consists of needles. Nonetheless, our main concern is still false positives. You'd far rather have to delete [a Nigerian 419 come-on](#)⁴ from your inbox every

⁴Even one from a Harvard faculty member.

4.6 Applications of Bayesian probability

day than miss the notification of your Harvard admission interview. So a good spam filter must be extremely **sensitive**, even if this means being a bit less **specific**.

Many very effective spam filters use Bayesian probability to identify spam emails by the words they contain. We'll look at a very simplified version of how this works that shows the main strategy but leaves out a lot of refinements needed to produce a highly sensitive, adequately specific filter. Let's use S and H to denote the events of seeing emails that are spam and that are ham—of course, $H = S^c$ —and W and N for seeing the words we're interested in, or not seeing them. The idea is to look for a collection W of words for which the conditional probabilities $\Pr(S|W)$ is very close to 1—hence $\Pr(H|W)$ is very close to 1. This ensures high sensitivity (an email containing W is almost never ham and can be safely discarded) and it's what we'll look at.


Real filters are considerably more complex. A lot of spam does not contain *any* words that are “smoking gun” markers for spam. It's often necessary to combine look for related markers: for example, in [one well-known database](#) only about 65% of emails that contain the word ‘free’ are spam, but 99.99% of emails that contain ‘free!!’ are. Often filters try to identify other words that identify an email containing a suspicious word as ham: for example, emails that contain ‘free’ but also contain ‘election’ or ‘trade’ are probably not spam. Further, to be specific, and succeed in throwing away most spam, W must contain words that *are* found in most spam emails. Finding the right combination of techniques to achieve very high sensitivity with good specificity is what makes designing a good filter hard. We won't say any more here.

Here are the questions I do want to look at: “How do we decide what W characterize spam?”, and “How can we calculate $\Pr(S|W)$?” The answer to both questions is to start with a database of consisting of a large number—say 10,000— of emails. We want to know which

4.6 Applications of Bayesian probability

emails in the database are in H and which are in S, and we want the makeup of the database to be similar to that of the flow of emails we're going to filter. For example, if we expect 70% of emails to be spam then we'd want about 7,000 pieces of spam and 3,000 pieces of ham in our database.

Next, we make a sort of dictionary of possible W's. A dictionary entry could be a single word like 'Free!!', or a short phrase like 'low interest rate'. Then, for each W in the dictionary we count how many spam and how many ham emails contain W. At this point we have the information needed to fill in the following tree diagram for each W.



Then from the diagram, we compute $\Pr(S|W)$ as $\Pr(S \cap W)$ divided by $\Pr(W)$. But how did we find the last numerator? We used $\Pr(S \cap W) = \Pr(S) \cdot \Pr(W|S)$. So what the tree diagram is doing is guiding us to

$$\Pr(S|W) = \frac{\Pr(S) \cdot \Pr(W|S)}{\Pr(W)},$$

which is exactly the [Time Reversal Formula 4.6.2](#). In other words, our tabulation of how often spam emails contains W can be turned around to predict how often emails containing W are spam.

This method has a few other nice features. First, you don't need to know in advance what words to look for. As long as you provide enough examples of spam and ham, the analysis identifies good candidates. This also makes it possible to train a filter for an individual user by having him or her add emails to the spam or ham portions of the database.

4.6 Applications of Bayesian probability

Here’s a problem to illustrate how combining terms can improve accuracy.

PROBLEM 4.6.3: The following table shows the frequency with which S and H emails contained the individual words “diploma” and “apply”, and both of these words.

	diploma	apply	diploma and apply	Total
S	960	845	840	7694
H	22	66	1	1870
Total	982	911	841	9564

- i) How likely is an email that contains the word “apply” but not “diploma” to be H?
- ii) How likely is an email that contains the words “apply” and “diploma” to be S?

“\$7’ll get you \$12” Revisited

We’re now in a position to explain the counterintuitive results we got when we played the game of “\$7’ll get you \$12”. Let’s recall how it’s played. To begin, the grifter shuffles the three cards face down and the mark puts her finger on one of the three cards. The grifter then looks at the two *other* cards and turns up, or *exposes* a King, reducing the number of face down cards to 2. Note that whether or not the mark has her finger on a King, the grifter can always find a King to expose. The mark bets \$7 on a face-down card, and if the card she bets on is the Queen, she receives \$12 (\$5 plus her bet of \$7), if not she loses her \$7.

It seems clear that this is a losing proposition. The mark has a 50-50 chance of picking the Queen, and, if so, she’ll win \$5 half the time and lose \$7 the other half averaging a loss of \$1 every time she plays. And since it’s a toss-up, it doesn’t seem to matter what strategy the mark uses to pick her card.

4.6 Applications of Bayesian probability

Our first set of experiments, in which mark made her final choice at random, by flipping a coin, confirmed this. However, we found that when the mark always picked the card she had not fingered initially, she *won* an average of \$1 per play, and when she always left her finger where it started, she lost an average of \$3 a play. It seems impossible to reconcile these results with each other, let alone with our intuition. But a little Bayesian probability will make everything clear.

Let's make a tree diagram of what happens, assuming the queen is the left card. The first stage of experiment is for the mark to choose a card (Left, Middle or Right) at random giving 3 first level branches. At the second stage, the grifter exposes a King, but not the card chosen in the first step. So there are 2 second level branches when the mark picked the Queen on the Left, but only 1 when she picked one of the Kings (Middle or Right) At the third stage, the mark either leaves her finger in place or moves it to the other unexposed card. So there are always 2 third level branches. We need to GROUP LIKE LEAVES in *two* ways in this problem: First, according to whether the mark picked the queen or a King (*Q* or *K*)—she'd did if the leaf ends in an *L*, and didn't if it ends in an *M* or *R*; Second, by whether she Stayed with her finger on her original card or Changed to the other card (*S* or *C*). In the diagram, we'll do this in two stages, first grouping by the 4 intersections such, then grouping these again by whether she stayed or changed. Here then is the tree:

I have marked the probabilities on the branches of this tree, but at the third level (Stay or Change), I have used the letter p to denote the mark's probability of Staying and q her probability of changing. Of course, we must have $p + q = 1$.

To give a full description of “\$7'll get you \$12”, we should really have 2 more copies of this tree, corresponding to starting with the Queen in the middle and on the right. But we don't need to show these, precisely because they *are* just copies with the main branches

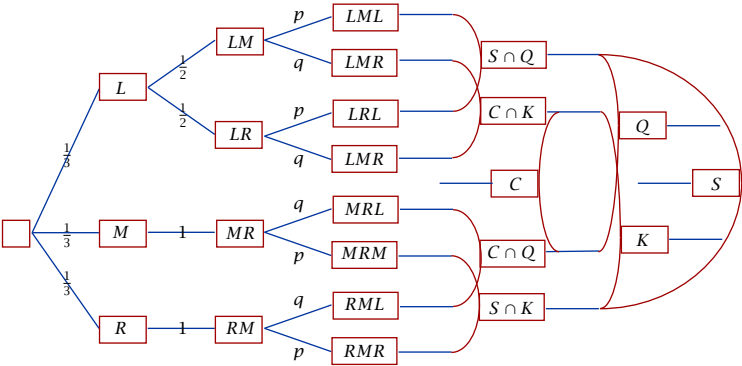


FIGURE 4.6.4: Playing “\$7’ll get you \$12”

repositioned by the same probabilities everywhere. For example, the tree we’d see if the Queen was the Right card is just the one in [FIGURE 4.6.4](#) turned upside down.

PROBLEM 4.6.5: Fill out the remainder of the tree 3 times with the values of p and q suggested below. As check of your calculations, you should find, in each case, that $\Pr(S)$ and $\Pr(C)$ are the values of p and q , respectively, that you are using.

- i) The first time use the value $p = q = \frac{1}{2}$: this models the experiment where the mark just flips a coin to decide whether to stay or change.
 - a. The mark’s chance of winning.
 - b. The mark’s chance of winning if she stayed with her original card.
 - c. The mark’s chance of winning if she changed to the other unexposed card.
- ii) The second time once use the value $p = 0$ and $q = 1$: this models the experiment where the mark always changes her card.
 - a. The mark’s chance of winning.
 - b. The mark’s chance of winning if she originally chose the Left

4.6 Applications of Bayesian probability

card.

c. The mark's chance of winning if she originally chose the Middle card.

d. The mark's chance of winning if she originally chose the Right card.

iii) The third time once use the value $p = 1$ and $q = 0$: this models the experiment where the mark never changes her card.

a. The mark's chance of winning.

b. The mark's chance of winning if she originally chose the Left card.

c. The mark's chance of winning if she originally chose the Middle card.

d. The mark's chance of winning if she originally chose the Right card.

Let's summarize what we have learned. When the mark always Stays she wins only when she originally selected the Queen (the left card) and so wins $\frac{1}{3}^{\text{rd}}$ of the time. This explains the loss of \$3 a play we saw in [EXPERIMENT 2.1.2](#). If she plays 3 times, she can expect to win \$5 once and lose \$7 twice for a net loss of \$9 or \$3 a play.

When the mark always Changes she wins when she originally selected a King (the middle or right card) and so wins $\frac{2}{3}^{\text{rds}}$ of the time. Again, this explains the gain of \$1 a play we saw in [EXPERIMENT 2.1.2](#). If she plays 3 times, she can expect to win \$5 twice and lose \$7 once for a net gain of \$3 or \$1 a play.

When she Stays and Changes $\frac{1}{2}$ the time each, she wins $\frac{1}{2}$ the time. But this $\frac{1}{2}$ decomposes as a $\frac{1}{3}$ chance of winning on plays where she Changes plus a $\frac{1}{6}$ chance of winning on plays where she Stays. So once again, she wins on $\frac{2}{3}^{\text{rds}}$ of plays where she Changes and $\frac{1}{3}^{\text{rd}}$ of plays where she stays. From what we've seen just above, if the mark plays 6 times, Staying 3 times and Changing 3, then she'll expect to lose \$9 on the plays where she stays and win \$3 on those where she

4.6 Applications of Bayesian probability

changes for a net loss of \$6 or \$1 per play. Once again this is very much what we saw in [EXPERIMENT 2.1.2](#).

PROBLEM 4.6.6: Fill out the remainder of the tree one more time but this time leave p and q as variables. As check of your calculations, you should find that $\Pr(S) = p$ and $\Pr(C) = q$. Now calculate:

- i) The mark's chance of winning.
- ii) The mark's chance of winning if she stayed with her original card.
- iii) The mark's chance of winning if she changed to the other unexposed card.

How should the mark play if she wants to break even on the average? Hint: If the mark wins $\frac{7}{12}$ ^{ths} of the time, then if she plays 12 times, she will expect to win \$5 on 7 trys and lose \$7 on the other 5, so breaking even.

I hope that many of you have realized by now that the game of “\$7'll get you \$12” is really just the famous Monty Hall problem (based on the game show “Let's Make a Deal”). At the end of each episode a lucky contestant would be faced with three doors. Behind one was a valuable prize (the Queen) and behind the other two were zonks (the Kings), much less valuable products placed there to get a plug from show announcer Jay Stewart. The contest (the mark) chose a door, then host Monty Hall had his assistant, the lovely Carol Merrill (the grifter), open one of the doors not chosen to reveal a zonk. The contestant then had the choice of Staying with the door originally chosen, or Changing⁵ to the other unopened door, after which she won whatever was behind her chosen door.

As with finding the lady in “\$7'll get you \$12”, Changing wins the valuable prize $\frac{2}{3}$ ^{rds} of the time, Staying only $\frac{1}{3}$ rd. One quick way to see this, is to change the rules slightly. Suppose the contestant is offered

⁵Called “switching” on the show. I used Changing here to have a different first letter from Staying.



FIGURE 4.6.7: The set and cast of “Let’s Make a Deal”

the choice between Staying with her original pick, or Changing to *both* of the other doors but losing the right to any zonks behind them. It’s clear this time that Staying wins the valuable prize $\frac{1}{3}^{\text{rd}}$ and Changing $\frac{2}{3}^{\text{rds}}$ of the time. Does it matter if one of the other doors is opened to reveal a zonk? No! Nothing has changed: the originally chosen door still hides the prize $\frac{1}{3}^{\text{rd}}$ of the time. The other two doors still hide it $\frac{2}{3}^{\text{rds}}$ of the time.

Here’s a very easy way to think about such problems. Suppose that a door is selected and *another* door is opened to reveal a zonk. Then the chance the selected door hides the prize remains fixed, and the chance the opened door was hiding the prize gets shared out equally between any the doors that are neither selected nor opened. In the case of 3 doors, the $\frac{1}{3}^{\text{rd}}$ chance that the prize was behind the opened door is inherited by the other unselected door, giving it a $\frac{2}{3}^{\text{rds}}$ chance of hiding the prize.

CHALLENGE 4.6.8: Suppose that “Let’s Make a Deal” were played with 4 doors, one hiding a valuable prize and 3 hiding zonks. You pick a door, then Carol Merrill opens one of the 3 doors not selected to reveal a zonk and Monty gives the choice of Staying or Changing to one of the two other doors. After you decide, Carol Merrill opens one

4.6 Applications of Bayesian probability

of the other 2 doors you didn't pick to reveal a zonk. Finally, Monty gives you the choice of Staying or Changing to other unopened door. Find the probability of winning the valuable prize if you adopt each of the four strategies below:

- i) Stay both times.
- ii) Stay the first time and switch the second
- iii) Switch the first time and stay the second.
- iv) Switch both times.

Hint: The four probabilities, in increasing order of *size*, are 0.250, 0.375, 0.625 and 0.750.

The Hardy-Weinberg Principle

Bayesian ideas are also widely applied in genetics. Here's an example that continues [PROBLEM 4.5.17](#). The 9 interior cells in the table you built there correspond to pairs of **zygosities** (for the husband and wife in a marriage), so they can be viewed as a couple-zygosity or type. We want to use the frequencies in these cells to understand what fraction of children born to each type of couple will be of each **zygosity**.

A few preparations will make this much easier. First, we expect that a child will get each of the mother's two genes half the time—i.e. equally often—and both of these will pair half the tie with each of the father's two genes. So we expect to see each of the 4 possible pairs of maternal and paternal genes in $\frac{1}{4}$ of children. This let's save a lot of labor, since it implies that switching the zygosities of husband and wife does not affect the fraction of children of each type. For example, a $GG \times Gg$ couple will contribute genes GG , Gg , GG , and Gg equally often so half the children will be GG and half Gg . Likewise the children of a $Gg \times GG$ couple will be half GG and half " gG ". But " gG " is the same as Gg , heterozygous.

4.6 Applications of Bayesian probability

PROBLEM 4.6.9: Check that the following types of couples have offspring in the proportions listed.

- i) $GG \times GG$: all GG .
- ii) $GG \times Gg$ or $Gg \times GG$: $\frac{1}{2} GG, \frac{1}{2} Gg$.
- iii) $GG \times gg$ or $gg \times GG$: all Gg .
- iv) $Gg \times Gg$: $\frac{1}{4} GG, \frac{1}{2} Gg, \frac{1}{4} gg$.
- v) $Gg \times gg$ or $gg \times Gg$: $\frac{1}{2} Gg, \frac{1}{2} gg$.
- vi) $gg \times gg$: all gg .

PROBLEM 4.6.10: Consider the experiment of first picking a couple at random and recording which of the six sets of couples in [PROBLEM 4.6.9](#) they belong to, then picking a random child of such a couple and recording its zygosity.

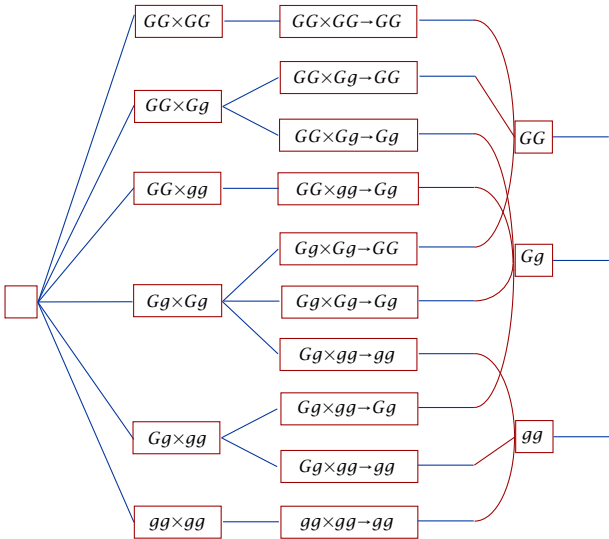


FIGURE 4.6.11: Zygosityes of marriages and offspring

- i) Draw a tree diagram of this compound experiment assuming that the 3 zygosityes have frequencies as given in [PROBLEM 4.5.17.i](#)).

4.6 Applications of Bayesian probability

I have provided you with the first step, the tree itself. The six first level branches correspond to the cases in [PROBLEM 4.6.9](#), and I have listed them in “upper-case before lower-case” order. The number of second level branches depends on which case you have reached. For example, out of $GG \times GG$ there is only one branch because all children are GG , but out of case $GG \times Gg$ there are 2 branches and out of case $Gg \times Gg$ there are 3.

The probabilities on the first level branches can be read off (or summed) from the table in [PROBLEM 4.5.17.i](#)). For example, case ii) occurs $.2058 + .2058 = .4116$ of the time.

After you GROUP LIKE LEAVES, you should recognize the probabilities of the events GG , Gg and gg . What are they? Could we have predicted these values?

- ii) Next use your tree diagram to answer a few questions.
 - a. Suppose that a randomly chosen blue-eyed girl has a younger brother. Show that the probability that her younger brother also has blue eyes is exactly 64%.
 - b. What is the chance that a sibling of brown-eyed child is brown eyed?
- iii) Draw a tree diagram of this compound experiment assuming that the 3 zygositys have frequencies as given in [PROBLEM 4.5.17.ii](#)). Before you start, predict what values you’ll obtain for GG , Gg and gg after you GROUP LIKE LEAVES. Then fill out the tree diagram and compare the values you obtain to your prediction.

What went wrong in the [iii](#))? In [i](#)), our tree diagram predicted that the proportion of children with each zygosity GG , Gg and gg would be exactly the same as those of their parents, on which we based out table in [PROBLEM 4.5.17.i](#)). Surely we should expect this in [iii](#)) as well. But now the tree diagram gives very different proportions—30.25%, 49.50% and 20.25% for the children versus 20%, 70% and 10% for the parents.

What’s going on here? Evolution predicts that distribution of genes in



4.6 Applications of Bayesian probability

a population *will* change under selective pressure. Genes that confer on those who carry them a greater ability to reproduce will become more common at the expense of less “fit” genes. But that’s not what we’re seeing here. The numbers here come purely from the conditional probability of sexual reproduction. If this process causes gene distributions to fluctuate by large amounts from one generation to the next, won’t this wipe out any effects of natural selection, which produces gradual changes over periods of many generations? The answer is, “Yes, it would”.

Fortunately, it doesn’t. What the two examples above show that that some zygotic frequencies (like the one in [PROBLEM 4.5.17.i](#)) are **stable**, meaning that the zygotic distribution of parents is the same as that of their children, and some (like the one in [PROBLEM 4.5.17.iii](#)) are *not*. How can we tell which are which and how can we tell what types we should expect to see in the real world? The answer is to do one more example using the tree in [FIGURE 4.6.11](#) but using variable rather than numerical frequencies. The next problem guides you to do this.

PROBLEM 4.6.12: Recompute the tree diagram in [FIGURE 4.6.11](#) but this time assuming that the 3 zygosities have frequencies given by variables x , y and z .

- i) The probabilities you need for the first level branches can be read off the table in [PROBLEM 4.5.17.iii](#)) and those for the second level branches are the same as in [PROBLEM 4.6.9](#).
- ii) In steps [ix](#)) and [ix](#)) of [TREE DIAGRAM METHOD 4.5.22](#) (FILL IN THE LEAVES and GROUP LIKE LEAVES), your probabilities will be quadratic polynomials in x , y and z . However, if we plug in $x = .49$, $y = .42$ and $z = .09$, the values we obtain should match the probabilities of the tree in [PROBLEM 4.6.10.i](#)). Use this to check your polynomials.
- iii) What should the sum of the GG , Gg and gg probabilities equal? Use this to give another check for your answers.



4.6 Applications of Bayesian probability

OK. Now, how do we tell from the tree diagram whether zygotic frequencies x , y and z are **stable**? Remember this just means that the each of the probabilities x , y and z that give the zygotic frequencies of the parents equals the corresponding probability (at the GG , Gg and gg vertices) that gives the frequency for the children. Since these probabilities are quadratic polynomials, what we get when we equate them are quadratic *equations*! For example, when we equate the GG frequencies, we get $x = x^2 + x \cdot y + \frac{1}{4}y^2$. For the Gg frequency, we get $y = \frac{1}{2}xy + 2xz + \frac{1}{2}y^2 + \frac{1}{2}yz$ and for the gg we get $z = \frac{1}{4}y^2 + y \cdot z + z^2$.

These equations are not hard to unwind, even if they look a bit scary, but there's an easier way to clear the clutter. We can get rid of all the fractions if we just set $Y = \frac{1}{2}y$ (so $2Y = y$ and $Y^2 = \frac{1}{4}y^2$). In terms of Y , the probability we found for GG children is $x^2 + x \cdot y + \frac{1}{4}y^2 = x^2 + x \cdot (2Y) + Y^2 = x^2 + 2x \cdot Y + Y^2 = (x + Y)^2$. Likewise, the probability for gg is $\frac{1}{4}y^2 + y \cdot z + z^2 = Y^2 + (2Y) \cdot z + z^2 = (Y + z)^2$. Two such perfect squares must be telling us something, so let's give names to the square roots, say $p = x + Y$ and $q = Y + z$. We can now re-express the conditions that the GG and gg probabilities are the same in the parent and child populations as $x = p^2$ and $z = q^2$.

The key observation is that $p + q = 1$. This follow directly by plugging in the definitions $p + q = (x + Y) + (Y + z) = x + 2Y + z = x + y + z = 1$. Squaring this $(p + q)^2 = 1$ too, so $p^2 + 2pq + q^2 = 1$. But $x = p^2$ and $z = q^2$ so $x + 2pq + z = 1$. Since $x + y + z = 1$ too, we must have $y = 2pq$.

PROBLEM 4.6.13: If our zygotic frequencies are stable then we know, as above, that y equals $\frac{1}{2}xy + 2xz + \frac{1}{2}y^2 + \frac{1}{2}yz$ by equality of the proportions of Gg parents and children.

- Replace y with $2Y$ in this expression.
- Plug the definitions of p and q into $2pq$ and expand.
- Show that these two expressions are indeed equal, as predicted.

Let's sum up what the algebra has told us.



HARDY-WEINBERG PRINCIPLE 4.6.14: We say that a triple x, y and z of non-negative zygotic frequencies is *stable* if, when a population of parents have these zygotic frequencies, we expect their children to have these frequencies too.

The *Hardy-Weinberg principle* says that, when mating and fertility are random, a triple is stable if and only if there is a number p with $0 \leq p \leq 1$ such that, if we set $q = 1 - p$, then $x = p^2$, $y = 2pq$ and $z = q^2$. We call the curve traced by the points (p^2, pq, q^2) as p runs from 0 to 1 the *Hardy-Weinberg curve*.




FIGURE 4.6.15: The Hardy-Weinberg curve

In [FIGURE 4.6.15](#), the origin $(0,0,0)$ is bottom rear corner. The set of zygotic frequencies in (x,y,z) -space is the gray triangle and the set of stable frequencies is the orange curve. The endpoints of the curve are $(1,0,0)$ corresponding to $p = 1$ and $(0,0,1)$ corresponding to $p = 0$. The point where the curve has largest y -value and turns back on itself is $(\frac{1}{4}, \frac{1}{2}, \frac{1}{4})$ corresponding to $p = \frac{1}{2}$. The frequencies $(.49, .42, .09)$ from [PROBLEM 4.5.17.i](#) lie on the curve, and those from [PROBLEM 4.5.17.ii](#)— $(.2, .7, .1)$ —are well off it. Biologists often call a flattened version of this picture a *de Finetti diagram*.

This principle we have rediscovered gets its name from the fact that it was originally discovered independently in 1908 by G. H. Hardy,

4.6 Applications of Bayesian probability

one of the great mathematicians of the early 20th century (as an answer to a question asked by a biologist colleague) and by Wilhelm Weinberg, a German doctor.

There are two more questions worth asking about zygotic frequencies. First, how are they inherited? We know what it means for a triple of frequencies to be stable: children inherit the frequencies of their parents. But what about the child frequencies when the parent frequencies are *unstable*? Surprisingly, the first generation of children *always* have stable frequencies and therefore these are inherited by the grandchildren and all later generations. In fact, the calculation above also show this, because we only used the formulae for the child frequencies in terms of the parent ones to find p and q . That is, the child frequencies are *always* on the Hardy-Weinberg curve and hence are stable, no matter what the parent frequencies were.

This means that, if the randomness assumptions of the [Hardy-Weinberg Principle 4.6.14](#) hold, we won't have to wait long to observe stable zygotic frequencies. Conversely, if they observe frequencies not on the Hardy-Weinberg curve, biologists suspect that something is not random. Perhaps mating is assortative (people prefer partners with similar genes), as is the case with genes for height (spouses tend to be of similar height). Perhaps some zygotypes are more or less fertile due to **natural selection**. But, it's hard to be sure that the deviation from stability is not due to random variation, especially if the population is small.

The second question is, “Does the magic probability p have an interpretation in terms of genetic frequencies?” The answer is yes, it's simply the proportion of chromosomes that carry the G allele of our gene—and q is just the complementary proportion of g alleles. With this interpretation, it's easy to see why the [HARDY-WEINBERG PRINCIPLE 4.6.14](#) holds. To produce a GG offspring each parent must contribute a G gene, with probability p on each side. Since the events are independent, they both happen pp of the time giving $x = p^2$. Like-

wise, considering gg zygotes we get $y = q^2$. For Gg zygotes, we can either we can get the G from either parent (and the g from the other) giving $y = 2 \cdot p \cdot q$.

4.7 The dice don't talk to each other

In this section, we explore the concept of **independence**. Independence is deceptively simple. On the one hand, when it holds, it is a very powerful tool—so powerful that many, indeed most, standard tools in statistics can only be applied when independence is assumed to hold. On the other, our intuition about independence is extremely unreliable. Many of the most widely and firmly held error in thinking about probabilities are due to incorrect intuition about independence.

Our plan here is to start with the easy part: defining independence and explaining how to test for whether or not it holds. Then we'll tackle a range of examples that illustrate common pitfalls in working with it.

Independence

Informally, two events are **independent** if knowing whether or not one occurred does *not* affect the probability that the other did. To take the simplest example, suppose you flip a quarter and a dime. You expect the quarter to come up heads half the time. Suppose I show you that the dime came up tails. Does this change your expectation for the quarter? No, you still expect the quarter to come up head half the time. Equally, if I'd told you the dime came up heads you'd still expect the quarter to come up head half the time. We express this expectation mathematically by saying that the side that come up on the quarter and the dime are **independent** of each other.

4.7 The dice don't talk to each other

We can recast this example in a slightly different way, Suppose that instead of tossing a quarter and a dime, I toss a quarter twice. You expect the quarter to come up heads on the second toss half the time. If I tell you that the first toss was a tail, does this change your expectation for the second toss? No, you still expect the second toss to be a head half the time. Equally, if I'd told you the first toss was a head, you'd continue to expect the quarter to come up heads half the time on the second toss. And vice-versa, you'd expect the first toss to come up heads and tails equally often whether you knew nothing about the second toss, or whether you knew it was a head. This time we say that the sides that come up on the first and second tosses of the quarter are **independent**. A more graphic way to express this idea is to say that “The coin has no memory”.

For a slightly more complex example, consider rolling our usual blue die and red die. We expect the blue die to show a 2 one-sixth of the time. Does that expectation change if I tell you that the red die showed a 5? Once again, no. In fact, nothing I tell you about what the red die showed seems to give any information about what blue die showed. It wouldn't make any difference if I changed the number 5 on the red die to a 3 or a 6 or any other number. Nor if, instead of specifying the exact number on the red die, I told you that some other event had happened—such as say a number less than 4, or an even number showing. Again we express this mathematically by saying that the numbers showing on the two dice are independent events. The title of this section—“The dice don't talk to each other”—is the graphic way of expressing this intuition.

INDEPENDENT EVENTS 4.7.1: *Fix a sample space S and a probability distribution Pr . The following 3 equations relating the probabilities of 2 events E and F are equivalent (that is, all hold or all fail). If they hold, we say that E and F are **independent events**. If they do not hold, we say that E and F are **dependent events**.*

i) $\text{Pr}(E) \cdot \text{Pr}(F) = \text{Pr}(E \cap F)$.



4.7 The dice don't talk to each other

ii) $\Pr(E|F) = \Pr(E)$.

iii) $\Pr(F|E) = \Pr(F)$.

The second and third equalities capture the informal concept of independence in the examples above. The second says that the probability of observing E *knowing that F occurred* is no different from the probability of observing E with no information about F . Likewise the third says that knowing that E occurred has no effect on the chance of observing F . Put this way, the last two equalities have a definite “direction”—which event do we ask about and which is either given or not—so it's not obvious why they should be equivalent.

On the other hand, the first equality is symmetric: if we swap the order of E and F we get $\Pr(F) \cdot \Pr(E) = \Pr(F \cap E)$ and the probabilities on both sides are unchanged. For this, and other reasons, we almost always want to check whether independence holds using the first equation.

METHOD FOR CHECKING INDEPENDENCE 4.7.2: *To check whether events E and F in an equally likely outcomes probability space S are independent or dependent:*

i) Find $\Pr(E)$, $\Pr(F)$ and $\Pr(E \cap F)$ using the **EQUALLY LIKELY OUTCOMES FORMULA 4.3.2**.

ii) Check, by substituting these 3 values, whether $\Pr(E) \cdot \Pr(F) = \Pr(E \cap F)$ holds or not.

Why did I bother setting out a method when all that's required is to plug into **INDEPENDENT EVENTS 4.7.1.i**? To emphasize one point and help you steer clear of a very common pitfall. The point of emphasis is that you must find $\Pr(E \cap F)$ in **i**) by counting $E \cap F$ and then plugging this count into $\Pr(E \cap F) = \frac{\#E \cap F}{\#S}$. The pitfall is to try to avoid making this count by applying some other formula. The common mistake is to try to use the **INTERSECTION PROBABILITY FORMULA 4.4.3**— $\Pr(E \cap F) = \Pr(F) \cdot \Pr(E|F)$ —but to use $\Pr(E)$ instead of $\Pr(E|F)$ in this formula. Of course, if you do this you'll find, when you plug



4.7 The dice don't talk to each other

in ii), that $\Pr(E \cap F) = \Pr(F) \cdot \Pr(E)$ because you *used the right side to compute the left*. And if you don't, you'll get stuck in a vicious circle when you try to find $\Pr(E|F)$ since the numerator of this conditional probability is $\Pr(E \cap F)$. Let's try this out in the examples above.

EXAMPLE 4.7.3: When we toss the quarter and the dime, we get a sample space $S = \{\text{Hh, Ht, Th, Tt}\}$ in which I use upper case for the quarter and lower case for the dime. If we let $E = \text{"quarter heads"} = \{\text{Hh, Ht}\}$ and $F = \text{"dime tails"} = \{\text{Ht, Tt}\}$ (so $E \cap F = \text{"quarter heads and dime tails"} = \{\text{Ht}\}$), then $\Pr(E) = \frac{2}{4}$, $\Pr(F) = \frac{2}{4}$, and $\Pr(E \cap F) = \frac{1}{4}$. Since $\Pr(E) \cdot \Pr(F) = \frac{2}{4} \cdot \frac{2}{4} = \frac{1}{4} = \Pr(E \cap F)$, E and F are independent.

PROBLEM 4.7.4: Show that $E = \text{"first toss H"}$ and $F = \text{"second toss H"}$ are independent events .

EXAMPLE 4.7.5: When we roll blue and red dice, we get the usual sample space S of 36 ordered pairs of numbers from 1 to 6. The events $E = \text{"blue die shows 2"}$ and $F = \text{"red die shows 1"}$ both have probability $\frac{6}{36}$ and $E \cap F = \text{"blue 2 and red 5"}$ has probability $\frac{1}{36}$. Since $\Pr(E) \cdot \Pr(F) = \frac{6}{36} \cdot \frac{6}{36} = \frac{1}{36} = \Pr(E \cap F)$, E and F are independent.

If we let G be the event "red die is less than 4", which has probability $\frac{18}{36}$, then $E \cap G = \{(2, 1), (2, 2), (2, 3)\}$ and has probability $\frac{3}{36}$. This time $\Pr(E) \cdot \Pr(G) = \frac{6}{36} \cdot \frac{18}{36} = \frac{3}{36} = \Pr(E \cap G)$, so E and G are also independent.

What about F and G ? If I tell you that the red die showed a 1, then you *know for sure* that it's less than 4. If I tell you the red die is less than 4, you don't know it's a 1 but you do know that there are only 3 possible rolls instead of 6. Either way knowing about event *changes* our expectation for the other. That's what it means for the events to be *dependent*. Let's check: $F \cap G$ is the event "red die is 1" and "red die is less than 4". Thus just means that the red die equals 1 so $\Pr(F \cap G) = \frac{6}{36}$. This not the same as $\Pr(F) \cdot \Pr(G) = \frac{6}{36} \cdot \frac{18}{36} = \frac{3}{36}$, confirming our guess that F and G are dependent.

4.7 The dice don't talk to each other

Let's ask about one more pair: F and the event $H =$ “red die is even”; I'll leave you to check that $\Pr(H) = \frac{18}{36}$. Now if we know F happens then the red die is odd and H cannot, and conversely if H happens we cannot see F because 1 is odd. So we guess that this pair is dependent. In fact, F and H are mutually exclusive: if one happens, then the other can't. So $\Pr(F \cap H) = 0$. We don't even need to calculate $\Pr(F) \cdot \Pr(H)$ to see it can't be 0 and confirm our guess of dependence.

You may remember that when we were looking at [OR ELSE FORMULA FOR PROBABILITIES 4.2.6](#). I had you chant “The checks for ‘mutually exclusive’ and for ‘independent’ are *completely* different” a few times. Say it again, please. Now I can explain it. You check that E and F are mutually exclusive by checking that $\Pr(E \cap F) = 0$. You check that E and F are independent by checking that $\Pr(E \cap F) = \Pr(E) \cdot \Pr(F)$. Many students want to confound the two. Don't! With that warning, I hope it is safe to back off a bit and note that the last example generalizes.

MUTUALLY EXCLUSIVE EVENTS ARE DEPENDENT 4.7.6: If E and F are mutually exclusive events with non-zero probabilities, then E and F are dependent. The reason is simple: $\Pr(E) \cdot \Pr(F)$ is *not* 0 because neither factor is, but $\Pr(E \cap F)$ is 0—that's what mutually exclusive means.

Here's a couple of easy ones for you to try.

PROBLEM 4.7.7: Consider the events $E =$ “blue die shows 2”, $F =$ “red die shows 4” and $G =$ “red die is at most 3”. Show that E and F are independent and that E and G are independent but that F and G are dependent.

So far there have been no surprises. Our intuition is nicely captured by the saying that “The dice don't talk”. So if one event involves only information about the blue die and the other about the red die, we expect the events to be independent. Likewise, what we know



4.7 The dice don't talk to each other

about the quarter tells us nothing about the dime. On the other hand, if both events deal with the we expect that information about the blue die tells us nothing involve only the red die, we expect to learn something about both events from the outcome of either—that is we expect dependence.

Now let's turn to some more subtle examples in which we ask about events which involve both dice or all the coins.

EXAMPLE 4.7.8: We toss a coin m times and record which side comes up on each roll, and we ask whether the events E =“first toss is a head” and F =“exactly two tosses are tails” are independent. Intuitively, we'd guess they are not. Knowing that the first toss was a head should make the number of heads we expect to see go up and the number of tails go down. So it ought to affect the chance for a specific number of tails. This guess is *almost* right!

In the table below, I have recorded, for m from 2 to 6, the probabilities of $\Pr(E)$, $\Pr(F)$ and $\Pr(E \cap F)$, as well as the product $\Pr(E) \cdot \Pr(F)$ and the conditional probability $\Pr(F|E)$. To find (the numerator of) $\Pr(F)$ we need to count the ways of picking 2 of m tosses to come up tails, which is $C(m, 2)$ as in [PROBLEM 3.8.27](#). For $\Pr(E \cap F)$, we must pick the two tosses from amongst the last $(m - 1)$ tosses (since the first is a head) getting $C(m - 1, 2)$. I'll leave you to check the other entries.

m	$\Pr(E)$	$\Pr(F)$	$\Pr(E \cap F)$	$\Pr(E) \cdot \Pr(F)$	$\Pr(F E)$	Independent?
2	$\frac{2}{4}$	$\frac{1}{4}$	$\frac{0}{4}$	$\frac{1}{8}$	$\frac{0}{2}$	No
3	$\frac{4}{8}$	$\frac{3}{8}$	$\frac{1}{8}$	$\frac{3}{16}$	$\frac{1}{4}$	No
4	$\frac{8}{16}$	$\frac{6}{16}$	$\frac{3}{16}$	$\frac{3}{8}$	$\frac{3}{8}$	Yes!
5	$\frac{16}{32}$	$\frac{10}{32}$	$\frac{6}{32}$	$\frac{5}{32}$	$\frac{6}{16}$	No
6	$\frac{32}{64}$	$\frac{15}{64}$	$\frac{10}{64}$	$\frac{15}{128}$	$\frac{10}{32}$	No

TABLE 4.7.9: INDEPENDENCE OF “FIRST H” AND “EXACTLY 2 T’S”

4.7 The dice don't talk to each other

As you can see from the $\Pr(F|E)$ column, knowing E makes F *less* likely when the number m of tosses is small, and *more* likely when m is big. This is the dependence we expected. But $m = 4$ is an exception and here the two events just happen to be independent. It's this kind of pitfall that makes trusting your intuition about independence so dangerous.

Here are a couple of similar exercises involving dice. In both cases, the events involve the total on the two dice and the number on one die. We expect the two to be dependent since having a small or large number showing on one die ought to raise or lower the expected total. Once again, this is usually, but not always, correct. But in one case there's an exceptional total for all numbers, and in the other there's a different exceptional total for each number. So not only are there exceptions, but there's not even any clear pattern to the exceptions.

PROBLEM 4.7.10: We roll two dice.

- i) Consider the events E = “blue die shows 5”, F = “total is 4”, G = “total is 7” and H = “total is 10”. Which of the events F , G and H are independent of E ?
- ii) Which of the events F , G and H above are independent of $E' =$ “blue die shows 3”?

PROBLEM 4.7.11: We roll two dice.

- i) Consider the events E = “total is at most 9”, F = “red die is 2”, G = “red die is 4” and H = “red die is 6”. Which if the events F , G and H are independent of E ?
- ii) Which of the events F , G and H above are independent of $E' =$ “total is at least 4 and at most 10”?

Knowing that events are independent can be very useful, because then we can use [INDEPENDENT EVENTS 4.7.1.i](#)) to calculate intersection probabilities:



4.7 The dice don't talk to each other

INTERSECTION PROBABILITY OF INDEPENDENT EVENTS 4.7.12: *If E and F are known to be independent, then $\Pr(E \cap F) = \Pr(E) \cdot \Pr(F)$. More generally, the probability of the intersection of any number—3, or 4, or n —of independent events, is the product of the probabilities of the individual events.*

Warning: You must somehow *know in advance* that E and F are independent to apply this formula. If not, you *must* use **INDEPENDENT EVENTS 4.7.1.i**), $\Pr(E \cap F) = \Pr(F) \cdot \Pr(E|F)$. Ignoring this warning is a very common cause of mistakes, especially in problems that ask to decide whether or not two events are independent, as I noted after the **METHOD FOR CHECKING INDEPENDENCE 4.7.2**.

Here are a couple easy exercises in using the **INTERSECTION PROBABILITY OF INDEPENDENT EVENTS 4.7.12**. The whole subsection that follows can be viewed as a much more important application.

PROBLEM 4.7.13: Suppose that we know that student's performance in math courses is independent of his or her gender. If 20% of students in math classes get A grades and 60% of students in math classes are women, what percent of students are women with A grades? What percent are men with grades of B or lower?

The **INTERSECTION PROBABILITY OF INDEPENDENT EVENTS 4.7.12** can also be used, though less frequently to solve for an unknown simple probability. Here's an example.

PROBLEM 4.7.14: Suppose that we know that student's performance in math courses is independent of his or her gender. If a class of 60 students contains 40 women and 10 of these women got A's, how many A's were awarded in total? If 8 men got B's, how many B's were awarded in total?

Here's a problem that uses independence of more than 2 events.

PROBLEM 4.7.15: Suppose we perform the experiment of tossing a fair coin 3 times and record which side comes up on each toss. Let

4.7 The dice don't talk to each other

H_i denote the event “heads on the i^{th} toss” and T_i denote the event “tails on the i^{th} toss”.

Showing, by counting outcomes in each case, that:

i) $\Pr(H_i) = \Pr(T_i) = \frac{1}{2}$.

ii) $\Pr(H_1 \cap H_2) = \Pr(H_1 \cap T_2) = \frac{1}{4}$.

iii) if i and j are different, then $\Pr(H_i \cap H_j) = \Pr(H_i \cap T_j) = \frac{1}{4}$.

We can sum up what's been checked so far as saying that, when i and j are different tosses, the events H_i and H_j , or H_i and T_j , are independent.

Show that $\Pr(H_1 \cap H_2 \cap T_3) = \Pr(T_1 \cap H_2 \cap T_3) = \frac{1}{8}$,

i) by showing there is 1 outcome in each event.

ii) by using independence to write the probability as a product of 3 simple probabilities.

PROBLEM 4.7.16: Suppose we perform the experiment of tossing a fair coin 5 times and record which side comes up on each toss. Let H_i denote the event “heads on the i^{th} toss” and T_i denote the event “tails on the i^{th} toss”.

Show that $\Pr(H_1 \cap H_2 \cap T_5) = \Pr(T_1 \cap H_2 \cap T_5) = \frac{1}{8}$

i) by counting outcomes

ii) by using independence to write the probability as a product of 3 simple probabilities.

Show that $\Pr(H_1 \cap H_2 \cap T_3 \cap T_5) = \frac{1}{16}$

i) by counting outcomes

ii) by using independence to write the probability as a product.

PROBLEM 4.7.17: Suppose that we know that student's performance in math courses is independent of his or her gender. If a class of 60 students contains 40 women and 10 of these women got A's, how many A's were awarded in total? If 8 men got B's, how many B's were awarded in total?

Applying independence: the binomial distribution

In this section, we draw together standard counting techniques and the notion of independence to understand very completely, a special but very important class of probability spaces, the binomial distributions. In later sections, we'll go on to use these to peek across the boundary between probability and statistics. To take such a peek for most distributions, we'd have to do a lot more work. You'll find out what's involved if you ever take a statistics course. But for the binomial distributions, our understanding of the probability side of the picture will be so complete that we'll be able to answer interesting questions without these statistical foundations.

We start from something that couldn't be simpler, called a **Bernoulli trial**. You can think of this as a flip of an *unfair* coin—meaning that the two sides do not come up equally often—and where, instead of denoting the two outcomes by H and T, we call them “**success**” and “**failure**”, abbreviated **s** and **f**. The notions of success and failure do not imply any kind of value judgement, but are in a metaphorical way, simply to have a standard way of distinguishing the 2 outcomes. For example, when our trial consists of performing a BSE test on a cow, we'll call the test “successful” if the cow tests **positive** for BSE. So,

BERNOULLI TRIAL 4.7.18: *A Bernoulli trial B is an experiment whose sample space consists of just 2 outcomes, **s** and **f**.*

Of course the trial B , like tossing a single coin, is too simple to be interesting. But, as with coin tossing, we can use it to build an interesting experiment simply by repeating the trial.

BINOMIAL EXPERIMENT 4.7.19: *For $n > 0$, a binomial experiment is just a sequence of n equivalent Bernoulli trials B and the binomial sample space \mathcal{B}^n the set of possible outcomes of such an experiment. In other words, the sample space of \mathcal{B}^n consists sequences or words of length n in the 2 outcomes, **s** and **f**, so $\#\mathcal{B}^n = 2^n$.*

4.7 The dice don't talk to each other

Note that this sample space only depends on the number n of trials. We have yet to introduce any probabilities in these trials, although we'll do so in a moment. Why use n instead of our standard letter l for the number of trials—or the length of the sequences? This letter n is the one you'll see used to describe binomial experiments in pretty much all references, so it's easier to get used to it right away. Some of the events (and counts) we have looked at in tossing coins transfer without change to any binomial experiment.

BINOMIAL EVENTS 4.7.20: For $0 \leq k \leq n$, we denote by $S^{i,n}$ (and $F^{i,n}$) the binomial event of observing s (and f) on the i^{th} of the n trials, and by $E^{k,n}$ the of observing s in exactly k of the n trials. What about the other trials? For $S^{i,n}$ (and $F^{i,n}$), their outcomes can be either s or f . Since there are only 2 outcomes, they must be f 's for $S^{i,n}$ and s 's for $F^{i,n}$. So we could equally well say that outcomes in $E^{k,n}$ are those where we observe exactly k s 's and $(n - k)$ f 's.

It's easy to count these events:

PROBLEM 4.7.21:

- Show that $\#S^{k,n} = \#F^{k,n} = 2^{n-1}$. Hint: See [PROBLEM 4.3.6](#).
- Show that $\#E^{k,n} = C(n, k)$. Hint: The reasoning in the coin tossing [EXAMPLE 3.6.15](#) applies unchanged (and gives the same answer). Hint: The reasoning in the coin tossing [EXAMPLE 3.6.15](#) applies unchanged (and gives the same answer).

We're now ready to introduce probabilities on these sample spaces. Unlike pretty much all the distributions we have worked with up to this point, outcomes are *not* going to be equally likely. How is this possible given that the spaces S_n get big very quickly? Since $\#B^n = 2^n$, already B^{20} contains over a million outcomes—and we're going to want to think about spaces $B^{40,000}$ whose order has over 10,000 digits!

This is where we use **independence** is an essential, and very powerful, way. We'll be able to completely specify the probability distributions that come up by a *single number* p , even though individual



4.7 The dice don't talk to each other

outcomes have lots of *different* probabilities. The upshot is that binomial distributions are not much harder to work with than equally likely outcomes distributions—which, remember, are also specified by a single number, the order of the sample space.

What is this magic number p and how can it tell us about all the outcomes? For the single trial B , it's easy. We ask that $\Pr(\mathbf{s}) = p$ —so our magic p is just the probability of success in a single trial—and we're done. The total probability 1 condition [PROBABILITY MEASURE 4.2.1.ii](#) forces us to have $\Pr(\mathbf{s}) + \Pr(\mathbf{f}) = 1$, so we have to set $\Pr(\mathbf{f}) = (1 - p)$. It's standard to set $q := 1 - p$ to simplify the formulae that are coming, but this is just a convenience: if we know p , we know q .

We'll write B_p to denote the sample space B with this choice of probability distribution. Notice that in B_p , we *almost never* have $\Pr(\mathbf{s}) = \Pr(\mathbf{f})$ —that is, equally likely outcomes. In fact, this *only* happens when $p = q = \frac{1}{2}$, the case of tossing a fair coin. So we can't hope to have equally likely outcomes for any related probability on the bigger binomial sample spaces B^n .

This is where we need to appeal to independence. Suppose we know that the events $S^{i,n}$ and $S^{j,n}$ (\mathbf{s} on the i^{th} of our n trials and on the j^{th}) are all **independent**? Then it turns out that we can easily find the probability of any of the 2^n outcomes—sequences of n \mathbf{s} 's and \mathbf{f} 's—in B^n . Further, this probability turns out to depend *only* on how many of the n trails are successes. This makes it easy for find probabilities of events like $E^{k,n}$ where this number of successes is fixed to be k . Let's put off writing down these formulae for a moment and focus on the assumption of independence.

Why—or better, when—should we think the events $S^{i,n}$ (or $F^{i,n}$) and $S^{j,n}$ (or $F^{j,n}$), \mathbf{s} (or \mathbf{f}) on the i^{th} and j^{th} of our n trials, are all independent? In mathematical situations, independence is often clear.

If our trials consist of repeatedly tossing a coin, then the fact that “the coin has no memory” makes the outcomes of any two tosses independent. It doesn't matter what happened on the first toss; on

4.7 The dice don't talk to each other

the second, we expect the coin to come up heads half the time and tails the other half. This is the conditional probability way of saying the tosses are independent in part ii) of the definition of **INDEPENDENT EVENTS** 4.7.1: $\Pr(\text{"heads on second"}|\text{"heads on first"}) = \Pr(\text{"heads on second"}|\text{"tails on first"}) = \Pr(\text{"heads on second"}) = \frac{1}{2}$.

This applies equally to any two tosses: $\Pr(\text{"heads on } j^{\text{th}}|\text{"heads on } i^{\text{th}}}) = \Pr(\text{"heads on } j^{\text{th}}|\text{"tails on } i^{\text{th}}}) = \Pr(\text{"heads on } j^{\text{th}}}) = \frac{1}{2}$. In terms of S and F events, this is just $\Pr(S^{j,n}|S^{i,n}) = \Pr(S^{j,n}|F^{i,n}) = \Pr(S^{j,n}) = \frac{1}{2}$.

For a second example—and one that will show that there's no need to have $p = \frac{1}{2}$ as with the coins—consider a trial consisting of rolling a single die in which we define S to be rolling a 2, so $p = \frac{1}{6}$, and failure to be observing any other roll and where our experiment consists of rolling n dice of different colors. Once again, it doesn't matter what happened on the blue (or i^{th}) roll. Because, "the dice don't talk", we expect to see a 2 on $\frac{1}{6}^{\text{th}}$ of the red (or j^{th}) rolls and this is again the conditional probability way of saying the outcomes of the rolls or trials are independent.

This dice example highlights two other ideas that come up in using binomial distributions. First, we asked for our trials to be **equivalent** rather than identical: what this means is is most easily explained by referring to these examples. Tosses of same fair quarter or rolls of the same fair die or BSE tests of the same cow are identical. Tosses of two different fair quarters or rolls of two different fair dice or BSE tests of the two different cows are equivalent: we view what happens as unchanged by the substitution of coin, die or cow.

Second, we often start with an experiment with a sample space S with many outcomes (so *not* a Bernoulli trial), and concoct a Bernoulli trial from it by picking an event E and equating S with observing E and f with not observing E (i.e. observing E^C). If, as is common, S comes equipped with a probability distribution, then we also get our parameter p by setting $p = \Pr_S(E)$.



4.7 The dice don't talk to each other

For example, suppose we are interested in a game like craps which involves rolling a *pair* of dice and want to focus on the event E of rolling a total of 7 or 11. Then observing either of these totals constitutes s in our trial (any other totals are f) and $p = \frac{6}{36} + \frac{2}{36} = \frac{2}{9}$. Once again, we expect the outcomes of successive trials to be independent but this time, not because “the dice don’t talk” (since our definition of success in a single trial involves the total), but because “they have no memory” (outcomes of prior rolls have no effect on outcomes of future ones).

Let’s return to our assumptions of independence. This becomes harder to justify when our trials have a more practical character. Often we will not know the probability p of success in a single trial and we’ll want to use the formulae below to estimate this from a number of trials. For example, if we are testing cows for BSE (and consider finding an infected cow to be an s), then p is the fraction of cows with the disease. We want to estimate this by testing many cows and looking at how many infected ones we find. But to do so we must assume that the probability that the next cow I test will be infected is *not* affected by the fact that the last cow I tested *was*. Now there’s no mathematical “coin has no memory/dice don’t talk” reason to expect this. We are making some choices and the question is whether or not we are truly making them “at random”—with each cow equally likely to be picked—or whether our method for selecting cows for testing has some bias which we may not understand or even suspect.

Moreover, in such real world cases, it’s often hard to know that the trials we perform really represent choices made “at random” from the population we’d like to survey. To get a feel for the way biases can creep in uninvited, consider trying to decide the fraction of electors (our p) who will vote in an upcoming election by conducting a survey. We can view each survey participant as a trial with s defined as “plans to vote”. Here’s one simple way in which your trials might

4.7 The dice don't talk to each other

be dependent.

EXAMPLE 4.7.22: You conduct your survey by going door-to-door and questioning anyone who responds to your knocks. Suppose that the survey population lives in two towns E and F with equal populations but that 60% of voters in E plan to turn out while only 40% of voters in F plan to turn out. Then the p we are looking for is 50% or $\frac{1}{2}$. If our trials were dependent then the probability of find two voters who plan to turn out in a row (S in 2 consecutive trials) would be 25% or $\frac{1}{4}$, as for tossing two coins. But if we are canvassing “at random” in town E , we are, in effect, making independent choices with a $p = .6$ and we’ll see S twice in a row $.6 \times .6$ or 36% of the time while if we are canvassing “at random” in F , we’ll see it $.4 \times .4$ or 16% of the time.

Even if successive trials are independent they can lead to incorrect guesses for p . Suppose in our example that you’re aware of the difference between the populations of towns E and F so you decide to conduct a phone survey instead, randomly choosing whether to dial a number from the E or F directory each time. What if 60% of people with 9 to 5 jobs plan to vote, while only 40% of those who are at home during the day do? If you simply tabulate the results of calls made from 9 to 5, what you’ll think is that 40% of voters plan to turn out. But when you check the chance that two successive respondents plan to vote, you’ll get $.4 \times .4$ or 16%. So your trials this time are independent, but your guess for p is wrong because of your method produced a sample that was not truly random in a way that you were not able to predict.

The last problem is a real one for those conducting phone surveys. Typically, such surveys attempt 15 calls (at varying times and days) to a number before it is discarded. And even then, there are biases from people with no phones, or who have only a cell phone and do not appear in phone directories.

The bad news is that there’s no magic bullet for ensuring that a se-



4.7 The dice don't talk to each other

quence of trials is truly being made “at random” or that the results of different trials are independent. One of the aims of statistics is to provide methods for designing trials to try to avoid such biases and tools for checking assumptions of independence. But, even using these, it's often hard to really be sure that such assumptions about your data are justified.

There are 2 pieces of good news. First, here, we only want to understand what we can learn when such assumptions *do* hold, so you can simply take them on trust in the examples we'll work. Second, when they do hold, the formulas they lead to are very simple.

BINOMIAL DISTRIBUTION FORMULA 4.7.23: *The binomial distribution with probability p is the unique probability distribution on the sample space \mathcal{B}^n for which $\Pr(S^{i,n}) = p$ (and $\Pr(F^{i,n}) = q$) and for which the events $S^{i,n}$ (or $F^{i,n}$) and $S^{j,n}$ (or $F^{j,n}$), s (or f) on the i^{th} and j^{th} of our n trials, are all independent. For this probability distribution:*

- i) *The probability of every individual outcome (sequence of n trials), in which we observe exactly k s s and $(n - k)$ f s equals $p^k \cdot q^{n-k}$.*
- ii) *The probability of the event $E^{k,n}$ that we observe some outcome with exactly k s s and $(n - k)$ f s is given by $\Pr(E^{k,n}) = C(n, k) \cdot p^k \cdot q^{n-k}$.*

The formula for outcomes in i) is where we rely on independence. For example, suppose that we're rolling dice and s is a total of 7 or 11, so $p = \frac{2}{9}$ and $q = \frac{7}{9}$, and we roll the dice $n = 5$ times.

What's the chance we'll observe the outcome $ffsf s$ (that is, a total of 7 or 11 on the third and fifth rolls and not on the first and second and fourth)? What all those “ands” mean is that this outcome is just the intersection of the events $F^{1,5} \cap F^{2,5} \cap S^{3,5} \cap F^{4,5} \cap S^{5,5}$. Since the events being intersected are independent, the probability of the intersection is just the product of the individual probabilities $p = \frac{2}{9}$ for each S and $q = \frac{7}{9}$ for each F . So we get

$$\Pr(ffsf s) = \frac{7}{9} \cdot \frac{7}{9} \cdot \frac{2}{9} \cdot \frac{7}{9} \cdot \frac{2}{9} = \left(\frac{7}{9}\right)^3 \left(\frac{2}{9}\right)^2 = p^2 \cdot q^3.$$

4.7 The dice don't talk to each other

But the same reasoning says that

$$\Pr(\text{fsfsf}) = \frac{7}{9} \cdot \frac{2}{9} \cdot \frac{7}{9} \cdot \frac{2}{9} \cdot \frac{7}{9} = \left(\frac{2}{9}\right)^2 \left(\frac{7}{9}\right)^3.$$

In fact, anytime we observe an outcome that involves 2 ss and 3 fs, we'll get a product with 2 factors of $p = \frac{2}{9}$ and 3 of $q = \frac{7}{9}$ and these will collect to give $\left(\frac{2}{9}\right)^2 \left(\frac{7}{9}\right)^3 = p^2 \cdot q^3$.

If, instead of 2 ss and 3 fs, we have k and $(5 - k)$, then we'll see k factors of $p = \frac{2}{9}$ and $5 - k$ of $q = \frac{7}{9}$ and these will collect to give $p^k \cdot q^{5-k}$.

PROBLEM 4.7.24: Show that, for our dice example, $\Pr(\text{s s s f s}) = \left(\frac{2}{9}\right)^4 \left(\frac{7}{9}\right)^1 = p^4 \cdot q^1$.

Finally, if we see k ss in n rolls rather than 5, then we'll see $n - k$ fs. So we'll get k factors of $p = \frac{2}{9}$ and $n - k$ of $q = \frac{7}{9}$ that will collect to give $p^k \cdot q^{n-k}$.

PROBLEM 4.7.25: Show that, for our dice example:

- i) $\Pr(\text{ssff}) = \left(\frac{2}{9}\right)^2 \left(\frac{7}{9}\right)^2 = p^2 \cdot q^2$.
- ii) $\Pr(\text{sffsff}) = \left(\frac{2}{9}\right)^2 \left(\frac{7}{9}\right)^4 = p^2 \cdot q^4$.

The formula in ii) for $\Pr(E^{k,n})$ now follows easily. Because $E^{k,n}$ consists of *all* outcomes with exactly k ss and these outcomes all have the *same* probability $p^k \cdot q^{n-k}$, we know that $\Pr(E^{k,n}) = \#(E^{k,n}) \cdot p^k \cdot q^{n-k}$. So we just have to check that the *number* of n letter sequences in s and f with exactly k ss is $C(n, k)$. We've seen this on many occasions (for example, [EXAMPLE 3.6.15](#) or [PROBLEM 3.6.16](#) or [PROBLEM 3.8.27](#)).

Before we start using the formula, I should point out that, although the answers we'll get could be written as fractions, both the numerators and denominators get pretty big. Too big, for most calculators—though we could work around this using symbolic calculation software. Further, it's easier to compute the powers in these formulas in decimal form. Finally, and most importantly, we're going to be more interested in how big, or how small, these probabilities are than in

4.7 The dice don't talk to each other

the counts that give numerator and denominator. So we'll make all our calculations with this formula in decimal form. How should we round these answers? I have chosen to keep 3 significant digits in my *final* answers, and you can do the same: this is plenty of accuracy to answer “how big? how small?” questions and saves writing down a sea of decimals. There's one proviso: when you enter p and q into your calculator, enter them as fractions—for example enter $\frac{1}{3}$ as (1/3) and *not* as .333. Why? So we will not violate the [FIRST RULE OF ROUNDING 1.2.4](#) by rounding an intermediate value. If you don't take this precaution, your final answer can often be substantially off (and, of course, you won't know it)..

Let's do a few easy problems just to get used to the formula. We could have worked the following example with coins using just the [EQUALLY LIKELY OUTCOMES FORMULAE FOR PROBABILITIES 4.3.3](#): since $p = q = \frac{1}{2}$ all sequences of n tosses have probability $\frac{1}{2^n}$. But it's a lot easier to just plug into the [BINOMIAL DISTRIBUTION FORMULA 4.7.23](#).

EXAMPLE 4.7.26: If we toss a fair coin repeatedly, what is the chance of seeing exactly:

- i) 5 heads in 10 tosses.
- ii) 50 tails in 100 tosses.

Solution

Our Bernoulli trial has $p = q = \frac{1}{2}$. Let's say that a head is a S.

- i) Here we want to know $\Pr(E^{5,10})$ —the chance of see exactly 5 heads or Ss in 10 tosses or trials so we plug in to find

$$\Pr(E^{k,n}) = C(n, k) \cdot p^k \cdot q^{n-k} = C(10, 5) \cdot \frac{1}{2}^5 \frac{1}{2}^{(10-5)} = 0.246.$$

Just in case you had difficulty keying those powers in here are the keystrokes I used to get $\frac{1}{2}^5 \frac{1}{2}^{(10-5)}$: $((1/2)^5)*((1/2)^{(10-5)})$. Did really need *all* those parentheses? No, but I needed most of them. If you're *sure* which ones are superfluous, feel free to leave them out. But it's much safer to remember and follow the [PAREN-](#)

4.7 The dice don't talk to each other

THESES RULE! 1.1.2. For the record, you need the blue ones in $((1/2)^5) \cdot ((1/2)^{(10-5)})$.

ii) Here we want to know $\Pr(E^{50,100})$ and we get

$$\Pr(E^{k,n}) = C(n,k) \cdot p^k \cdot q^{n-k} = C(100,50) \cdot \frac{1}{2}^{50} \frac{1}{2}^{(100-50)} = 0.0796.$$

In other words, while expect to get *about* 5 heads in 10 tosses and about 50 in 100 tosses, the chances of getting exactly that many heads are not that high—25% and 8%—and get smaller the *more* tosses we perform. For example, in 10,000 tosses we will only see exactly 5000 heads about 2.5% of the time (this one really needs a symbolic algebra system).

We can also use the **BINOMIAL DISTRIBUTION FORMULA 4.7.23** to find the chance of “at least” or “at most” a specified number of wins by the familiar device of summing the probabilities for each exact number of wins covered. Here’s an example.

EXAMPLE 4.7.27: We are playing craps and repeatedly roll two dice to start a game repeatedly. Find the probability of seeing the winning total of 7 or 11:

- i) exactly 2 times in 9 games.
- ii) exactly 8 times in 9 games.
- iii) at most 2 times in 9 games.
- iv) at least 8 times in 9 games.
- v) at most 8 times in 9 games.

Solution

In this example, we call a total of 7 or 11 a **s** so $p = \frac{2}{9}$ and $q = \frac{7}{9}$.

- i) Here we just plug in: $\Pr(E^{2,9}) = C(9,2) \cdot \frac{2}{9}^2 \frac{7}{9}^{(9-2)} = 0.3061020142$ or about 30%.
- ii) Again, we want $\Pr(E^{8,9}) = C(9,8) \cdot \frac{2}{9}^8 \frac{7}{9}^{(9-8)} = 0.0000416291872$, or barely 1 in 25,000.
- iii) Here we need to sum the chances of seeing exactly 0 or 1 or 2 successes, getting $C(9,0) \cdot \frac{2}{9}^0 \frac{7}{9}^{(9-0)} + C(9,1) \cdot \frac{2}{9}^1 \frac{7}{9}^{(9-1)} + C(9,2) \cdot \frac{2}{9}^2 \frac{7}{9}^{(9-2)} = 0.6781009898$.

4.7 The dice don't talk to each other

iv) Here we need to sum the chances of seeing exactly 8 or 9 successes, getting

$$C(9, 8) \cdot \frac{2^8 7^{(9-8)}}{9^9} + C(9, 9) \cdot \frac{2^9 7^{(9-9)}}{9^9} = 0.00004295074853.$$

v) Here we *could* sum the chances of seeing exactly k successes for k running from 0 to 8, but it's a lot less work to subtract the complementary chance of getting 9 successes from 1, getting $1 - C(9, 9) \cdot \frac{2^9 7^{(9-9)}}{9^9} = 0.9999986784$.

Here are a few for you to try.

PROBLEM 4.7.28: You are taking a test with 20 multiple choice questions, each of which has 4 answers. If you guess answers at random, what's the chance that:

- i) you'll get all 20 of the questions right.
- ii) you'll get none of the questions right.
- iii) you'll get exactly 5 of the questions right. (This is the number of right answers we'd "expect" to guess).
- iv) you'll get at most 5 of the questions right.
- v) you'll get at least 5 of the questions right. Hint: This is not quite complement to "at most 5 right", but if you use the answer to [iii](#)) as well, you do not need any further work.

One way the [BINOMIAL DISTRIBUTION FORMULA 4.7.23](#) is often used is to decide how much testing is needed to uncover flaws in a process. Here's an example that illustrates how this works.

You are responsible for quality control at a flat-screen TV factory. You know that fewer than 1 in a 5000 TVs have this problem, but you also know that, on rare occasions, TVs with bad pixels occur in big bunches—1 defect in 20 or 25—because one machine used to make the LCDs has a calibration problem. You need to design a testing protocol that catches these calibration problems.

Each day your factory produces 5,000 TVs. Of course, you could assure 0 defects by simply testing *all* the TVs you make. But it takes a worker 4 minutes to test each TV so each can only test about 120

4.7 The dice don't talk to each other

a day, and you can't spare the 40+ workers it would take to test every TV, especially you would only catch one defect on an average day. If you could be 99% certain that at most 2% of the TVs in each day's run have a dead pixel, you wouldn't be worried about calibration problems. How many TVs do you need to test each day to be this confident if no defects are found?

EXAMPLE 4.7.29: We can think of each test as a Bernoulli trial where p is the chance of finding a dead pixel. We want to pick the number of trials n (how many TVs we test) so that the chance of seeing at least 1 S in n trials is at least 99%. Equivalently, we want to chance of seeing 0 TVs with defects to be less than $1\% = .01$. This is just $\Pr(E^{0,n}) = C(n, 0) \cdot p^0 \cdot q^n = q^n$. As p gets bigger, q gets smaller and so does q^n . So if testing n TVs and seeing no defects gives us enough confidence when $p = .02$, we'll be even more certain of catching days when p is bigger.

So we want to need to find an n for which $.98^n < .01$. We can either do this by repeated guessing, or use the \ln function. Since [ln IS INCREASING 1.4.44](#), we can ask $\ln(.98^n) < \ln(.01)$ and since $\ln(x^n) = n \cdot \ln(x)$, this means $n > \frac{\ln(.01)}{\ln(.98)} = 227.9481712$ so $n = 228$ works. As a check, $.98^{228} = 0.009989534656$. In other words, we just need 2 testers.

Here are some claims about how such a protocol would function that I'll leave for you to check. Most days (roughly 19 in every 20) will pass with no defects being found. If you *do* find a defective TV, you'll need to start worrying about calibrations. But if you find just 1 defective TV, you can still be pretty sure that things are running normally. This happens less than 6% of the time when the defect rate is at least 2% and almost 99.9% of the time when the defect rate is $\frac{1}{5000}$. Therefore, if you find 2 or more defective TVs, you can be pretty sure some calibration is needed. You'll observe this only 1 day in a 1000 when the defect rate is $\frac{1}{5000}$. In other words, 0, 1 and "more than 1" defects found correspond with high probability to "no problem", "problem possible but not too likely" and "problem almost certain". Moreover,

4.7 The dice don't talk to each other

almost 95% of the time that a machine needs calibration, your testing program will reveal at least 2 defects.

PROBLEM 4.7.30: Use the [BINOMIAL DISTRIBUTION FORMULA 4.7.23](#) to check each of the following:

- i) If the probability of a defect is $\frac{1}{5000}$ then the chance of seeing 0 defects when 228 TVs are tested is 0.9554196981.
- ii) If the probability of a defect is $\frac{1}{5000}$ then the chance of seeing at most 1 defect when 228 TVs are tested is 0.9989955515.
- iii) If the probability of a defect is $\frac{1}{5000}$ then the chance of seeing 2 or more defects when 228 TVs are tested is 0.0010044485—barely 1 in a 1000.
- iv) If the probability of a defect is 2% then the chance of seeing at most 1 defect when 228 TVs are tested is 0.05647145102.
- v) If the probability of a defect is 2% then the chance of seeing 2 or more defects when 228 TVs are tested is 0.9435285489.

BSE Testing

We're now ready to tackle a more substantial application of the [BINOMIAL DISTRIBUTION FORMULA 4.7.23](#), using it to confirm the analysis of the U.S.D.A.'s BSE testing program that we gave in [SECTION 2.2](#). Our goal here is to check all the numbers in [TABLE 2.2.2](#) so you might want to reread this section before starting. I have broken this up into three problems, one for each of the three numbers of cows tested per year. I'll do the first as an example and leave the other two as problems for you.

The basic idea is the same in all three parts. We view each BSE test as a Bernoulli trial in which finding an infected cow is s . The probability p of success is then prevalence of BSE in the American herd—that is, the fraction of cattle infected with BSE. We don't know what this fraction is: the main point I tried to make in [SECTION 2.2](#) is that, while getting a handle on this number is *supposed* to be the purpose

4.7 The dice don't talk to each other

of the test, they are very unlikely to provide good estimate. But, we can ask what results we'd expect if we *assume* that p takes various values. The two that I chose to work with are $p = 0.000001$, the 1 in 1,000,000 prevalence that the testing program is supposed to rule out, and $p = 0.00001$, the much higher 1 in 100,000 prevalence at which we'd be consuming about 150,000 pounds of meat from infected cattle every year.

To apply the [BINOMIAL DISTRIBUTION FORMULA 4.7.23](#), we also need to know the number of trials n —this is just the number of cows tested each year and is the only ingredient that will vary in analyzing the 3 subtables—and the number of successes k , which is given in the first column of each table. We can then find the entries in the “exactly k infected” columns by just plugging in to $\Pr(E^{k,n}) = C(n,k) \cdot p^k \cdot q^{n-k}$.

And we can find the entries in the “more than k infected” columns, by applying the [COMPLEMENT FORMULA FOR PROBABILITIES 4.2.3](#). In other words, starting at $k = 0$ we successively subtract the probability of seeing “exactly k ” infected cattle from the probability of seeing “more than $(k - 1)$ ” to get the probability of seeing “more than k ”. This is so straightforward I won't bother to give the calculations.

EXAMPLE 4.7.31: Here's the first set of probabilities with $n = 20,000$.

Prevalence	1 in 100,000		1 in 1,000,000	
20,000 cattle tested	Probability that number of infected cattle is			
k	exactly k	more than k	exactly k	more than k
0	0.818730	0.181270	0.980199	0.019801
1	0.163748	0.017522	0.019603	0.000197
2	0.016374	0.001148	0.000196	0.000001
3	0.001092	0.000057	0.000001	0.000000006

TABLE 4.7.32: TESTING 20,000 CATTLE YEARLY FOR BSE

4.7 The dice don't talk to each other

For $p = 0.00001$ and $k = 0$, we can just plug in to check that $\Pr(E^{k,n}) = C(n,k) \cdot p^k \cdot q^{n-k} = C(20000,0)(0.00001)^0(0.99999)^{20000} = 0.818730$. The second and third numbers in this column are

$$C(20000,1)(0.00001)^1(0.99999)^{19999} = 0.163748$$

and

$$C(20000,2)(0.00001)^2(0.99999)^{19998} = 0.016374$$

respectively.

Likewise we can get the first 2 numbers in the column where $p = 0.000001$ (and $q = 0.999999$) as

$$C(20000,0)(0.000001)^0(0.999999)^{20000} = 0.980199$$

and

$$C(20000,1)(0.000001)^1(0.999999)^{19999} = 0.019603.$$

PROBLEM 4.7.33: Verify the remaining values 0.000196, 0.001092 and 0.000001 in [TABLE 4.7.32](#).

PROBLEM 4.7.34: Proceed as in [Example 4.7.31](#), but with $n = 360,000$ to verify the values below:

Prevalence	1 in 100,000		1 in 1,000,000	
360,000 cattle tested	Probability that number of infected cattle is			
k	exactly k	more than k	exactly k	more than k
0	0.027323	0.972677	0.697676	0.302324
1	0.098365	0.874312	0.251164	0.051160
2	0.177058	0.697255	0.045209	0.005950
3	0.212470	0.484784	0.005425	0.000526

TABLE 4.7.35: TESTING 360,000 CATTLE YEARLY FOR BSE

PROBLEM 4.7.36: Verify the final set of values:

4.7 The dice don't talk to each other

Prevalence	1 in 100,000		1 in 1,000,000	
40,000 cattle tested	Probability that number of infected cattle is			
k	exactly k	more than k	exactly k	more than k
0	0.670319	0.329681	0.960789	0.039210
1	0.268130	0.061551	0.038431	0.000778
2	0.053625	0.007926	0.000769	0.000010
3	0.007150	0.000776	0.000010	0.0000001

TABLE 4.7.37: TESTING 40,000 CATTLE YEARLY FOR BSE

What's wrong with this and how can we fix it? To have a sensible testing program, we first need to decide what prevalence we wish to rule out and how certain we demand to be about the actual prevalence being lower. We could require complete certainty but then we'd need to test almost every cow. Instead, our goal here is to test a few as possible, while providing the necessary certainty.

Let's say, to keep things simple, that we want a program in which not detecting a case of BSE in a calendar year makes us 95% certain that the prevalence of BSE is *less* than 1 in a 100000. The theme of the preceding problems is that you can't make such a claim by giving just 40,000 tests. Quite the opposite. If tests are independent and the prevalence equals 1 in a 100000, there's still a 67% chance that we'll see no positive tests.

How do we decide how many tests *are* enough? The key is to restate our demand for certainty: "If the prevalence equals 1 in a 100000, how many cows do we need to test to ensure that the chance of finding 0 infected cows is *less* than 5%?" I'll call this an "uncertainty threshold" of 5%

We can answer this question the same way we made the tables above, with the [BINOMIAL DISTRIBUTION FORMULA 4.7.23](#) formula. Write n for the number of cows. Then we're asking about $Pr(E)$ where

4.7 The dice don't talk to each other

the event E is “exactly 0 successes (infected cows) in n binomial trials (BSE tests), when $p = \frac{1}{100,000}$ ”. Assuming that the tests are independent—a question about how the cows are selected that we won't consider—it's given by $C(n, 0)p^0q^n$. But $C(n, 0) = p^0 = 1$ and $q = 1 - p$. So we can rewrite this probability as $(1 - p)^n$.

Now comes a point at which a formula from an apparently unrelated area of mathematics can help. It's a good example that mathematics involves knowledge as well as thought. The formula is the **BERNOULLI'S LIMIT FOR exp 1.4.56**: $(1 + \frac{1}{n})^n \rightarrow e$.

PROBLEM 4.7.38: Use **BERNOULLI'S LIMIT FOR exp 1.4.56** and **RULES OF EXPONENTS 1.4.10** to show that $(1 - \frac{1}{n})^n \rightarrow e^{-1} = \frac{1}{e}$.

We can match this last form up by equating $p = \frac{1}{n}$ getting $(1 - \frac{1}{n})^n$. If n is big—we're testing a lot of cows—then this probability is close to $e^{-1} \approx 0.36787944$. To see what this means, let's take $p = \frac{1}{100,000}$ and hence $n = \frac{1}{p} = 100,000$. Suppose we test 100,000 cows for BSE and find 0 cases. Then we have seen an event that happens in only 36.7% of all years if that the current prevalence is greater than $\frac{1}{100,000}$. That meets our criterion for a testing program, except for the fact that 36.7% is still way above our 5% uncertainty threshold. We need to test more cows.

But in fact for *any* small p we get a 36.7% level of uncertainty by testing $n = \frac{1}{p}$ cows. To see how far to raise n , we introduce a parameter a in the numerator: $n = \frac{a}{p}$. This time the **BINOMIAL DISTRIBUTION FORMULA 4.7.23** gives us an uncertainty threshold of $(1 - p)^{\binom{a}{p}}$ and algebra like that in **PROBLEM 4.7.38** tells us that, if, as with BSE, p is small then this is close to $e^{(-a)}$.

So we'll meet our 5% uncertainty threshold if we have $e^{(-a)} < 0.05$. To solve for a we apply **exp AND ln ARE INVERSES 1.4.53** to get in turn $\ln(e^{(-a)}) < \ln(0.05)$, $-a < \ln(0.05)$ and $a > -\ln(0.05) = 2.99573$. So if $p = \frac{1}{100,000}$, we need to test at least 300,000 cows because $n = \frac{a}{p} = 2.99573 \cdot 100,000 = 299,573$. But because the bound for a does not

4.7 The dice don't talk to each other

depend on p we can say much more. If I want 5% or lower uncertainty and $p = \frac{1}{1,000,000}$, I need to test about 3,000,000 cows.

In fact, we can say even more by introducing another variable u for our uncertainty threshold. Replacing the 0.05's in the algebra above with u 's leads to $a > -\ln(u)$. Since the natural logarithm is a very slowly growing function, this means that it's relatively cheap to promise lower uncertainties.

Taking $p = \frac{1}{100,000}$ as an example, at $u = 10\%$ uncertainty we get $a = -\ln(0.1) = 2.30258$ so we need to test 230,000 cows. By testing another 70,000 (total 300,000), we lower the uncertainty to 5%. At 1% uncertainty, $a = -\ln(0.01) = 4.60517$.

PROBLEM 4.7.39: . Show that testing an additional 160,000 cows, or 460,000 in all, suffices to reach the 1% uncertainty level when $p = \frac{1}{100,000}$.

We need to test a few million cows before not seeing positive tests makes us confident that the prevalence is below the $p = \frac{1}{1,000,000}$ that seems to be the testing goal in most public health circles. This is a lot more than the U.S.D.A tests but only about 10% of the cattle slaughtered annually.

Binomial distributions are also a perfect place to get acquainted with a few of the basic tools of statistics: as we'll see in the next section, the [BINOMIAL DISTRIBUTION FORMULA 4.7.23](#) makes it possible to give formulas for basic quantities that can usually call for numerical calculations lengthier—much lengthier—than those done here.

Common misconceptions about independence

I just want to make a few comments here about the examples we looked at in [HE'S ON Fire!](#). If you skim over that discussion again, you'll remember that it dealt with **runs** in sequences in Bernoulli trials. A run of length ℓ is just a set of ℓ consecutive trials with



4.7 The dice don't talk to each other

the *same* outcome. For example, if we're tossing coins and we get 4 heads in a row, that's a run of length 4; likewise, 5 tails in a row is a run of length 5. Our first goal is to see how likely runs are to continue.

PROBLEM 4.7.40: First, let's consider runs of successive ss.

- i) Observing a run ℓ ss means we have a sequence of $n = \ell$ trials in which we see $k = \ell$ ss. Use this to show that the chance of seeing a run of length ℓ in the next ℓ trials is p^ℓ .
- ii) Suppose we have just observed a run of ℓ ss. Show that the probability that the run will continue on the next trial is p . Hint: The chance that the run will continue is $\Pr(S | \text{"run of length } \ell")$ and the event (" S on the $(\ell + 1)^{\text{st}}$ trail" \cap "run of length ℓ ") just means "run of length $(\ell + 1)$ " so you can use i) to find its probability.
- iii) If we have just observed a run of ℓ ss, show that the probability that the run will stop on the next trial is q . Hint: If it doesn't stop, it will continue.

PROBLEM 4.7.41: First, let's consider runs of successive fs. Observing a run ℓ fs means we have a sequence of $n = \ell$ trials in which we see $k = \ell$ fs. Argue as in the preceding problem to show that:

- i) The chance of seeing a run of ℓ fs is q^ℓ .
- ii) If we have just observed a run of ℓ fs, the probability that the run will continue on the next trial is q and the probability that the run will stop on the next trial is p .

What does this problem tell us? Simply that the chances of success and failure on the next trial are *always* p and q , regardless of how long (or short) a run has preceded this trial and regardless of whether this was a run of ss or fs. Another way to put this is that the results of the next trials are unaffected by the run leading up to it. And that's just what it means to say that the outcome of the next trial is independent of outcomes of the preceding ones.

When we're tossing coins, for example, the next toss is equally likely to be a head or a tail whatever the run or streak that preceded it. So



4.7 The dice don't talk to each other

streaks are *equally* likely to continue and to stop. When the chance of success is not $\frac{1}{2}$ —for example, in craps where $p = \frac{2}{9}$ —this last statement is no longer true without modification. The chance that a run of **ss** will continue is $p = \frac{2}{9}$ which is less than the chance $q = \frac{7}{9}$ that it will end. Conversely, the chance q that a run of **fs** will continue is more than the chance p that it will end. But it remains true that following a run of either **ss** or **fs**—of *any* length—the chances of observing **s** or **f** on *next* trial are always p and q respectively.

Very few people agree. Most people feel that the more consecutive heads we have tossed, the more we should expect a tail on the next throw. The belief that independence means “runs tend to stop” is what causes us to avoid long streaks when trying to simulate random sequences of **Hs** and **Ts**. But “runs tend to stop” is wrong because “The coin has no memory”: regardless of *what* has happened on the previous tosses, we should expect the next one to come up heads (and tails) half the time.

What's really odd about our intuition is that most people simultaneously believe even more strongly in a converse fallacy: when we observe runs, we conclude that successive trials must *not* be independent. This is the mechanism behind belief in the “hot hand”. When you watch a basketball game, you see players make (and miss) runs of shots in streaks, just as, when you toss a coin, you see runs of heads and tails. The parallel is not quite exact because the probability p of making a shot varies from player to player⁶, but that just means we're observing an unbalanced trial like craps instead of a coin toss.

Does observing these runs indicate that there's a “hot hand”? Is a player who's made his last shot (or two, or three), more likely to make his next. No. Many studies looking at whole seasons of shot by shot records solidly confirm this. A shooter's current run has no

⁶Most college and pro players make close to 45% of their shots—90% averaged between 36% and 54% in 2008-2009.

4.8 Who would have expected ...

influence on his chance of making his next shot. A 40% shooter will make his next shot 40% of the time; a 50% shooter will make hers 50% of the time.

The common argument is that shooters aren't coins and that it's the role of confidence in shooting that affects their success. Wrong. The math and the data both show this. I know many of you don't believe me. I won't try to convince you further, but if you'd like to buy a bridge...

4.8 Who would have expected ...

An example that we've tossed around many times in this chapter is that of an experiment in which a fair coin is tossed 100 times. We then expect to observe about 50 heads, because heads have probability $\frac{1}{2}$ and $50 = \frac{1}{2} \cdot 100$. In this section, we're going to see how to calculate such **expected values** systematically. This will allow us to deal both with more complicated experiments with big sample spaces, and to make predictions for more complicated quantities than just the frequency of a single outcome. Such predictions are fundamental in virtually all more advanced applications of probability. This is the subject of the first subsection, **RANDOM VARIABLES AND EXPECTED VALUES**.

We've also noted many times that actually observing exactly 50 heads in 100 tosses is rather rare. Instead, we expect the observed number of heads to be "close to" 50. But this answer only more questions. How close to 50 is "close"? In what sense is 56 "close" but 35 "far" from 50? In general, we do not expect observed values to exactly match expected values, and we'd like to know how small (or large) the gap between an observed value and an expected one needs to be to view the observation as confirming (or contradicting) the ex-

4.8 Who would have expected ...

pectation. This is a question of enormous practical importance. Indeed, the whole science of statistics is devoted to studying its many ramifications, and probability derives most of its application from its use in statistics. The second subsection provides just a first glimpse of how these questions are answered: it turns out that expected values are the key tool here as well.

Random variables and expected values

A **random variable** on a sample space S is just a function $Y : S \rightarrow \mathbb{R}$ that assigns a real number $Y(x)$ to every outcome x in S . We already know one important example: the probability distribution \Pr on S is a random variable. It assigns the number $\Pr(x)$ to each x . General random variables, however, are much more flexible: they need not obey the restrictions we impose on a probability distribution. We require that $\Pr(x)$ lie between 0 and 1 (because we think of it as the fraction of trials when we'll observe x) but we allow the value $Y(x)$ to be *any* real number—it can be negative, or bigger than 1 (even very large, as we'll see), and correspondingly we can't interpret a random variable in terms of frequencies. Similarly, while we require that the values $\Pr(x)$ for all $x \in S$ total exactly 1, there's no restriction on the sum of the values $Y(x)$. The fact that random variables can have arbitrary values means that, as we'll see in the examples to come, lots of quantities that we're interested in studying are random variables in disguise.

RANDOM VARIABLE 4.8.1: *A random variable Y on a sample space S is a real-valued function $Y : S \rightarrow \mathbb{R}$. That is, we associate to each outcome $x \in S$, a real number—any real number— $Y(x)$.*

Next, we want to look at how to condense the many values $Y(x)$ —there can be millions of them since there's one for each x in S —into a single expected value $\mathcal{E}(Y)$ that's a kind of executive summary of all this information. The way to think of an **expected value** $\mathcal{E}(Y)$ is



4.8 Who would have expected ...

as a probability average of Y . Informally, $\mathcal{E}(Y)$ should be value we'd obtain if we performed a large number n of trials of our experiment, recorded the n outcomes x that we observed, and averaged the corresponding n numbers $Y(x)$. Even more loosely, $\mathcal{E}(Y)$ should be the value of Y we'd observe on a mythical “average trial”. The formula for $\mathcal{E}(Y)$ involves a term $Y(x)$ for each $x \in S$, but instead of just summing $Y(x)$ we sum the products $\text{Pr}(x)Y(x)$. The factor $\text{Pr}(x)$ is what's needed to encode for the fact that if $\text{Pr}(x)$ is big (meaning x is observed often) and $\text{Pr}(x')$ is small (meaning that we seldom observe x'), then in a series on n trials there'll be many $Y(x)$ s and few $Y(x')$ s contributing to the average $\mathcal{E}(Y)$. Formally,

OUTCOMES EXPECTED VALUE FORMULA 4.8.2: *If Y is a random variable on the probability space S , then the expected value $\mathcal{E}(Y)$ is equal to the sum over outcomes x in S of the probability-weighted values $\text{Pr}(x)Y(x)$:*

$$\mathcal{E}(Y) := \sum_{x \in S} \text{Pr}(x)Y(x).$$

EXAMPLE 4.8.3: If our experiment consists of rolling a single die (so S is the set of numbers from 1 to 6 each having probability $\frac{1}{6}$), and $Y(x) = x$ (that is, we think of each roll as a real number rather than a side of the die), then

$$\mathcal{E}(Y) := \sum_{x \in S} \text{Pr}(x)Y(x) = \frac{1}{6} \cdot 1 + \frac{1}{6} \cdot 2 + \frac{1}{6} \cdot 3 + \frac{1}{6} \cdot 4 + \frac{1}{6} \cdot 5 + \frac{1}{6} \cdot 6 = \frac{21}{6} = \frac{7}{2}.$$

In other words, an “average” roll is $\frac{7}{2}$ —or better, if we roll many times, the average of the numbers we'll see will be close to $\frac{7}{2}$.

One point worth noting about this example is that we had a common factor of $\frac{1}{6}$ in the sum. That's because this was an equally likely outcomes probability space, so $\text{Pr}(x) = \frac{1}{6}$ for every outcome x and so is a common factor in the sum $\sum_{x \in S} \text{Pr}(x)Y(x)$. So we could have saved some arithmetic by writing

$$\mathcal{E}(Y) := \frac{1}{6}(1 + 2 + 3 + 4 + 5 + 6) = \frac{21}{6} = \frac{7}{2}.$$

4.8 Who would have expected ...

In this form, what we see is the *average* of the 6 values $Y(x)$, in the same way that the average of 2 numbers a and b is $\frac{1}{2}(a + b)$ and the the average of 3 numbers a , b and c is $\frac{1}{3}(a + b + c)$ and so on.

EQUALLY LIKELY OUTCOMES EXPECTED VALUE FORMULA 4.8.4: *If Y is a random variable on the equally likely outcomes probability space S , then the **expected value** $\mathcal{E}(Y)$ is equal to the average over all outcomes x in S of the values $Y(x)$:*

$$\mathcal{E}(Y) := \frac{1}{\#S} \sum_{x \in S} Y(x) .$$

Here are a couple of easy problems to get you used to finding expected values. First one with equally likely outcomes.

PROBLEM 4.8.5: Use an expected value to find how many sons are there in the average family with 3 children, assuming that each child is equally likely to be a son or a daughter.

- i) List the 8 elements in the equally likely sample space S given by recording the sexes of the 3 children. Hint: Use words in S and D .
- ii) Gives the values of the random variable Y whose value on each outcome is the number of sons.
- iii) Find $\mathcal{E}(Y)$ using the **OUTCOMES EXPECTED VALUE FORMULA 4.8.2**.

Now a problem with outcomes that are *not* equally likely to see that, in general the expected value is *not* the arithmetic average but a probability-weighted one.

PROBLEM 4.8.6: You have a black urn containing balls numbered 1, 2 and 3 and a white urn containing balls numbered 4 and 5. An experiment consists of picking an urn at random and then drawing a ball at random from the selected urn. We'll consider the sample space S for this experiment to be the numbers from 1 to 5, since the number let's us determine the urn if we want.

- i) Use a tree diagram to find the probability of each of the 5 outcomes in S .

4.8 Who would have expected ...

ii) Show that the average of the numbers from 1 to 5 is 3 5 outcomes in S .

iii) Let Y the random variable whose value on each outcome is the number on the chosen ball. Show that $\mathcal{E}(Y) = \frac{13}{4} = 3 \frac{1}{4}$ using the [OUTCOMES EXPECTED VALUE FORMULA 4.8.2](#).

In other words, in this example, the expected value is *more* than the average.

PROBLEM 4.8.7: You have a black urn containing balls numbered 1 and 2 and a white urn containing balls numbered 3, 4 and 5. An experiment consists of picking an urn at random and then drawing a ball at random from the selected urn. Let Y the random variable whose value on each outcome is the number on the chosen ball. Show that $\mathcal{E}(Y) = \frac{11}{4} = 2, \frac{3}{4}$ using the [OUTCOMES EXPECTED VALUE FORMULA 4.8.2](#). In other words, in this example, the expected value is *less* than the average.

In both cases, the explanation is clear. In [PROBLEM 4.8.6](#), the larger numbers 4 and 5 are more likely than the smaller ones. In [PROBLEM 4.8.7](#), it's the smaller numbers 1 and 2 that are more likely.

OK. These examples show that it's pretty straightforward to find expected values—at least when $\#S$ is small—but, they don't explain what's so interesting about such unconstrained functions Y and their averages $\mathcal{E}(Y)$. And the whole point of our study of counting was to be able to work with big sample spaces without listing their elements. How will we ever find expected values for such S if what's involved is not just counting but *totaling* over outcomes?

Let's look at one more complex example in detail to get a feel for the answers. Then we'll use this example to find a way to efficiently calculate expected values that we can apply when $\#S$ is big. Finally, we'll look at a variety of questions whose answers involve random variables and their expected values.

Our example involves the New York state lottery game *Sweet Millions*, described on [the lottery site](#). To play this “game”, you pay \$1 to buy



4.8 Who would have expected ...

a card with the numbers from 1 to 40 on it, and pick 6 (distinct) numbers from the card. If all 6 of your picks match a set of 6 numbers drawn at random by the lottery, you win \$1,000,000—hence the name. Matching 5 of the 6 numbers drawn wins you \$500, matching 4 wins you \$40, and matching 3 wins you \$3. The question we’ll use an expected value to answer is the obvious one. How much do you expect to win each time you play *Sweet Millions*?

One answer to this question is, “Nada, zip, bupkis”. That’s because of the 3,838,380 possible ways to fill in a card, only 128,300 will win any prize—so you have a 96.7% of winning nothing. Before going further, let’s check these numbers and find a few other counts that we’ll need to complete our analysis.

PROBLEM 4.8.8: Consider the sample space S whose outcomes are all possible *Sweet Millions* game cards. Use standard counting techniques to verify that:

- i) The number of ways of filling in a *Sweet Millions* game card—that is $\#S$ —is $C(40, 6) = 3,838,380$.
- ii) The number of ways of filling in a *Sweet Millions* card to match exactly i of the 6 numbers drawn by the lottery—let’s call this event E_i —is $C(6, i) \cdot C(34, 6 - i)$. Hint: To match exactly i of the numbers drawn, i of the 6 numbers on your card must be from the subset of 6 numbers drawn by the lottery, and the rest must be from the subset of 34 numbers *not* drawn.

Number i matched	Matching cards	Probability
exactly 6	1	$\frac{1}{3,838,380}$
exactly 5	204	$\frac{204}{3,838,380}$
exactly 4	8,415	$\frac{8,415}{3,838,380}$
exactly 3	119,680	$\frac{119,680}{3,838,380}$
at least 3	128,300	$\frac{128,300}{3,838,380}$
2 or fewer 3	3,710,080	$\frac{3,710,080}{3,838,380}$

TABLE 4.8.9: WINNING CARDS IN *Sweet Millions*

4.8 Who would have expected ...

iii) Verify the counts in TABLE 4.8.9.

OK so our numbers check, but, of course, no one would play *Sweet Millions* if they really expected to win nothing. The reason they play is for the small chance of winning a million (and for the other smaller prizes too). Because that chance is so small but the prize is so big, it's not clear what this chance is worth. One way to get an idea is to ask what would happen if we bought 3,838,380 cards and filled out one in each of the 3,838,380 possible ways. We'd then win 1 prize of \$1,000,000, plus 204 prizes of \$500, plus 8,415 of \$40, plus 119,680 of \$3, for a total winnings of

$$1 \cdot \$1,000,000 + 204 \cdot \$500 + 8,415 \cdot \$40 + 119,680 \cdot \$3 = \$1,797,640.$$

An easier way to grasp what this means is to divide this by the number of tickets—the size of our sample space—to get an average amount won per card of $\frac{1,797,640}{3,838,380} \approx 0.4683329946$. Thus, on the average (playing a great many cards indeed), we expect to gross about 47¢ in winnings each time we play—or more accurately, taking account of the dollar it costs to buy a card, to net about 53¢ in losses.

We've just found the expected value of two random variables. The first is the random variable Y that assigns to each *Sweet Millions* card (each outcome on our sample space) the gross amount it would win—that is, if $x \in E_6$ and matches all 6 numbers, $Y(x) = 1,000,000$, if $x \in E_5$ and matches exactly 5 numbers, $Y(x) = 5000$ and so on. The variable Y illustrates why we want to allow random variables to have large values. The second is the function Z that assigns to each assigns to each *Sweet Millions* card, the net amount it would win. We just subtract the cost of the card so $Z(x) = Y(x) - 1$ for every $x \in S$. The variable Z shows why we want to allow random variables to have negative values.

Let's focus on Y for a moment. The expected value $\mathcal{E}(Y) = \frac{1,797,640}{3,838,380} \approx 0.4683329946$ and represents the average amount we'd expect to win each time we play *Sweet Millions*. The usual disclaimer that goes



4.8 Who would have expected ...

with probability expectations applies: we have to play *Sweet Millions* a *great many* times before we can expect our *observed* average winnings to settle down anywhere near 47¢ a play. Even if we, say, play 100,000 times, we probably won't win a million dollar prize. If we don't, our average winnings will only be about 21¢ (because about $\frac{1,000,000}{3,838,380} \approx 27$ of the 47¢ come, on the average, from the million dollar prize). And, if we get lucky and win a million dollar prize, then our average winnings will be more than \$10 per card.⁷

Why do I say we have calculated $\mathcal{E}(Y)$ (and $\mathcal{E}(Z)$)? The sample space for *Sweet Millions* is an equally likely outcomes one, but even so, we never computed any sum, or average, over the 3,838,380 possible game cards. Indeed, while the “average of all the values” viewpoint is intuitively helpful, as a practical matter, we'd *never* want use it to calculate $\mathcal{E}(Y)$ because we'd have to add up 3,838,380 values $Y(x)$. OK, it's true that 3,710,080 of these values (corresponding to card where 2 or fewer numbers match) are 0, but that still leaves 128,300 non-zero values to total. That's be one nasty homework problem.

What we did was use a second way of thinking of expected values that is usually much better suited for calculating them. Instead of totally over outcomes, we total over *values of Y*.

RANGE 4.8.10: *The set of values of a function is usually called its range. Since random variables on S are just functions with domain S , we'll use this term. We'll denote by $R(Y)$ the range, or set of values $y = Y(x)$ of the random variable Y .*

$E_{Y,y}$ **OR** E_y **4.8.11:** *We define the event $E_{Y,y} := \{x \in S | Y(x) = y\}$ to be the set of outcomes on which Y takes on the value y . Of course, unless $y = Y(x)$ for at least one outcome, the event $E_{Y,y}$ is the empty set and we can ignore it. This leaves us with one event $E_{Y,y}$ for each*

⁷Across the state, the game nearly lost money in its first month—there were two million dollar winners on sales of less than \$3,000,000—but then went another month with no jackpot winner.



4.8 Who would have expected ...

value that Y actually takes on—that is for each $y \in R(Y)$. When the event Y is understood, we'll usually just write E_y for simplicity.

It's very common that the range $R(Y)$ —number of values y we need to consider—is quite small. In the example of *Sweet Millions*, only 5 values arise and $R(Y) = \{1000000, 5000, 40, 3, 0\}$. Correspondingly, there are only 5 events E_y corresponding to winning these amounts, or more directly, to matching 6, 5, 4, 3, or fewer than 3 numbers. That's a whole lot less than 3,838,380 outcomes.

I should disclaim that there *are* lots of random variables that take on different values on every outcome. If we need to find the expected value of one of these, we are up the creek: it's no easier to total over values than over outcomes. But, in practice, such random variable just don't come up very often and the strategy of totally over values is very effective.

How do we take a total over values? Essentially, we convert the “retail total” $\mathcal{E}(Y) = \sum_{x \in S} \Pr(x)Y(x)$ over outcomes in S into a “wholesale total” over values in $R(Y)$. The key point to notice is that every outcome x lies in the event E_y for which $y = Y(x)$ and no other. So we can rewrite $\mathcal{E}(Y) = \sum_{x \in S} \Pr(x)Y(x)$ as a double sum $\mathcal{E}(Y) = \sum_{y \in R(Y)} \sum_{x \in E_y} \Pr(x)Y(x)$.

This looks like we're losing ground not gaining it, until we remember that, by definition, the factor $Y(x)$ is equal to y for every $x \in E_y$. This lets us write $\sum_{x \in E_y} \Pr(x)Y(x) = \sum_{x \in E_y} \Pr(x)y = y(\sum_{x \in E_y} \Pr(x))$ by pulling out the common factor y . The final observation is that the last sum $\sum_{x \in E_y} \Pr(x)$ is nothing more than $\Pr(E_y)$ by the formula for the [PROBABILITY OF AN EVENT 4.2.2](#).

So we wind up with the formula $\mathcal{E}(Y) = \sum_{y \in R(Y)} y \cdot \Pr(E_y)$. In other words, we can find the expected value $\mathcal{E}(Y)$ by summing over values y of Y , the product of y with the probability of seeing an outcome where Y has the value y . In the example of *Sweet Millions*, where there are only 5 values we have a sum with only 5 terms—really 4

4.8 Who would have expected ...

that we need to worry about because a term where $y = 0$ can be ignored, as we did above.

When we are dealing with equally likely outcomes probability spaces, we can make one further simplification which again, we have already seen in our example. Since $\Pr(E_y) = \frac{\#E_y}{\#S}$ for every event E_y by the **EQUALLY LIKELY OUTCOMES FORMULA 4.3.2**, we can pull out a common denominator of $\#S$ and write $\mathcal{E}(Y) = \frac{1}{\#S} \sum_{y \in R(Y)} y \cdot \#E_y$. That is, we sum products of values and counts and then divide by $\#S$ to get $\mathcal{E}(Y)$. Let's record these formulae for future reference.

VALUES EXPECTED VALUE FORMULA 4.8.12: *If Y is a random variable on the probability space S , then the **expected value** $\mathcal{E}(Y)$ is equal to the sum over the values y in the range $R(Y)$ of Y of the probability-weighted values $\Pr(E_{Y,y}) \cdot y$:*

$$\mathcal{E}(Y) := \sum_{y \in R(Y)} \Pr(E_{Y,y}) \cdot y.$$

EQUALLY LIKELY OUTCOMES EXPECTED VALUE FORMULA 4.8.13: *If Y is a random variable on the equally likely outcomes probability space S , then the **expected value** $\mathcal{E}(Y)$ is equal to:*

$$\mathcal{E}(Y) := \frac{1}{\#S} \sum_{x \in S} \#E_{Y,y} \cdot y.$$

Calculating expected values even by this simpler method usually involves a fair bit of arithmetic. The computation is usually much easier (and less error-prone) if you organize it using the following method.

METHOD FOR FINDING EXPECTED VALUES 4.8.14: I'll first give a general method that works in all problems, then explain how it can be simplified .

Step 1: First find the sample space S and its size $\#S$.

Step 2: Identify the random variable Y whose expected value you want to find and determine its range or set of values $R(Y)$.

4.8 Who would have expected ...

Step 3: Next make a table with one row for each value in $R(Y)$, and one row for totals, and with columns for: the values y ; the probabilities $\Pr(E_y)$ of observing the value y ; and the products $\Pr(E_y)y$.

Step 4: Fill in the values column of your table.

Step 5: Fill in the probabilities column of your table.

Step 6: Fill in each row of the products column with a by multiplying the value y and probability in the same row.

Step 7: Total the products column to find the expected value $\mathcal{E}(Y)$.

This may seem a bit complicated but it mostly pretty easy in practice. Usually the only step that calls for thought rather than mere arithmetic, is Step 5. Often, however, there is a pattern to the different probabilities needed and the calculation that gives one entry can be used to find all the rest by changing a single parameter in some count or formula. To take advantage of this, it's best to enter both an unevaluated shorthand form and a numerical (fraction or decimal) value. The former let's you see the pattern, the latter is needed in the next step

When outcomes are equally likely, we can avoid decimal or fractions and save some arithmetic using a counts column instead of a probabilities column. We then just enter the count $\#E_{Y,y}$ beside each value in Step 5. However, to convert these counts back to probabilities, we need to divide the total in the products column obtained in Step 7 by $\#S$ to get $\mathcal{E}(Y)$.

EXAMPLE 4.8.15: Consider a 5 member committee chosen at random from the 59 Democrats and 41 Republicans in the U.S. Senate. How many Republicans will such a committee contain on the average?

Solution

Here we'll use the simpler, equally likely outcomes variant with a count column.

4.8 Who would have expected ...

- Step 1: As in [PROBLEM 3.8.20](#), the sample space S consists of subsets of 5 members of the 100 member Senate and has order $C(100, 5) = 75287520$.
- Step 2: The random variable Y is the number of Republicans on the committee and its range (set of possible values) is $R(Y) = \{0, 1, 2, 3, 4, 5\}$.
- Step 3: The completed table is [TABLE 4.8.16](#)
- Step 4: Duh!
- Step 5: Let's pick the row corresponding to the value $y = 2$. Again, as in [PROBLEM 3.8.20](#), there are $C(59, 3) \cdot C(41, 2)$ committees that contain 3 Democrats and 2 Republicans. To find the count with any number y of Republicans we just replace the 2 by y and the 3 by $5 - y$. This lets us fill the whole column as shown. Notice how I included both the shorthand count (to see the pattern) and the value (used in the next step). Having both seems like more work but, in the end, is the easiest way. I also totaled this column as a check—the total should always equal $\#S$ since every outcome is associated to exactly one value y
- Step 6: Simple arithmetic. I have shown the calculation of each entry, just to be sure things are clear, but we'd usually just write the products here.
- Step 7: Ditto.

Value y	Count	$\#E_y$	Product	$\#E_y \cdot y$
0	$C(59, 5) \cdot C(41, 0) =$	5006386	$0 \cdot 5006386 =$	0
1	$C(59, 4) \cdot C(41, 1) =$	18660166	$1 \cdot 18660166 =$	18660166
2	$C(59, 3) \cdot C(41, 2) =$	26657380	$2 \cdot 26657380 =$	53314760
3	$C(59, 2) \cdot C(41, 3) =$	18239260	$3 \cdot 18239260 =$	54717780
4	$C(59, 1) \cdot C(41, 4) =$	5974930	$4 \cdot 5974930 =$	23899720
5	$C(59, 0) \cdot C(41, 5) =$	749398	$5 \cdot 749398 =$	3746990
Total		75287520		154339416

TABLE 4.8.16: EXPECTED NUMBER OF REPUBLICANS

4.8 Who would have expected ...

Finally we divide the total 154339416 by 75287520 to get the Expected value $\mathcal{E}(Y) = \frac{154339416}{75287520} = \frac{41}{20} = 2.05$ for the number of Republicans.

Here are three examples for you to practice with that deal with the game Chuck-a-luck, discussed in [CHUCK-A-LUCK](#). If you get stuck you can peek at [EXAMPLE 3.8.36](#) for all of the necessary probabilities. In fact, in that example, we more or less computed the expected values that follow, without calling them by this name.

Two points to note: although we'll work with 3 different random variables Y , Z , and W , the sets of events E_y , E_z and E_w are the same—it's only the *values* of y , z and w associated to these events that are different. This means that you can reuse the same counts or probabilities column in all 3 parts of the problem. Second, it's exactly the difference between Z and W that's the key to the con in Chuck-a-luck.

PROBLEM 4.8.17: Consider the experiment of rolling 3 dice (one red, one blue and one green) and recording the number on each. Describe the sample space S and show that $\#S = 6^3 = 216$.

i) Consider the random variable Y whose value on each triple of numbers is the *number* of dice (from 0 to 3) that show a 6. Use the [Method for finding expected values 4.8.14](#) to show that $\mathcal{E}(Y) = \frac{1}{2}$.

ii) Consider the random variable W whose value on each triple of numbers is the *amount* you'd win or lose if you bet \$1 and win \$2 for each die that shows a 6. That is, Z is \$ -1 if you roll no 6s, \$ $+1$ if you roll one 6, \$ $+3$ if you roll two 6s, and \$ $+5$ if you roll three 6. Use the [Method for finding expected values 4.8.14](#) to show that $\mathcal{E}(W) = 0$.

iii) Consider the random variable Z whose value on each triple of numbers is the *amount* you'd win or lose if you bet the number 6 at a Chuck-a-luck booth. That is, Z is \$ -1 if you roll no 6s, \$ $+1$ if you roll one 6, \$ $+2$ if you roll two 6s, and \$ $+3$ if you roll three 6. Use the [Method for finding expected values 4.8.14](#) to show that $\mathcal{E}(Z) = -\frac{17}{216} \simeq -\0.07870370370 .



4.8 Who would have expected ...

When people gamble, they're naturally interested in whether they should model their expectation on something like the Z above, or on something like the W . The two situations have acquired standard names. In both cases, it is assumed that, whatever the game being played, outcomes are randomized so that expected values computed as we are doing represent actual expectations. This, of course, does not apply if we're playing roulette on a crooked wheel, or craps with shaved dice, or blackjack with a dealer who is dealing seconds ...

FAIR GAME 4.8.18: *A game of chance in which the expected value of a bettor's winnings is 0 is called a fair game.*

Of course, for this to happen, it must be possible that the winnings are negative—that is, losses!

HOUSE ADVANTAGE 4.8.19: *A game of chance in which the expected value of a bettor's winnings is negative is called a unfair game. The expected value of the players loss, expressed as a percentage of the total bet, is called the house advantage.*

We not allow for games where the expected winnings are positive, but please email me about any you may encounter.

So Chuck-a-luck (the variable W) is an unfair game and the house advantage is about 7.87%, but a similar game played with stakes described by the variable Z would be fair.

The next problem is a first look at betting “systems”. Gamblers often think that they have found a method to beat the house advantage, but they're almost always making an incorrect evaluation of their expected winnings. Our example will deal with roulette—a standard wheel is shown in [FIGURE 4.8.20](#).

There are 38 slots (our outcomes) of which 18 are red, 18 are black, and 2—the 0 and 00—are neither. There are lots of ways to bet, but we'll look only at the most common, betting on the color of the slot in which the ball comes to rest. Each dollar bet in this way returns \$2 if the selected color comes up and is lost otherwise.





FIGURE 4.8.20: A standard roulette wheel

PROBLEM 4.8.21: Assuming that the ball is equally likely to land in any slot (that is, the wheel is fair), show that the expected value of a \$1 bet on black is $-\$ \frac{2}{38} \approx -\0.05263 .

So, not surprisingly, roulette is an unfair game with a **house advantage** of 5.263%. How can we beat this house edge? The simplest system and one commonly favored by roulette players is called the **martingale system**. You leave the table as soon as you have made a winning bet and just double the amount of any losing bet. Suppose we start by betting \$1 on black. If black comes up, you quit and take your \$1. If not, we make a second bet of \$2; if this wins, we are again up \$1—\$2 from the second bet minus \$1 from the first. If the second bet loses, we place a third bet of \$4: if this wins, we are again up \$1—\$4 from the third bet minus \$3 = \$1 + \$2 from the first two. And so on. You can't lose, can you?

Not only can you lose, but you should expect to. The catch is that if red (or a 0) comes up often enough, you won't have the money needed to keep doubling your bet.⁸ The next problems let us quan-

⁸We'll ignore the fact that the casino may also limit the size of bets for reason we'll discuss later.



4.8 Who would have expected ...

tify this idea. The only probabilities we will need are given in:

PROBLEM 4.8.22: Assuming that the results of successive spins of the wheel are independent, show that the chance that none of the first $k - 1$ spins is black is $(\frac{20}{38})^{k-1}$.

Show that the chance of seeing $k - 1$ spins that are not black followed by a k^{th} black spin is $(\frac{20}{38})^{(k-1)} \frac{18}{38}$. Hint: most of the work is in the previous probability.

Confirm these by drawing a tree diagram, assuming that we quit after we win or after the third spin whether or not we win on that spin.

EXAMPLE 4.8.23: Here we'll find the expected value of playing the martingale strategy assuming we stop after the third spin, regardless of whether we have won or not. The first point to note is that we either win \$1 (if a ball lands on black on one of the three spins) or lose \$7 (if we see three non-black spins losing \$1 + \$2 + \$4). So we're looking for the expected value of a random variable Y with just 2 values 1 and -7 and we just need to find the probabilities of the two events E_1 and E_{-7} .

But, we lose \$7 only if we see no black spin in the first 3 so by **PROBLEM 4.8.22**, $\Pr(E_{-7}) = (\frac{20}{38})^3$. Moreover, E_1 and E_{-7} are complementary events so $\Pr(E_1) = 1 - \Pr(E_{-7}) = 1 - (\frac{20}{38})^3$. Applying the **VALUES EXPECTED VALUE FORMULA 4.8.12**—we don't even need to bother with a table—we find $\mathcal{E}(y) = 1 \cdot ((1 - (\frac{20}{38})^3)) + (-7) \cdot (\frac{20}{38})^3 = 1 \cdot \frac{46872}{54872} - 7 \cdot \frac{8000}{54872} = -\frac{9128}{54872} \simeq -\0.1664 .

So what? You haven't really played the system if you stop after 3 bets. The whole point is to keep playing. Nope. This just makes your expected losses bigger.

PROBLEM 4.8.24: Find the expected value of playing the martingale strategy assuming we stop when we win but also, regardless of whether we have won or not:

- You should have seen the expected loss grow to about \$0.2277 and \$0.2924 in these two examples. The next example uses different expected values to check this.

Value z	Probability $\Pr(E = z)$	Product $\Pr(E = z) \cdot y$
0	0.1	0.0
1	0.2	0.2
2	0.3	0.6
3	0.2	0.6
4	0.1	0.4
5	0.1	0.5

Value z	Probability	$\Pr(E_z)$	Product $\Pr(E_z) \cdot y$
1	$\frac{18}{38} =$	$\frac{684}{1444}$	$\frac{684}{1444}$
3	$\frac{20}{38} \frac{18}{38} =$	$\frac{360}{1444}$	$\frac{1080}{1444}$
7	$\frac{20}{38} \frac{20}{38} =$	$\frac{400}{1444}$	$\frac{2800}{1444}$
Total		$\frac{400}{1444}$	$\frac{4564}{1444}$

TABLE 4.8.26: TOTAL BET IN A THREE SPIN MARTINGALE

So $\mathcal{E}(Z) = \frac{4564}{1444} \simeq \3.1601 . But, if we bet an average of $\$ \frac{4564}{1444}$ each time we following the martingale strategy for three spins, and we lose an average of $\$ \frac{2}{38}$ of each dollar we bet, we should expect to lose $\$ \frac{4564}{1444} \frac{2}{38} = \$ \frac{9128}{54872}$ each time we bet this strategy. In decimals, $\$3.1601 \cdot \$0.0526 \simeq \$0.1664$. This exactly confirms [EXAMPLE 4.8.23](#).

4.8 Who would have expected ...

PROBLEM 4.8.27: Find the expected value number of dollars you will bet following the martingale strategy assuming we stop when we win but also, regardless of whether we have won or not:

i) after the fourth spin. Hint: Here you bet \$7 if you lose on the first 2 spins then win the third, and \$15 if you lose on the first 3, regardless of the outcome of the fourth; so you can again apply [PROBLEM 4.8.22](#).

ii) after the fifth spin.

Then use these two values to check your answers to [PROBLEM 4.8.24](#) as above.

Still, what these examples show is that if you bail out on a martingale strategy—that is, stop doubling up even though you haven’t won—you can’t expect to be a winner and that the longer you wait to bail, the more you can expect to lose. But isn’t that an unfair test? The whole idea behind a martingale strategy is that you *never* bail and always double up. There’s just one problem with this idea. If you lose too many spins, there will come a point when your bankroll will no longer be big enough to double up and you’ll be *forced* to bail whether you want to or not.

For example, suppose I come to the table with \$50. If I lose the first 5 spins, I am down to \$19 (bets of 1, 2, 4, 8 and 16) and I am forced to bail because I do not have \$32 to double up with. Most of the time (from [PROBLEM 4.8.24](#) we can see it’s about 95.96% of the time) you will win \$1. But the other 4.04% of the time you’ll lose \$31 and, on the average, that means you’ll lose \$0.2924 (again from [PROBLEM 4.8.24](#)).

PROBLEM 4.8.28: Show that if you come to the table with a bankroll of \$5,000, you’ll have to bail if you lose 12 times in a row—leaving a \$4,095 loser.

CHALLENGE 4.8.29: Show that if you come to the table with a bankroll of \$5,000, the chance you will have to bail is only about

4.8 Who would have expected ...

0.00045181 or less than 1 in 2,000 but that, nonetheless, your expected losses will be \$0.8506.

One further point to note is that the martingale strategy increases the average amount you wager *on each spin*. As we will see at the end of this section, you can't expect to beat the house, but you can hold out against it longer by keeping your bets small: so playing a the martingale strategy will usually cause you to tap out faster than repeatedly betting a single fixed sum. Worse, the more spins you stick with the strategy the bigger your average bet gets.

PROBLEM 4.8.30: Let's denote by W the average number of dollars on each spin by following the martingale strategy assuming you stop when you win but also after k spins, regardless of whether you have won or not.

i) Show that if we take $k = 3$ (stop after the third spin regardless), then Show that the expected value $\mathcal{E}(W) = \$\frac{3004}{1444} \simeq \2.08 .

Hint: TABLE 4.8.26 can be reused with only a little modification. To get the value bet per spin of W , you just divide the values column there by the corresponding number of spins, replacing $1/3/7$ by $\frac{1}{1}/\frac{3}{2}/\frac{7}{3}$. The probabilities remain the same (they just depend on how many spins you bet). So only the products column needs recalculating. You can use a common denominator of 38^2 .

ii) Show that if we take $k = 4$ and $k = 5$, we get expected average bets per spin of $\mathcal{E}(W) = \$\frac{146152}{54872} \simeq \2.66 and $\mathcal{E}(W) = \$\frac{6833776}{2085136} \simeq \3.28 .

Hint: Use common denominators of 38^3 and 38^4 .

Finally, one challenge. Have you ever had to try to find the right key to open a lock in the dark and noticed it can seem to take forever. How many tries should you expect to have to make? First a warm-up.

PROBLEM 4.8.31: Assume that you have 10 keys on your key ring. You repeatedly choose a key at random from the ring and try it in

4.8 Who would have expected ...

the lock. If it doesn't work, you put take the chosen key off the ring and try again. How many keys do you expect to have to try before you find the right one?

Hint: The probability you'll succeed on the first try is $\frac{1}{10}$ and that you'll fail is $\frac{9}{10}$. So the probability that you'll succeed on the second trial is $\frac{9}{10} \cdot \frac{1}{9} = \frac{1}{10}$ and that you'll fail the first *two* times is $\frac{8}{10}$.

CHALLENGE 4.8.32: Assume again that you have 10 keys on your key ring. You repeatedly choose a key at random from the ring and try it in the lock. This time, however, if it doesn't work, you simply pick another key at random from the key ring (possibly one you have already tried—remember it's dark). Show that you expect to have make 10 tries before hitting the right key.

Hint: Show that probability that you'll succeed on the n^{th} trial—that is, that the value of the number N of keys you need to try is n —is $\frac{9}{10}^{(n-1)} \cdot \frac{1}{10}$. Use this to write $\mathcal{E}(N)$ as an infinite sum S . Then consider $S - \frac{9}{10}S$ and use [GEOMETRIC SERIES FORMULA 1.3.6](#).

Relations between expected values

The examples of the previous subsection—especially, the check in [EXAMPLE 4.8.25](#) and [PROBLEM 4.8.27](#)—make it clear that related random variables have related expected values. It turns out that, if we take advantage of these relations, we can enormously simplify many expected value calculations.

Our goal here is to set down the simplest, but also the most useful, rules of this type. These rules all easy-to-remember because they are all [BANG ZOOM RULES 3.3.16](#) that say that operations of random variable are reflected in their expected values. The *bang* in these rules is “expected value” and we have a variety of *zooms*. To keep the statements of these rules short I'll abbreviate “expected value” as EV. So we'll be looking for rules that have the shape, “The EV of the *zoom* is the *zoom* of the EVs”.



4.8 Who would have expected ...

First the simplest of all. Let's write C_a ("constant a ") for the random variable that has only one value a . Thus $R(C_a) = \{a\}$, then the **VALUES EXPECTED VALUE FORMULA 4.8.12** $\sum_{y \in R(C_a)} \Pr(E_{C_a, y})y$ becomes just $\Pr(E_a) \cdot a$. But $E_a = S$ so $\Pr(E_a) = \Pr(S) = 1$ and $\mathcal{E}(C_a) = a$.

CONSTANTS ARE THEIR OWN EVS 4.8.33: *The expected value of the constant random variable C_a with value a for every outcome x is $\mathcal{E}(C_a) = a$. Informally, constants are their own expected values.*

Likewise, suppose that $Z = aY$ for some constant a —in other words, Z is a constant multiple of Y . Then the range of Z is just the set of values $z = ay$ where y is in the range of Y . Moreover, $Z(x) = z$ exactly when $Y(x) = y$ so $E_{Z, z} = E_{Y, y}$ and hence $\Pr(E_{Z, z}) = \Pr(E_{Y, y})$. Thus the **VALUES EXPECTED VALUE FORMULA 4.8.12** gives, $\mathcal{E}(Z) := \sum_{z \in R(Z)} \Pr(E_{Z, z}) \cdot z = \sum_{y \in R(Y)} \Pr(E_{Y, y}) \cdot ay = a \sum_{y \in R(Y)} \Pr(E_{Y, y}) \cdot y = a\mathcal{E}(Y)$ by pulling out the common factor of a .

THE EV OF THE MULTIPLE IS THE MULTIPLE OF THE EV 4.8.34: *If the random variable Z is a constant a times Y , then $\mathcal{E}(Z) = a\mathcal{E}(Y)$. Informally, the expected value of a constant multiple is the multiple of the expected value.*

Next, the analogous rule holds for sums:

THE EV OF THE SUM IS THE SUM OF THE EVS 4.8.35: *If we can write the random variable Y as the sum of two (or more) other random variables Z and W — $Y = Z + W$, then $\mathcal{E}(Y) = \mathcal{E}(Z) + \mathcal{E}(Y)$. Informally, the expected value of a sum is the sum of the expected values.*

Saying $Y = Z + W$ just means that for each outcome x the value $Y(x)$ is the sum of the values $Z(x)$ and $W(x)$: $Y(x) = Z(x) + W(x)$. This notion is a bit trickier than that of a multiple, so let's look at an example.

EXAMPLE 4.8.36: Consider the sample space for a sequence of 2 binomial trials. To be definite, I'll think of tossing a coin twice with heads as S . Consider the random variables L_1 , L_2 and K in the first

4.8 Who would have expected ...

three rows of the table below. We'll use the L and M rows later in the section.

Random variable	Outcome				Expected value
	HH	HT	TH	TT	
L_1	1	1	0	0	$\frac{1}{2}$
L_2	1	0	1	0	$\frac{1}{2}$
$K = L_1 + L_2$	2	1	1	0	1
$M = L_1 \cdot L_2$	1	0	0	0	$\frac{1}{4}$
$N = K \cdot K$	4	1	1	0	$\frac{3}{2}$

TABLE 4.8.37: SUMS OF RANDOM VARIABLES FOR TWO TOSSES

Informally, L_1 is the “number of heads on the first toss”—of course, this is either 1 if the first toss is a head or 0 if it’s a tail; likewise, L_2 is the “number of heads on the second toss”. As the table shows, $K = L_1 + L_2$: each $K(x)$ is obtained by summing the values $L_1(x)$ and $L_2(x)$ above it. But we can also think of K , in its own right, as the *total* number of heads in the 2 tosses. The expected values are easy to find using the [EQUALLY LIKELY OUTCOMES EXPECTED VALUE FORMULA 4.8.4](#) (all 4 outcomes have probability $\frac{1}{4}$) and the main point is that the value in K row is again the sum of those in the L_1 and L_2 rows.

For future reference, it’s useful to relate the expected values of L_1 and L_2 to the binomial parameter p which here is $\frac{1}{2}$. Note that $L_1(x) = 1$ if the outcome of the first trial is s and $L_1(x) = 0$ if the outcome of the first trial is f. Thus $\Pr(E_{L_1,1}) = p$ and $\Pr(E_{L_1,0}) = q$. So we can compute $\mathcal{E}(L_1) = p \cdot 1 + q \cdot 0 = p$ using the [VALUES EXPECTED VALUE FORMULA 4.8.12](#). Of course, we get back the value $\frac{1}{2}$. The same reasoning can be applied equally well to L_2 .

Simple as it is, this example is the prototype for the most important applications of the rule that, [THE EV OF THE SUM IS THE SUM OF THE](#)

4.8 Who would have expected ...

EVS 4.8.35. We can replace our experiment by any binomial one (with any probability p of success and any number n of trials). Instead of 2 random variables L_1 and L_2 , we define n —denoted L_1, L_2, \dots, L_n —just as above by setting $L_i(x)$ equal to the number of ss on the i^{th} trial (again, this always is either 1 or 0). Then we define K to be their sum: $K(x) = \sum_{i=1}^n L_i = L_1(x) + L_2(x) + \dots + L_n(x)$. Just as above, we can also think of K , in its own right, as the *total* number of heads in all n trials.

Again, just as above, each of the expected values $\mathcal{E}(L_i) = p$. Seeing s (or f) on the i^{th} trial is the same as having $L_i(x) = 1$ (or $L_i(x) = 0$) and has probability p (or q). Since the only values of L_i are 1 and 0, we again find $\mathcal{E}(L_i) = p \cdot 1 + q \cdot 0 = p$ using the **VALUES EXPECTED VALUE FORMULA 4.8.12**. But, applying **THE EV OF THE SUM IS THE SUM OF THE EVS 4.8.35**, this means that

$$\mathcal{E}(K) = \mathcal{E}(L_1) + \mathcal{E}(L_1) + \dots + \mathcal{E}(L_n) = \underbrace{p + p + \dots + p}_{n \text{ terms}} = n \cdot p.$$

This example is so important we record it.

SUCCESSSES IN THE i^{th} BINOMIAL TRIAL 4.8.38: *We denote by L_i be the random variable whose value is the number of success on the i^{th} trial alone. The variables L_i all have expected value $\mathcal{E}(L_i) = p$, where p is the probability of success in a single trial.*

INDICATOR EXPECTED VALUE FORMULA 4.8.39: Random variables I that only take the values 0 and 1—like the L_i above—come up frequently enough to have a name. Such a random variable is called an **indicator** variable. Every event E determines a unique indicator variable I_E : the variable that is 1 on outcomes in E and 0 on all others. We can reverse this by associated to an indicator variable I the event $E_I := E_{I,1}$ of outcomes where I has value 1. Moreover, $\mathcal{E}(I) = \Pr(E_{I,1}) \cdot 1 + \Pr(E_{I,0}) \cdot 0 = \Pr(E_I)$.

TOTAL SUCCESSSES IN BINOMIAL TRIALS 4.8.40: *Let K be the random variable on the sample space of a sequence of n binomial trials*

4.8 Who would have expected ...

whose value on each outcome is the total number of successes. Then the variable K is the sum of the n variables L_1, L_2 to L_n : $K = \sum_{i=1}^n L_i$ and, by THE EV OF THE SUM IS THE SUM OF THE EVS 4.8.35, the variable K has expected value $\mathcal{E}(K) = n \cdot p$.

Strictly speaking we should write K_n to make the number of trials completely clear, but, as above, the number n is usually clear from the context—so we'll only add the subscript when it's needed to avoid confusion.

The most important part of TOTAL SUCCESSES IN BINOMIAL TRIALS 4.8.40 is that the total successes K is the sum of the indicator variables L_i . We will use this in the last subsection to understand spreads or variances for binomial distributions. The formula $\mathcal{E}(K) = n \cdot p$ is very useful, but also very intuitive. When we say that if we toss a fair coin 100 times we expect to see about 50 heads—which we can do without much thought—we're really using this formula: here $n = 100$ and $p = \frac{1}{2}$ and we mentally multiply $100 \cdot \frac{1}{2}$ to get 50.

OK. So we're going to have lots of applications for THE EV OF THE SUM IS THE SUM OF THE EVS 4.8.35 and we can how it works in simple examples. But why is it always true? I'll stick to explaining it for a sum of 2 random variables; the case of more than 2 summands requires a bit of extra fiddling but the central idea is the same. This is one of the few cases where the general OUTCOMES EXPECTED VALUE FORMULA 4.8.2 is needed. We use it to write compute $\mathcal{E}(Y) = \sum_{x \in S} \Pr(x) Y(x) = \sum_{x \in S} \Pr(x) (Z(x) + W(x))$ —the last is just the definition of $Y(x)$. Next we distribute the $\Pr(x)$ and divide into two sums: $\mathcal{E}(Y) = \sum_{x \in S} (\Pr(x) Z(x) + \Pr(x) W(x)) = \sum_{x \in S} \Pr(x) ZY(x) + \sum_{x \in S} \Pr(x) W(x)$. Finally, we use OUTCOMES EXPECTED VALUE FORMULA 4.8.2 again to recognize the first sum as $\mathcal{E}(Z)$ and the second as $\mathcal{E}(W)$ getting, $\mathcal{E}(Y) = \mathcal{E}(Z) + \mathcal{E}(W)$ as desired.

I hope than many of you have guessed what's coming next. If THE EV OF THE SUM IS THE SUM OF THE EVS 4.8.35, shouldn't "The EV



4.8 Who would have expected ...

of the product equal the product of the EVs”? Unfortunately, this is one of those **BANG ZOOM RULES 3.3.16** that’s just wrong. We can see this by going back to **TABLE 4.8.37** and considering the random variable N that is the product of the variable K with itself. That is $N(x) = K(x) \cdot K(x)$ for each outcome x as you can easily check. But $\mathcal{E}(N) = \frac{3}{2}$ while $\mathcal{E}(K) \cdot \mathcal{E}(K) = 1 \cdot 1 = 1$. So here, and in general, “The EV of the product is *not* the product of the EVs”.

On the other hand, it *is* correct in many cases. An example is the random variable M from the same table which is equal to the product of the L_1 and L_2 rows. Here $\mathcal{E}(M) = \frac{1}{4}$ which agrees with $\mathcal{E}(L_1) \cdot \mathcal{E}(L_2) = \frac{1}{2} \cdot \frac{1}{4}$. So the real question we need to understand is “*When* is the EV of the product equal to the product of the EVs?”. The answer is when the random variables being multiplied are **independent**.

THE EV OF THE PRODUCT OF INDEPENDENT VARIABLES IS THE PRODUCT OF THE EVS 4.8.41: *If we can write the random variable Y as the product of two (or more) **independent** random variables Z and W — $Y = Z \cdot W$, then $\mathcal{E}(Y) = \mathcal{E}(Z) \cdot \mathcal{E}(W)$. Informally, the expected value of a product of independent random variables is the product of the expected values of the factors.*

There’s just one problem here. We know what independence means for events. What does it mean to say two random variables are independent? Let’s postpone worrying about this for a while and try to get a feel for what’s going on.

One analogy that’s helpful is to think of multiplication of random variables as the analogue of intersection of events. Now recall **INTERSECTION PROBABILITY OF INDEPENDENT EVENTS 4.7.12:** If E and F are independent, then $\Pr(E \cap F) = \Pr(E) \cdot \Pr(F)$. So if we have an intersection product formula for independent events, it’s reasonable to at least hope for one for independent random variables. We just need to come up with the right notion of independence.

We can start by looking more closely at the random variables L_1 (and L_2). In particular, let’s ask “What are the events $E_{L_1,1}$ and $E_{L_1,0}$?” A

4.8 Who would have expected ...

glance at [TABLE 4.8.37](#) makes the answer clear: $E_{L_1,1}$ = “heads on the first toss” and $E_{L_1,0}$ = “tails on the first toss”. Likewise, $E_{L_2,1}$ = “heads on the second toss” and $E_{L_2,0}$ = “tails on the second toss”. In particular, any of the four pairs of events we get by specifying a value (0 or 1) for L_1 and another value (0 or 1) for L_2 are *independent*. This is why $\mathcal{E}(L_1 \cdot L_2) = \mathcal{E}(L_1) \cdot \mathcal{E}(L_2)$ as we can see by evaluating the right hand side using the [VALUES EXPECTED VALUE FORMULA 4.8.12](#):

$$\begin{aligned}\mathcal{E}(L_1) \cdot \mathcal{E}(L_2) &= (\Pr(E_{L_1,1}) \cdot 1 + \Pr(E_{L_1,0}) \cdot 0) \cdot (\Pr(E_{L_2,1}) \cdot 1 + \Pr(E_{L_2,0}) \cdot 0) \\ &= \Pr(E_{L_1,1})\Pr(E_{L_2,1}) \cdot 1 \cdot 1 + \Pr(E_{L_1,1})\Pr(E_{L_2,0}) \cdot 1 \cdot 0 \\ &\quad + \Pr(E_{L_1,0})\Pr(E_{L_2,1}) \cdot 0 \cdot 1 + \Pr(E_{L_1,0})\Pr(E_{L_2,0}) \cdot 0 \cdot 0 \\ &= \Pr(E_{L_1,1} \cap E_{L_2,1}) \cdot 1 \cdot 1 + \Pr(E_{L_1,1} \cap E_{L_2,0}) \cdot 1 \cdot 0 \\ &\quad + \Pr(E_{L_1,0} \cap E_{L_2,1}) \cdot 0 \cdot 1 + \Pr(E_{L_1,0} \cap E_{L_2,0}) \cdot 0 \cdot 0\end{aligned}$$

by using the [INTERSECTION PROBABILITY OF INDEPENDENT EVENTS 4.7.12](#) on each term.

Now we collect the probabilities that are multiplied by 0 and 1 together, obtaining $\Pr(E_{L_1,1} \cap E_{L_2,1}) \cdot 1 + (\Pr(E_{L_1,1} \cap E_{L_2,0}) + \Pr(E_{L_1,0} \cap E_{L_2,1}) + \Pr(E_{L_1,0} \cap E_{L_2,0})) \cdot 0$ and use the [OR ELSE FORMULA FOR PROBABILITIES 4.2.6](#), to write the sum of the probabilities of the three *disjoint* events in the second term as

$$\Pr((E_{L_1,1} \cap E_{L_2,0}) \dot{\cup} (E_{L_1,0} \cap E_{L_2,1}) \dot{\cup} (E_{L_1,0} \cap E_{L_2,0}))$$

obtaining, finally,

$$\begin{aligned}\mathcal{E}(L_1) \cdot \mathcal{E}(L_2) &= \Pr(E_{L_1,1} \cap E_{L_2,1}) \cdot 1 \\ &\quad + \Pr((E_{L_1,1} \cap E_{L_2,0}) \dot{\cup} (E_{L_1,0} \cap E_{L_2,1}) \dot{\cup} (E_{L_1,0} \cap E_{L_2,0})) \cdot 0.\end{aligned}$$

Here’s why I went to all this trouble to rewrite a probability that is getting multiplied by 0 anyway. The event $E_{L_1,1} \cap E_{L_2,1}$ —“heads on both tosses”—is also a value event for $M = L_1 \cdot L_2$. It’s just the event that the product variable $L_1 \cdot L_2$ has value 1: $E_{L_1,1} \cap E_{L_2,1} = E_{M,1}$. Likewise the event $((E_{L_1,1} \cap E_{L_2,0}) \dot{\cup} (E_{L_1,0} \cap E_{L_2,1}) \dot{\cup} (E_{L_1,0} \cap E_{L_2,0}))$ (when either factor is 0) is equal to the event $E_{L_1 \cdot L_2,0}$ (when the product is



4.8 Who would have expected ...

0). Plugging these in and then using the **VALUES EXPECTED VALUE FORMULA 4.8.12** for $L_1 \cdot L_2$ we see that

$$\mathcal{E}(L_1) \cdot \mathcal{E}(L_2) = \Pr(E_{L_1 \cdot L_2, 1}) \cdot 1 + \Pr(E_{L_1 \cdot L_2, 0}) \cdot 0 = \mathcal{E}(L_1 \cdot L_2).$$

Of course, this was a lot more work than just computing all the expected values and comparing. But we'll have applications in which using this product relation is the only way to get the answer. Moreover, this toy calculation contains all the ideas needed to show **The EV of the product of Independent Variables is the product of the EVs 4.8.41** in general. The key point is a value event $E_{Z \cdot W, y}$ for the product $Y = Z \cdot W$ is the disjoint union of the intersection events $E_{Z, z} \cap E_{W, w}$ for all solutions z and w of the equation $y = z \cdot w$. If we know the pairs of events $E_{Z, z}$ and $E_{W, w}$ are *all* independent, then this is enough to give $\mathcal{E}(Z \cdot W) = \mathcal{E}(Z) \cdot \mathcal{E}(W)$ by imitating the argument given above for L_1 and L_2 . The details are a bit messy, so I won't give them, but the ideas needed are *exactly* the same. So we've found out what the right definition of independent random variables is.

INDEPENDENT RANDOM VARIABLES 4.8.42: *We say the random variables Z and W are independent if for every value z of Z and every value w of W the value events $E_{Z, z}$ and $E_{W, w}$ are independent events—that is, $\Pr(E_{Z, z} \cap E_{W, w}) = \Pr(E_{Z, z}) \cdot \Pr(E_{W, w})$.*

As we have seen above, any two of the variables L_i and L_j that count **SUCCESSSES IN THE i^{th} BINOMIAL TRIAL 4.8.38** are independent in this sense. Hence

SUCCESSSES PRODUCT FORMULA 4.8.43: *When i and j are not equal,*

$$\mathcal{E}(L_i \cdot L_j) = \mathcal{E}(L_i) \cdot \mathcal{E}(L_j) = p \cdot p = p^2.$$

Here are a couple of problems that show how relations amongst can simplify expected value calculations.

PROBLEM 4.8.44: First, let's look again at **EXAMPLE 4.8.15** involving the number of Republicans Y on a 5 member committee chosen at

4.8 Who would have expected ...

random from the 59 Democrats and 41 Republicans in the U.S. Senate. We first define some indicator random variables, then use them to recompute $\mathcal{E}(Y) = 2.05$.

- For each Senator s , show that the probability that s is on the committee is $\frac{5}{100} = 0.05$. Hint: Show that 3,764,376 of the possible committees contain s .
- Let Z_s be the random variable that is 1 if s is on the committee and 0 otherwise. Show that $\mathcal{E}(Z_s) = \frac{5}{100} = 0.05$.
- Show that Y is the sum of the 41 variables Z_s for which s is a Republican.
- Use the previous two parts and [THE EV OF THE SUM IS THE SUM OF THE EVS 4.8.35](#) to show that $\mathcal{E}(Y) = 41 \cdot 0.05 = 2.05$.

Consider a bus that makes 10 stops after it leaves the terminal. If there are 5 passengers on the bus, each is equally likely to get off at each stop and they choose their stops independently of each other, how many stops can you expect the bus to make?

What's the sample space S ? Outcomes—all equally likely—are just choices of 8 stops from 10 and since **R?** is “Yes”, $\#S = 10^5 = 100,000$. The number of stops is a random variable—call it Y —which can take the values 1 to 5 and we just want to find $\mathcal{E}(Y)$. So one way to proceed is just to count the events E_y and apply the [EQUALLY LIKELY OUTCOMES EXPECTED VALUE FORMULA 4.8.13](#). Let's try.

PROBLEM 4.8.45: First we'll get some numbers that are not quite the right counts.

- Show that the number of outcomes in which all the passengers get off at the first 4 stops is 4^5 .
- Show that the number of outcomes in which all the passengers get off at the first y stops is y^5 .
- Show that the number of different sets of exactly y stops is $C(10, y)$.
- How many outcomes are there for which there are y stops where all the passengers get off?



4.8 Who would have expected ...

I'm guessing that your answer to iv) was $y^5 \cdot C(10,y)$. We need to choose which y stops the passengers get off at, then count the ways they can get off at these stops. Sadly, there's a lot of overcounting in that number and it's wrong. We can see this by evaluating it when $y = 5$ —all the passengers get off at different stops—when it gives $5^5 \cdot C(10, 5) = 787,500$ which is much bigger than the 100,000 outcomes in the sample space. The right answer in this case is $P(5, 5) \cdot C(10, 5) = 30,240$.

The problem is that a count like 4^5 in i) includes the cases where some passenger got off at each of the first 4 stops, but also cases where passengers only got off at the first 3 stops, or where they all got off at the second stop, and so on. Moreover, these cases are often multiply overcounted so unwinding this overcounting is quite tricky, and I'm not even going to try it here. I'll just give you the answer below.

PROBLEM 4.8.46: Complete the table below to show $\mathcal{E}(Y) = 4.0951$: this confirms the Murphy's Law feeling the bus always stops pretty much as many times as possible.

Stops y	Count	Product
1	10	
2	1350	
3	18000	
4	50400	
5	30240	
Total	100000	

TABLE 4.8.47: EXPECTED NUMBER OF TIMES THE BUS STOPS

Now we're going to find $\mathcal{E}(Y)$ without doing any tricky counting, by relating Y to simpler random variables. The idea is to let Z_i be the **indicator** variable whose value is 1 on the event E_i that “some passenger(s) get off at the i^{th} stop”. (and 0 on any outcome where

4.8 Who would have expected ...

no one gets off at this stop). The variable Y is just the total number of stops where someone gets off—in other words Y is the sum of the 10 variables Z_1 to Z_{10} . Moreover, the probabilities that someone gets off at stop i are clearly the same for all i so all 10 of these expected values are the same. So, by the [INDICATOR EXPECTED VALUE FORMULA 4.8.39](#), all we have to do is find, say $\Pr(E_1)$, and multiply by 10 to get $\mathcal{E}(Y)$.

Whether any given passenger does or doesn't get off at the first stop is a binomial trial. since stops are chosen at random. Since the 5 passengers make independent choices, we have a binomial experiment with $n = 5$ and $p = \frac{1}{10}$ (since there are 5 passengers and 10 stops). By the [BINOMIAL DISTRIBUTION FORMULA 4.7.23](#), the chance that *no* passenger will get off at the first stop is $C(5, 0) \left(\frac{1}{10}\right)^0 \left(\frac{9}{10}\right)^5 = \frac{59049}{100000}$. We want the probability of the complementary event E_1 that at least one passenger *does* get off at the first stop, so $\Pr(E_1) = 1 - \frac{59049}{100000} = \frac{40951}{100000}$. Thus $\mathcal{E}(Y) = 10 \cdot \frac{40951}{100000} = 4.0951$ as above.

PROBLEM 4.8.48: At how many stops, should we expect exactly 2 passengers to get off? Hint: You can use the [BINOMIAL DISTRIBUTION FORMULA 4.7.23](#) to find the chance of the event E that “exactly 2 passengers get off at the first stop”. Use the indicator variable Z_E (and 9 analogues for the other stops).

PROBLEM 4.8.49: If the bus route has 20 stops and there are 30 passengers how many stops should the driver expect to make? At how many stops should she expect more than 1 passenger to get off?

PROBLEM 4.8.50: This is a follow up to [PROBLEM 3.8.43](#), where we saw that in a group of 23 people with randomly chosen birthdays, it's more likely than not that there'll be two people with the same birthday.

Suppose that a club has 50 members whose birthdays fall on random days of the year and independently of each other (and none on

4.8 Who would have expected ...

February 29th). Find the expected number of days of the year that are the birthday of:

- i) 0 members;
- ii) exactly 1 member;
- iii) at least 2 members;

Hint: In each case, use 365 indicator variables, one for each day, determined by having the desired birthday(s) on that day.

Simpson's Paradox

What we've been studying are relations between the expected values of *different* random variables on the *single* sample space S . Before we turn to the most critical applications of the relations between expected values developed in the preceding subsection, a few cautions are in order.

One way to think of the relations we use is as ways of dividing a single random variable Y into smaller, hopefully simpler, pieces. A typical example is [TOTAL SUCCESSES IN BINOMIAL TRIALS 4.8.40](#) where we divide the “total successes” binomial random variable K_n into n “bite-size” Bernoulli pieces L_i .

Let's ask, instead “Can we understand the expected value of a random variable Y on a sample space S more easily by dividing S into smaller subsets T and studying the expected values of Y on these pieces?” As I have stated it, this question makes no sense.

A random variable Y on S is a way of assigning a number $Y(x)$ to every outcome x in S , so we have such numbers for the outcomes in each of the smaller pieces T . But a probability distribution on S —meaning numbers $\Pr(x)$ for each outcome x —won't give us a probability distribution on the subset T . We have probabilities $\Pr(x)$ for x in T because T is a subset of S , so each x is also in S , but these only sum to 1 when we total over *all* of S , not over the subset T .

4.8 Who would have expected ...

If we stick to the equally likely outcomes case, this isn't so hard to get around. All the outcomes x in S have the same probability $\Pr(x) = \frac{1}{\#S}$. Can we find a new probability \Pr_T on T so that all the outcomes in T have the same probability too and so that these new probabilities total to 1 over outcomes in T ? Sure, we simply set $\Pr_T(x) = \frac{1}{\#T}$ for each x in T .

Now we can at least ask our question sensibly. Let's write $\mathcal{E}_S(Y)$ for the expected value of Y on S : i.e. $\sum_{x \in S} Y(x)\Pr(x)$. Likewise for a subset T of S , write $\mathcal{E}_T(Y)$ for the expected value $\sum_{x \in T} Y(x)\Pr_T(x)$ on T . What we want to know is whether we can expect there to be useful relations between $\mathcal{E}_S(Y)$ and the expected values $\mathcal{E}_T(Y)$ for various subsets?

The short answer is "Usually not" and we've already seen what can go wrong in "WE WUZ ROBBED". In the rest of this section, I want to look at better known example. Let's take as our sample space S the set of applicants to the graduate and professional schools at a large university with the [EQUALLY LIKELY OUTCOMES PROBABILITY MEASURE 4.3.1](#). Although I have made up the data to keep things simple, this example is based on a celebrated 1973 class action suit involving the graduate programs at Berkeley summarized in [this paper](#).

What random variable Y on S do we want to consider? To make everything as simple as possible, I'll take Y to be the indicator variable for the subset A of admitted students: that is $Y(x)$ is 1 if applicant x was admitted, and is 0 otherwise. The expected values $\mathcal{E}_S(Y)$ is just the overall fraction of all applicants who are admitted as the next problem shows.

PROBLEM 4.8.51: se the [EQUALLY LIKELY OUTCOMES EXPECTED VALUE FORMULA 4.8.4](#) to show that $\mathcal{E}_S(Y) = \frac{\#A}{\#S}$.

Similarly, for a subset T of S the expected value $\mathcal{E}_T(Y)$ is just the fraction of the applicants in T who get admitted: $\mathcal{E}_T(Y) = \frac{\#A \cap T}{\#T}$.

There's an even simpler way to think of these fractions. The first $\frac{\#A}{\#S}$



4.8 Who would have expected ...

is just the unconditional probability $\Pr(A)$ of the event A and the second $\frac{\#A \cap T}{\#T}$ is the conditional probability $\Pr(A|T)$.

The last element in the setup is to decide what subsets T of the full pool of applications S we want to consider. In the Berkeley case, the pool was divided in two ways.

First, we divide by the gender of the applicant. That is, we write $S = F \dot{\cup} M$ as the disjoint union of the subsets of female and male applicants. So $\mathcal{E}_F(Y) = \Pr(A|F)$ is just the chance that a female applicant was admitted and $\mathcal{E}_M(Y) = \Pr(A|M)$ the chance that a male applicant was admitted. The fact that these numbers were unequal—the proportion of female applicants admitted was substantially lower than the proportion of male applicants—was the basis of the suit. It alleged that this demonstrated discrimination on admission on the basis of sex. [TABLE 4.8.52](#) shows simplified data to illustrate this claim.

	Female	Male	Total
Admitted	800	1500	2300
Not admitted	3200	4500	7700
Total	4000	6000	10000

TABLE 4.8.52: OVERALL ADMISSION RATES BY GENDER

The case seems pretty clear. Overall, 25% of male applicants were admitted but only 20% of female applicants. The surprise came in the University’s defense. It acknowledged that discrimination on the basis of sex *was* taking place, but claimed that all of the University’s post-graduate programs were actually *favoring* women applicants.

The University’s claim was supported by data like that in [TABLE 4.8.53](#). I’ll let you check that the overall totals in each category match those in [TABLE 4.8.52](#). In this table we have divided the applicants by the school—Business (B), Law (L), Engineering (E) or Health Sciences (H)—to which each applied.

4.8 Who would have expected ...

Business	Female	Male	Total
Admitted	200	850	1050
Not admitted	200	1150	1350
Total	400	2000	2400

Engineering	Female	Male	Total
Admitted	200	400	600
Not admitted	400	1100	1500
Total	600	1500	2100

Law	Female	Male	Total
Admitted	200	150	350
Not admitted	800	850	1650
Total	1000	1000	2000

Health Sciences	Female	Male	Total
Admitted	200	100	300
Not admitted	1800	1400	3200
Total	2000	1500	3500

TABLE 4.8.53: PROGRAM BASED ADMISSION RATES BY GENDER

In other words, we are writing $S = B \cup E \cup L \cup H$ and then dividing each of programs by sex

$$\begin{aligned}
 S &= F \cup M \\
 &= ((B \cap F) \cup (E \cap F) \cup (L \cap F) \cup (H \cap F)) \cup \\
 &\quad ((B \cap M) \cup (E \cap M) \cup (L \cap M) \cup (H \cap M)).
 \end{aligned}$$

The numbers support the University's claim. In each of the four programs, the admission rate of female applicants was significantly higher than that of males: 50.0% versus 42.5% in Business, 33.3% versus 26.7% in Engineering, 20.0% versus 15.0% in Law and 10.0% versus 6.7% in Health Sciences.

4.8 Who would have expected ...

How is it possible that every program favors women and yet women have a lower overall chance of being admitted? The answer is that the programs to which most women were applying were very selective (half the applications were in Health Sciences where the admissions rate was less than 10% and another quarter were in Law where it was less than 20%) while most men were applying to much less competitive schools (here three-quarters of the applications were to Business and Engineering which had admission rates of over 40% and 25% respectively).

Let's restate this in terms of expected values. For example, 200 of the 1000 female applicants to the Law school, or 20%, were admitted. This 20% is the expected value $\mathcal{E}_{(F \cap L)}(Y)$, or as above, the conditional probability $\Pr(A \cap F \cap L \mid F \cap L)$. Likewise, the 50.0% of women admitted to the Business school is $\mathcal{E}_{(F \cap B)}(Y)$. The point is that there's no way to relate the overall chance of 25% that a woman is admitted to the four program based chances 50.0%, 33.3%, 20.0% and 10.0%. To reconstruct the overall chance $\mathcal{E}_{(F)}(Y)$ we need to know just these component expectations but also the sizes of the events (like $F \cap L$) associated to each. It's the different sizes of the program based events (like $F \cap L$ versus $M \cap L$) that allow the uniformly *higher* program based chances that a woman is admitted to combine to a *lower* overall chance.

The moral is that, even if we only have a single random variable in mind, we cannot use an expected value taken over an entire sample space S to draw conclusions about expected values from subsets T of our sample space S , nor can we hope, in general, to combine expected values from subsets T that fill out S to tell us much about the overall expected value.

Here's a problem along the same lines that deals with averages that are close to home for many of you. What average do students pay most attention to? Their GPA, of course. Like the overall admissions figures at Berkeley, the GPA is an average of averages. your GPA is

4.8 Who would have expected ...

built up from your grades in individual courses, each of which is an average of your work in that course—each course grade plays a role like that of a programs at Berkeley. The general view is that a student with a higher GPA is a better student. Does this view hold water? The following problem shows that it does not. Once again, I have made up data to keep the numbers simple, but the conclusion of the problem—a better student can have a lower GPA because he or she takes harder courses—holds in real life, where, for example, the GPA penalizes students majoring in the sciences.

PROBLEM 4.8.54: Consider a University with two classes of courses. In science courses, science majors average a grade of B- (GPA 2.7) and humanities majors average a C- (GPA 1.7). In humanities courses, science majors average a grade of A- (GPA 3.7) and humanities majors average a B+ (GPA 3.3). A science major takes 30 science courses and 10 humanities courses while at the University while a humanities major takes 35 humanities courses and 5 science courses.

- i) Why does this data demonstrate that science majors at this University are better students than humanities majors?
- ii) Show that the GPA of a typical science major is 2.95 while that of a typical humanities major is 3.10.
- iii) Explain how to reconcile the two answers above and discuss how they illustrate the same paradox as the Berkeley discrimination data.
- iv) What GPA would a science major who took the same mix of courses as a humanities major expect? What GPA would a humanities major who took the same mix of courses as a science major expect?

Spreads: variance and standard deviation

In this section, we want to scratch the surface of a problem of critical importance in all kinds of practical testing in virtually every social and physical science. How can we decide when an observed value is

4.8 Who would have expected ...

sufficiently different from an expected one that we should doubt our expectation?

Suppose, as a very typical example, we want to test a new treatment for some form of cancer. The standard protocol is to enroll a set of patients in a trial and then to randomly give patients either the old and new treatment. At this point, you have no evidence that the new treatment is either better or worse than the old, so you start from the *expectation* that both groups of patients will fare equally well. This “makes no difference” starting point is usually called a **null hypothesis**.

The problem is that, even if the **null hypothesis** is true, it’s unlikely that you’ll actually *observe* both groups doing equally well (just as you probably won’t observe 50 heads in 100 tosses). So just observing that the patients receiving the new treatment do better (or worse!) does not contradict our expectation that neither treatment is more effective. For such a trial to be useful, you need to know going in, by *how much* the observed outcomes of the two groups must differ before you can be *confident* that the new treatment is better (or worse) than the old—that is, before you reject the **null hypothesis** that there’s no difference.

It turns out that **expected values** are the key to assessing observations in this way. For any degree of confidence, usually specified as a percentage, expected values can be used to predict how much the observed outcomes of the two groups must differ before you can conclude, with the specified degree of confidence, that one treatment is superior to the other.

It’s common, for example, to speak of a 95% **confidence interval**. This is a range or interval of possible observed values surrounding the expected one. By chance alone, an observed value will lie *inside* this interval in 95% of trials. So we’ll observe a value *outside* this interval only 5% of the time, or 1 time in 20. If the observed difference between the old and new treatments lies outside this interval, then

4.8 Who would have expected ...

we've either observed a relatively unlikely outcome, or the **null hypothesis** is wrong and the one treatment is more effective than the other. If we draw the latter conclusion, we will only be wrong about 1 time in 20.

If we are more skeptical, we can ask for a higher level of confidence—99% is a common one. There'll be a larger 99% confidence interval. Our observation will lie outside this interval only 1 time in 100. So if we conclude that the treatments differ, we'll only be wrong about 1% of the time.

As I've already indicated, to really understand how to make such judgements calls for lengthy study—a course in statistics, or even several. Here we will just study a single example where we have most of the necessary tools in hand. Suppose we are interested in studying the number of successes in a series of n binomial trials. By **TOTAL SUCCESSES IN BINOMIAL TRIALS 4.8.40**, we know that the expected number of successes is just $E(K) = n \cdot p$. How do we judge the likelihood of seeing a observed value of ℓ for this number?

First, let's describe the setup for a general random variable Y a bit. Then we'll get down to brass tacks with the number of successes K in a binomial distribution. It turns out we'll need to use the expected value $\mathcal{E}(Y)$ to create a new random variable related to Y and then calculate the expected value of this new variable. To avoid getting confused about which \mathcal{E} s we have calculated and which we are trying to calculate, it's standard to use the Greek letter μ —mu, pronounced like the cat not the cow, for mean—to indicate an expected value we already know. So we'll set $\mu_Y = \mathcal{E}(Y)$ in the

MEAN 4.8.55: *The mean μ_Y of a random variable is just an alternate term for the expected value $\mathcal{E}(Y)$. We use μ_Y when we want to think of this value as known constant, rather than as something to be computed.*

OK. The expected value Y tells us where we expect our observations to center. To assess how likely a given observation is, what we need

4.8 Who would have expected ...

is a measure of how widely the values are spread around. The plan is to cook up a random variable W whose values measure how far values of Y are from μ_Y , and then find $\mathcal{E}(Y)$.

The first idea you might try is setting $W = Y - C_{\mu_Y}$ —remember C_{μ_Y} is the constant random variable with value μ_Y on every outcome. The value of W is exactly how far Y deviates from its mean, which seems like just what we want. Unfortunately, $\mathcal{E}(W)$ is *always* 0: because [THE EV OF THE SUM IS THE SUM OF THE EVS 4.8.35](#), $\mathcal{E}(W) = \mathcal{E}(Y) - \mathcal{E}(C_{\mu_Y})$ and both these expected values equal μ_Y : the first, by definition, and the second because [CONSTANTS ARE THEIR OWN EVS 4.8.33](#).

We need to find some way to make all the deviations positive and prevent cancellations. One way to do this is by taking an absolute value and setting $W = |Y - C_{\mu_Y}|$. This approach is not used much for two reasons. At a practical level, the absolute value makes computing and working with the corresponding expected values very difficult, because relations involving Y will have no analogue for W . Theoretically, it turns out that this overemphasizes small deviations from μ_Y and underemphasizes big ones. The second way to guarantee that things are positive is to square them.

SQUARED DEVIATION 4.8.56: *The squared deviation of a random variable is the random variable $(Y - C_{\mu_Y})^2$.*

The squared deviation of Y turns out to capture perfectly how far it's values tend to deviate or spread away from μ_Y . It's expected value condenses these squared deviations down to a single number that carries such basic information about Y that it has its own name.

VARIANCE 4.8.57: *The variance $\text{Var}(Y)$ of a random variable is defined by either of the two equivalent formulae:*

i) (Variance First Definition) $\text{Var}(Y) := \mathcal{E}((Y - C_{\mu_Y})^2)$.

ii) (Variance Second Definition) $\text{Var}(Y) := \mathcal{E}(Y^2) - \mu_Y^2$.

It's usually best to think about $\text{Var}(Y)$ as the squared deviation of Y (that is, using the first formula) but to calculate it using the second.

4.8 Who would have expected ...

First, an easy, but useful, remark about this definition. For any Y , the square in $(Y - C_{\mu_Y})^2$ makes this variable greater than or equal to 0 and the same must hold for its expected value $\text{Var}(Y)$. But, looking a bit more closely, we can only get 0 for $\text{Var}(Y)$ if *all* the values of $Y - C_{\mu_Y}$ are 0: in other words, Y equals the constant C_{μ} . So

VARIANCE IS POSITIVE EXCEPT FOR CONSTANTS 4.8.58: If Y is a constant random variable, then $\text{Var}(Y) = 0$. If Y is *not* constant, then $\text{Var}(Y)$ is strictly positive.

Next let's calculate a few variances to get a feel for the two formulae.

EXAMPLE 4.8.59: First we continue [EXAMPLE 4.8.3](#) where our experiment consists of rolling a single die, $Y(x)$ is just the number that comes up, and $\mathcal{E}(Y) = \frac{7}{2}$. We can then calculate $\text{Var}(Y)$ using the [OUTCOMES EXPECTED VALUE FORMULA 4.8.2](#). Pulling out the common factor of $\frac{1}{6}$ from the $\text{Pr}(x)$, then [VARIANCE 4.8.57.i](#)) gives:

$$\begin{aligned} & \sum_{x \in S} \text{Pr}(x) (Y(x) - \mu_Y)^2 \\ &= \frac{1}{6} \left(\left(1 - \frac{7}{2}\right)^2 + \left(2 - \frac{7}{2}\right)^2 + \left(3 - \frac{7}{2}\right)^2 + \left(4 - \frac{7}{2}\right)^2 + \left(5 - \frac{7}{2}\right)^2 + \left(6 - \frac{7}{2}\right)^2 \right) \\ &= \frac{1}{6} \left(\frac{25}{4} + \frac{9}{4} + \frac{1}{4} + \frac{1}{4} + \frac{9}{4} + \frac{25}{4} \right) = \frac{35}{12}. \end{aligned}$$

On the other hand, [VARIANCE 4.8.57.ii](#)) gives $\text{Var}(Y)$ as

$$\begin{aligned} & \left(\sum_{x \in S} \text{Pr}(x) Y(x)^2 \right) - \mu_Y^2 \\ &= \frac{1}{6} (1^2 + 2^2 + 3^2 + 4^2 + 5^2 + 6^2) - \left(\frac{7}{2} \right)^2 = \frac{91}{6} - \frac{49}{4} = \frac{35}{12}. \end{aligned}$$

This illustrates why the second formula is preferred for calculations.

PROBLEM 4.8.60: Here we will use the probabilities you found in [PROBLEM 4.8.17](#), to further analyze the random variable Z that gives your winnings when you play Chuck-a-luck. First find the expected value of Z^2 , then use this and the value of μ_Z for find the variance Var_Z . Finally, show that the standard deviation $\sigma_Z \simeq 1.113$.

Here are a couple of more substantial examples for you to try.



4.8 Who would have expected ...

PROBLEM 4.8.61: This problem is a continuation of [PROBLEM 4.8.46](#) where we considered a bus carrying 5 passengers and with 10 stops on its route. Complete the table below to find $\mathcal{E}(Y^2)$.

Stops y	y^2	Count m	Product y^2m
1		10	
2		1350	
3		18000	
4		50400	
5		30240	
Total		100000	

TABLE 4.8.62: EXPECTED NUMBER OF TIMES THE BUS STOPS

Then use the fact, from [PROBLEM 4.8.46](#), that $\mathcal{E}(Y) = 4.0951$ to show $\text{Var}(Y) = 0.52825599$.

PROBLEM 4.8.63: This is a continuation of [PROBLEM 4.8.63](#), involving the number of Republicans Y on a 5 member committee chosen at random from the 59 Democrats and 41 Republicans in the U.S. Senate. As above add a “squared value” column to [TABLE 4.8.16](#) and use it to find $\mathcal{E}(Y^2)$. Then use $\mu_Y = 2.05$ to show that $\text{Var}(Y) = \frac{45961}{39600} \approx 1.1606\overline{31}$. That bar means a repeating 31.

Next, let’s check that the two formulae for $\text{Var}(Y)$ always agree. This gives the first confirmation that $\text{Var}(Y)$ has nice properties (and incidentally shows why we want to all the trouble to understand relations between expected values in the last subsection). First we expand the square and use the fact that [THE EV OF THE SUM IS THE SUM OF THE EVs 4.8.35](#) to get $\mathcal{E}(Y - C_{\mu_Y})^2) = \mathcal{E}(Y^2 - 2C_{\mu_Y} \cdot Y + C_{\mu_Y}^2) = \mathcal{E}(Y^2) - \mathcal{E}(2C_{\mu_Y} \cdot Y) + \mathcal{E}(C_{\mu_Y}^2)$.

Since C_{μ_Y} is constant, we can rewrite $\mathcal{E}(2C_{\mu_Y} \cdot Y)$, first as $\mathcal{E}(2\mu_Y \cdot Y)$, then as $2\mu_Y\mathcal{E}(Y)$ using the rule that [THE EV OF THE MULTIPLE IS THE MULTIPLE OF THE EV 4.8.34](#), and finally as $2\mu_Y^2$ since $\mathcal{E}(Y) = \mu_Y$ by definition. Likewise, the variable $C_{\mu_Y}^2$ is constant with value μ_Y^2 so

4.8 Who would have expected ...

using [CONSTANTS ARE THEIR OWN EVs 4.8.33](#) gives $\mathcal{E}(C_{\mu_Y}^2) = \mu_Y^2$. Plugging all this in and cancelling, we are left with $\mathcal{E}(Y^2) - \mu_Y^2$.

In these checks, we used heavily the relations between expected values from the previous subsection. It’s natural to ask whether there are similar relations for variances—after all variances are expected values. To get a feel for the answer, let’s compute a few more variances.

PROBLEM 4.8.64: This problem continues [EXAMPLE 4.8.36](#) and asks you to find the variances of the variables studied there by completing the last two columns of the table below.

Random variable Y	Outcome				μ_Y	$\mathcal{E}(Y^2)$	$\text{Var}(Y) = \mathcal{E}(Y^2) - \mu_Y^2$
	HH	HT	TH	TT			
L_1	1	1	0	0	$\frac{1}{2}$		
L_2	1	0	1	0	$\frac{1}{2}$		
$K = L_1 + L_2$	2	1	1	0	1		
C_1	1	1	1	1	1		
$2 \cdot K$	4	2	2	0	1		
$N = K + L_1$	3	2	1	0	$\frac{3}{2}$		

TABLE 4.8.65: VARIANCES FOR TWO TOSSES

The completed table makes it clear that, even though variances are expected values, most relationships between random variables are modified or eradicated in their variances. For example the variance of the constant variable C_1 is not 1 but 0. This holds for any constant variable because the rule that [CONSTANTS ARE THEIR OWN EVs 4.8.33](#) means that $\mu_{C_a} = a$ and hence $C_a - \mu_{C_a}$ is the zero random variable. Likewise, the $\text{Var}(2 \cdot K) = 4 \cdot \text{Var}(K)$. In general, if we multiply a variable by a , its variance gets multiplied by a^2 .

VARIANCE IS QUADRATIC 4.8.66: For any Y , $\text{Var}(aY) = a^2 \text{Var}(Y)$. The variance of a multiple get multiplied not by the multiple a but

4.8 Who would have expected ...

by the square of a .

Looking at the row for $N = K + L_1$ we can see that, in general, the variance of the sum is not the sum of the variances. We have $\text{Var}(N) = \frac{5}{4}$ but $\text{Var}(K) + \text{Var}(L_1) = \frac{1}{2} + \frac{1}{4} = \frac{3}{4}$. On the other hand, looking at the row for $K = L_1 + L_2$, we see that this sum is *sometimes* true: $\text{Var}(K) = \frac{1}{2} = \frac{1}{4} + \frac{1}{4} = \text{Var}(L_1) + \text{Var}(L_2)$.

Once again, the question we need to understand is *when* this relation holds and once again the key is **independence**. This is expressed in the following fundamental formula that is they key to almost all applications of variances.

VARIANCE OF INDEPENDENT SUMS 4.8.67: *If the random variables Z and W are independent, then $\text{Var}(Z + W) = \text{Var}(Z) + \text{Var}(W)$. Informally, the variance of a sum of independent variables is the sum of their variances.*

A moment's thought should make you wonder how this beautifully simple formula can possibly be true in general. After all, we've just proved that variance is more-or-less a "squaring" operation, and despite the fact that thousands of students use the rule that "the square of the sum is the sum of the squares" every day, this rule is just *wrong*! Try it with 1 and 2! So something very special (and important) is going on here. The proof of the formula is actually quite easy and we won't go beyond it to penetrate the mystery completely, so I'll content myself with emphasizing its presence.

So we start calculating. Using **VARIANCE 4.8.57.ii**), $\text{Var}(Y + Z) = \mathcal{E}((Y + Z)^2) - (\mu_{Y+Z})^2$. We'll simplify the two terms separately then combine them.

For the first, we expand and use the rule that **THE EV OF THE SUM IS THE SUM OF THE EVS 4.8.35** to get: $\mathcal{E}((Z + W)^2) = \mathcal{E}(Z^2 + Z \cdot W + W \cdot Z + W^2) = \mathcal{E}(Z^2) + \mathcal{E}(Z \cdot W) + \mathcal{E}(W \cdot Z) + \mathcal{E}(W^2)$. Then, and here's where the mystery hides and where the independence of Z and W is crucial, we use the fact that **THE EV OF THE PRODUCT OF**



4.8 Who would have expected ...

INDEPENDENT VARIABLES IS THE PRODUCT OF THE EVS 4.8.41 to write both $\mathcal{E}(Z \cdot W) = \mu_Z \cdot \mu_W = \mathcal{E}(W \cdot Z)$. Summing up, $\mathcal{E}((Z + W)^2) = \mathcal{E}(Z^2) + \mathcal{E}(W^2) + 2 \cdot \mu_Z \cdot \mu_W$.

The second term is easier. By definition, $\mu_{Z+W} = \mathcal{E}(Z + W)$; because **THE EV OF THE SUM IS THE SUM OF THE EVS 4.8.35**, $\mathcal{E}(Z + W) = \mathcal{E}(Z) + \mathcal{E}(W)$; and, again by definition $\mathcal{E}(Z) + \mathcal{E}(W) = \mu_Z + \mu_W$. So $(\mu_{Z+W})^2 = (\mu_Z + \mu_W)^2 = \mu_Z^2 + 2 \cdot \mu_Z \cdot \mu_W + \mu_W^2$.

Combining the two terms, $\mathcal{E}((Z + W)^2) - \mu_{Z+W}^2 = (\mathcal{E}(Z^2) + \mathcal{E}(W^2) + 2 \cdot \mu_Z \cdot \mu_W) - (\mu_Z^2 + 2 \cdot \mu_Z \cdot \mu_W + \mu_W^2) = (\mathcal{E}(Z^2) - \mu_Z^2) + (\mathcal{E}(W^2) - \mu_W^2)$. But the last two terms are just $\text{Var}(Z)$ and $\text{Var}(W)$ by **VARIANCE 4.8.57.ii**. So $\text{Var}(Z + W) = \text{Var}(Z) + \text{Var}(W)$.

To see how useful **VARIANCE OF INDEPENDENT SUMS 4.8.67** is, let's use it to find the variance of the variable $K = \sum_{i=1}^n L_i$ that counts successes in a sequence of n Bernoulli trials as the sum of the indicator variables L_i that count successes on the i^{th} trial. Since L_i has value 1 with probability p and value 0 with probability $q = 1 - p$ and mean $\mu_{L_i} = p$, **VARIANCE 4.8.57.ii** gives $\text{Var}(L_i) = \mathcal{E}(L_i^2) - \mu_{L_i}^2 = (1^2 \cdot p + 0^2 \cdot q) - p^2 = p - p^2 = p(1 - p) = pq$. Since the variables L_i are independent, **VARIANCE OF INDEPENDENT SUMS 4.8.67** then immediately gives $\text{Var}(K) = \sum_{i=1}^n \text{Var}(L_i) = \sum_{i=1}^n pq = n \cdot pq$. This is so important we record it.

VARIANCE OF SUCCESSES IN BINOMIAL TRIALS 4.8.68:

- i) $\text{Var}(L_i) = pq = p(1 - p)$.
- ii) $\text{Var}(K) = n \cdot pq = n \cdot p(1 - p)$.

PROBLEM 4.8.69: A better feeling for the power of **VARIANCE OF INDEPENDENT SUMS 4.8.67** is given by trying to calculate $\text{Var}(K)$ as

$$\mathcal{E}(K^2) - \mu_K^2 = \sum_{k=0}^n k^2 (C(n, k) p^k q^{(n-k)}) - (np)^2.$$

- i) Evaluate this sum for $n = 8$ and $p = \frac{1}{2}$. You should get 2 by **VARIANCE OF SUCCESSES IN BINOMIAL TRIALS 4.8.68.ii**.



4.8 Who would have expected ...

ii) Evaluate this sum for $n = 9$ and $p = \frac{1}{3}$. You should again get 2. Already just these two examples are much more work than the general derivation above. For a really hard challenge, try to evaluate this sum for general n and p .

PROBLEM 4.8.70: An experiment consists of rolling a die 12 times and recording all 12 outcomes.

i) Show that the sample space S consists of sequences of length 12 in the numbers from 1 to 6 and that $\#S = 6^{12}$.

ii) Consider the events $E_{2,3}$ = “the number on the 2nd roll is 3” and $E_{9,1}$ = “the number on the 9th roll is 1”. Show, *by counting outcomes*, that $\Pr(E_{2,3}) = \Pr(E_{9,1}) = \frac{1}{6}$ and that $\Pr(E_{2,3} \cap E_{9,1}) = \frac{1}{36}$. Conclude that $E_{2,3}$ and $E_{9,1}$ are independent.

What would happen if we changed the rolls from the 2nd and 9th to any pair of *different* rolls? What would happen if we changed the numbers that came up from 3 and 1 to any other values, equal or not?

iii) Consider the random variable Y given by totalling the numbers that come up on all 12 rolls. The variable Y is the sum of 12 variables Z_i whose values are the number on the i^{th} roll. Use this and [EXAMPLE 4.8.3](#) to find $\mathcal{E}(Y)$.

iv) Explain why [ii\)](#) shows that Z_i and Z_j are independent variables if i and j are different rolls.

v) Find $\text{Var}(Y)$ by using [EXAMPLE 4.8.59](#) to give $\text{Var}(Z_i)$ and applying [VARIANCE OF INDEPENDENT SUMS 4.8.67](#).

Find the variance and expected value of the analogous total when 2 dice are rolled and when 60 dice are rolled.

Here’s a problem by way of warning you not to skip a careful check that, when you use [VARIANCE OF INDEPENDENT SUMS 4.8.67](#), the variables you are summing are indeed independent.

PROBLEM 4.8.71: Let’s compare [PROBLEM 4.8.44](#) and [PROBLEM 4.8.63](#) involving the number of Republicans Y on a 5 member committee chosen at random from the 59 Democrats and 41 Republicans

4.9 Everything is back to normal

in the U.S. Senate. Let Z_s be the random variable of [PROBLEM 4.8.44](#) that is 1 if Senator s is on the committee and 0 otherwise. Recall that $\Pr(E_{Z_s,1}) = 0.05$ and hence that $\mathcal{E}(Z_s) = \frac{5}{100} = 0.05$.

Since whether Senator s is chosen for the committee has no effect on the possible choice of Senator t —regardless of the parties of s and t —the variables Z_s and Z_t (that is, the events $E_{Z_s,1}$ that s is selected and $E_{Z_t,1}$ that t is selected) are independent.

- i) Show that $\text{Var}(Z_s) = .0475$ using [VARIANCE 4.8.57.ii](#).
- ii) Use [VARIANCE OF INDEPENDENT SUMS 4.8.67](#) and the fact that Y is the sum of the 41 variables Z_s for which s is a Republican, to show that $\text{Var}(Y) = 41 \cdot .0475 = 1.9475$.

Why doesn't this agree with the answer of $\frac{45961}{39600} \simeq 1.1606\overline{31}$ we got in [PROBLEM 4.8.63](#)? Because the claim I made above that Z_s and Z_t are independent seems right but isn't! So [VARIANCE OF INDEPENDENT SUMS 4.8.67](#) did not apply above.

- iii) Show that for any two Senators s and t there are $C(98, 3) = 152096$ committees containing both s and t .

- iv) Conclude that $\Pr(E_{Z_s,1} \cap E_{Z_t,1}) = \frac{152096}{75287520} = \frac{4}{99} \simeq 0.00\overline{20}$ (that is, the 20 is a repeating decimal).

- v) Show, on the other hand, that $\Pr(E_{Z_s,1}) \cdot \Pr(E_{Z_t,1}) = 0.0025$.

Since these last two answers don't match up, we conclude that $E_{Z_s,1}$ and $E_{Z_t,1}$ —and likewise Z_s and Z_t —are dependent. That $\frac{4}{99}$ explains why: once we pick s , then t is one of 99 Senators vying for 4 spots—not one of 100 vying for 5—so picking s reduced t 's chances by partially filling the committee.

4.9 Everything is back to normal

In this section, we'll look what are probably the two most important theorems in probability, the [LAW OF LARGE NUMBERS 4.9.3](#) and the

4.9 Everything is back to normal

[CENTRAL LIMIT THEOREM 4.9.12](#), and see how these can be applied, using the test case of binomial distributions. We'll be especially interested in understanding the binomial random variable K_n that counts [TOTAL SUCCESSES IN BINOMIAL TRIALS 4.8.40](#) when the number n of trials is large.

These theorems will be exceptions to the general rule in **MATH⁴LIFE**, that we understand why every result we use is true. The reason is that what both theorems tell us about is not the value of a single K_n but the behavior of *all* the values K_n when n becomes large. We can state this more succinctly in language of calculus as saying they predict the limiting behavior of K_n as n goes to infinity. And quite a few tools from calculus are needed to verify these predictions.

Fortunately, applying these predictions to get answers to questions about binomial (and, although we won't consider them here, many other distributions) that come up in just about every area of the physical, medical, social and human sciences turns out to be elementary. We'll finally be able to say when we should think that the difference between an observed value and a probability expectation is small enough that it's probably just random, and when such a difference is large enough that it's highly unlikely to be random. Designing experiments that produce such unlikely differences is the most common way that economic trends are identified, drugs are tested, quality is controlled in manufacturing processes—the list is endless.

Let's start by describing the general setup, and pinning down what each component amounts to for our binomial example.

- i) We start with a trial probability space S and a random variable Z on S with expected value $\mathcal{E}(Z) = \mu$: we'll use the [BERNOULLI TRIAL 4.7.18](#) B_p as our trial space and the indicator variable L for success s on this trial so our mean value μ is just p .
- ii) Next, we consider a sequence of n **independent** trials with sample space \mathcal{T}^n the length n sequences of outcomes in S : our sample

4.9 Everything is back to normal

space will be the binomial distribution space \mathcal{B}^n of sequences of n Bernoulli trials.

iii) On \mathcal{T}^n , we have n independent random variables Z_1, Z_2, \dots, Z_n that are “copies of Z ”—the value of Z_i is just the value of Z on the i^{th} trial—and we can form a new random variable Y_n by summing all of these.

When we start from the indicator variable L , our L_i are the variables of [SUCCESSIONS IN THE \$i^{\text{th}}\$ BINOMIAL TRIAL 4.8.38](#) and their sum is the variable K_n whose value is the total number of successes in all n trials.

iv) By [THE EV OF THE SUM IS THE SUM OF THE EVs 4.8.35](#), the mean of Y_n is given by $\mathcal{E}(Y_n) = \sum_{i=1}^n \mathcal{E}(Z_i) = n\mu$ since each Z_n has expected value μ . By [THE EV OF THE MULTIPLE IS THE MULTIPLE OF THE EV 4.8.34](#), we can just divide both sides by n , getting $\mathcal{E}(\frac{Y_n}{n}) = \mu$. We can informally think of $\frac{Y_n}{n}$ as the *average value* of the variable Z on the sequence of n trials, and the equation $\mathcal{E}(\frac{Y_n}{n}) = \mu$ as saying that we expect this average value to be close to μ .

In our binomial example, $\frac{K_n}{n}$ is the *average proportion* of successes observed in our sequence of n trials and we expect this proportion to be close to p : $\mathcal{E}(\frac{K_n}{n}) = p$.

The statement, “We expect this average value to be close to μ .” comes with a lengthy voiceover of disclaimers. First, when we say we expect something in a probability sense, we only mean that it’s “likely” to happen. We always have to allow for “unlikely” possibilities where we observe something very different from what we expect.

EXAMPLE 4.9.1: If we toss a coin 10 times, we expect to see about 5 heads. This is our binomial examples with $p = \frac{1}{2}$ and $n = 10$. But, of the 1024 outcomes of such an experiment, there’s 1 where the observed number of heads is 10 (and another where it’s 0), and 10 each where it is 9 or 1. Even if we toss the coin 100 times we can’t be sure we won’t see 100 heads (or 1 or 99)—we can only say that these observations are *very* “unlikely”.

4.9 Everything is back to normal

Next, we only really expect the average to be “close” to the expected value μ when the number of trials n is “large”. Even this statement needs to be watered down: how close is “close” and how large is “large”? We can’t, in general, say.

EXAMPLE 4.9.2: Suppose for example that our basic variable Z equals 1,000,000 with probability $\frac{1}{1000000}$ and is otherwise 0. Then $\mu = \frac{1}{1000000} \cdot 1,000,000 = 1$. Now suppose we perform $n = 500,000$ trials. Then we’ll either see 0 outcomes in which Z equals 1,000,000—that is, each individual $Z_i = 0$ and the average value $\frac{Y_n}{n}$ is also 0—or, at least one Z_i will be 1,000,000—and then $\sum_{i=1}^n \mathcal{E}(Z_i) \geq 1,000,000$ so the average $\frac{Y_n}{n}$ is at least $\frac{1,000,000}{500,000} = 2$. To sum up, even though we made a “large” number of trials (half-a-million), our observed average and the expected average will never be as any “closer” than 1 apart.

What makes both the [LAW OF LARGE NUMBERS 4.9.3](#) and, even more, the [CENTRAL LIMIT THEOREM 4.9.12](#) so valuable is that they allow us to be more precise about what “unlikely”, “close” and “large” mean by comparing them to fixed quantities: in fancy terms, they let us *quantify* these terms.

Let’s look at the simpler [LAW OF LARGE NUMBERS 4.9.3](#) first. It introduces two quantities, a big whole number N and real numbers ε and τ . The positive whole number N provides a threshold for “largeness”. A number n of trials is “large” if $n \geq N$. The positive real number ε provides a threshold for “unlikelyness”. The symbol ε is the Greek letter epsilon—the equivalent of a Roman ‘e’—and is the standard way to denote a small error or other quantity. An event is E “unlikely” if $\Pr(E) < \varepsilon$ (that is, we’ll see E happen less than an ε -fraction of the time) and E is “likely” if $\Pr(E) > 1 - \varepsilon$. The positive real number τ (τ is the Greek letter t) provides a threshold for “closeness”. An observed average $\frac{Y_n}{n}$ is close to an expected value μ if $|\frac{Y_n}{n} - \mu| < \tau$. We use the absolute value signs here to make sure that y is within an error τ of μ whether y is bigger or smaller than μ .

4.9 Everything is back to normal

In terms of these, the statement—if the number of trials is “large”, then it’s “likely” that an observed average is “close” to the expected value—can be quantified as:

LAW OF LARGE NUMBERS 4.9.3: *For any positive ε and τ , there is an $N := N_{\varepsilon, \tau}$ for which:*

If $n \geq N_{\varepsilon, \tau}$, the probability that $|\frac{Y_n}{n} - \mu| < \tau$ is greater than $1 - \varepsilon$.

Informally: *By taking enough trials, we can it as likely as we wish that our observed average is as close as we want to the expected or mean value.*

The **LAW OF LARGE NUMBERS 4.9.3** asserts, in a more precise mathematical form, our intuitive sense of what we should expect when we run a lot of trials. But it’s still too vague to be of much practical value. In a practical situation, we’ll usually have an ε and τ in mind—we know how much error we’re willing to tolerate and how much certainty we want to have. If we pick, say, $\varepsilon = 0.01$ and $\tau = 0.2$, then we’re saying that we want the difference between our observed average and the mean to be less than 0.2 at least 99% of the time ($1 - 0.01 = 0.99$). The **LAW OF LARGE NUMBERS 4.9.3** tells we can have this if we run enough trials, but it gives us no clue as to how many trials is enough. In other words, we know there’s an N that works and we’re OK if we run at least N trials, but we have no idea how to find this N so we have no idea how many trials is enough.

What we need is an **effective** result: one that tells us how to find the $N_{\varepsilon, \tau}$ that works for any chosen ε and τ . The **CENTRAL LIMIT THEOREM 4.9.12** is the fundamental result of this type. It works by allowing us to compare the random variable Y_n we’re interested in to a *standard* random variable N called a standard **normal** (or **Gaussian**) random variable on a standard normal probability space or distribution \mathcal{N} .

What’s most amazing about the **CENTRAL LIMIT THEOREM 4.9.12** is that it’s **universal**. By this, I mean that for *any* sample space S and random variable Z that we want to start from, we can effectively compare the variable Y_n that we get by summing n independent copies

4.9 Everything is back to normal

of Z to the *same* variable N on the *same* space \mathcal{N} . No matter how different two starting points Z and Z' are, the corresponding sums Y_n and Y'_n will always match the same standard normal model.

What is the magic normal space \mathcal{N} ? On the one hand, it's something you're already familiar with. It's the famous **bell curve** that you're referring to when you ask instructors if they “grade on a curve”. On the other hand, it's very different from anything we have done to this point because it uses an infinite rather than a finite sample space. In fancier terms, we have been studying discrete probability distributions and \mathcal{N} is a **continuous** distribution. Such distributions are really the main subject studied in probability, but this study requires techniques from calculus. So we won't enter into any of the general theory. Instead, here I will explain just enough about \mathcal{N} for us to state and use the **CENTRAL LIMIT THEOREM 4.9.12**.

The sample space for \mathcal{N} is the set \mathbb{R} of real numbers. Not only is \mathbb{R} infinite, but as explained in **INFINITIES AND AN ARGUMENT FROM The Book**, it's infinite in a really big way. This means that we can't describe a probability distribution by assigning a number $\Pr(x)$ to each real x . There are just too many x for us to add these up and get 1 as required in **PROBABILITY MEASURE 4.2.1**. Well, we could set $\Pr(x) = 0$ for all but a few x , but that's not what we want. We want the probability to be spread out over *all* of \mathbb{R} .




FIGURE 4.9.4: Unscaled graph of $G(x) = \frac{1}{\sqrt{2\pi}}e^{(-\frac{x^2}{2})}$

To arrange this “spread out” probability, we introduce the Gaussian density $G(x) := \frac{1}{\sqrt{2\pi}}e^{(-\frac{x^2}{2})}$. I have included 3 graphs of G . In **FIGURE 4.9.4**, the graph is shown with equal x and y scales. In **FIGURE 4.9.5**, the y -scale has been magnified by a factor of 10 to make it easier

4.9 Everything is back to normal

to see values of G . Even so, the graph gets quite low at the sides, so the third graph shows the “right side” of the graph with the y -scale magnified by 100 and the region from $x = 3$ to $x = 4$ shaded.




FIGURE 4.9.5: Rescaled graph of $G(x) = \frac{1}{\sqrt{2\pi}}e^{(-\frac{x^2}{2})}$




FIGURE 4.9.6: Rescaled graph of “right side” of $G(x)$

Now we want to assign probabilities to events—that is, subsets—of the real numbers using G . As with the outcomes—point of \mathbb{R} —there are too many subsets for us to be able to work with them all. Instead, we’ll be able to get by with assigning a probability only to each closed interval event $[a, b]$. The idea for doing this, reminiscent of how we defined \ln in [AREA DEFINITION OF \$\ln\$ 1.4.21](#), is to use areas under the



4.9 Everything is back to normal

graph of G . So we define $\Pr_G([a, b])$ to be the area under the graph of $G(x)$ from x running from a to b . It turns out that the area under the whole (infinite graph)—that is the probability we assign to the interval $(-\infty, \infty)$ which is all of \mathbb{R} —is exactly the required value 1. In fact, arranging this is why we need the funny factor of $\frac{1}{\sqrt{2\pi}}$ —no other will work—in front of the exponential. From the symmetry of the graph, the probabilities of the positive and negative halves—the intervals $[0, \infty)$ and $(-\infty, 0]$ —are each $\frac{1}{2}$. The probabilities associated to all other intervals have to be found using calculus.




FIGURE 4.9.7: Graph of $G(x)$ with standard bands shaded

FIGURE 4.9.7 shows as shaded areas a few of these interval events. For example, $\Pr_G([-1, 1]) = 0.68268949$ is the area of the “central” band running from -1 to $+1$. Likewise, $\Pr_G([-2, -1]) = \Pr([1, 2]) = 0.13590512$ is the area of either of the lighter “middle” bands and $\Pr_G([-3, -2]) = \Pr_G([2, 3]) = 0.02140023$ is the area of either of the darker “outside” bands. Referring back to **FIGURE 4.9.6**, the shaded band has area 0.00131827 which is $\Pr_G([3, 4])$. Fortunately, we only

4.9 Everything is back to normal

need a very small number of these values to be apply to apply the [CENTRAL LIMIT THEOREM 4.9.12](#). They are listed—where they are needed—below in [TABLE 4.9.15](#), [TABLE 4.9.16](#) and [TABLE 4.9.17](#).

How can we compare a discrete random variable—think of the variable K_N that counts total successes in n binomial trials which takes values from 0 to n —with the standard normal graph G ? This turns out to be amazingly easy. We extend the general notion of expected value, and the special cases mean and variance to quantities like G . I'm not going to try to even give any details here—I'll just state that, once again, this comes down to using calculus to define (and compute) these quantities in terms of areas under graphs related to G . It turns out that $\mathcal{E}(G) = \mu_G = 0$ (not surprising, given the symmetry of the graph about 0) and that $\text{Var}(G) = 1$ —making this work out so simply is the reason for the 2 in the denominator of the exponential defining G .

Now we simply **normalize** Y_n so that it has the same mean and variance as G . To make the mean 0, we just need replace Y_n by $Y_n - nC_{\mu_Z}$. Since Y_n has mean $n\mu_Z$, this shifted version has mean 0 by [THE EV OF THE MULTIPLE IS THE MULTIPLE OF THE EV 4.8.34](#) and [THE EV OF THE SUM IS THE SUM OF THE EVS 4.8.35](#).

Matching the variance takes a moment's more thought. Now that we have the right mean, we need to adjust *without* altering the mean. The way to do this is to rescale $Y_n - nC_{\mu_Z}$ —that is, replace it by a multiple $a \cdot (Y_n - nC_{\mu_Z})$. Because [VARIANCE IS QUADRATIC 4.8.66](#), we need to scale not by the variance of y but by it's square root. This square root is so important in the sequel it has a name.

STANDARD DEVIATION 4.9.8: *For any random variable Z , the standard deviation σ_Z of Z is defined to be the square root of the variance of Z : $\sigma_Z = \sqrt{\text{Var}(Z)}$. For the random variable Y_n given by summing n independent copies of Z , we have $\text{Var}(Y_n) = n \text{Var}(Z)$ and hence, $\sigma_{Y_n} = \sqrt{n \text{Var}(Z)} = \sqrt{n} \sigma_Z$*

4.9 Everything is back to normal

For example, the standard deviation of the binomial variable L for success in a single Bernoulli trial is $\sigma_L := \sqrt{\text{Var}(L)} = \sqrt{pq}$, and the standard deviation of the binomial variable K_n for total successes in a sequence of n such trials is $\sigma_{K_n} := \sqrt{\text{Var}(K_n)} = \sqrt{n \cdot pq} = \sqrt{n}\sigma_L$.

The most important feature of σ_{Y_n} is the \sqrt{n} in the formula $\sqrt{n}\sigma_Z$. Although, on the one hand, it means that the absolute deviation between an observed value for Y_n and its expectation grows with n , it also means that, viewed this deviation gets very small *as a proportion* of μ_{Y_n} . This, as we'll soon see, is very powerful.

PROBLEM 4.9.9: Show that Z has standard deviation $\sigma_Z = 1$ if and only if it has variance $\text{Var}(Z) = 1$.

The \sqrt{n} that relates σ_{Y_n} to σ_Z may seem like an annoyance, but it's the key to the applications of the **CENTRAL LIMIT THEOREM 4.9.12**. The random variable

$$\overline{Y}_n = \frac{(Y_n - nC_{\mu_Z})}{\sqrt{n} \cdot \sigma_Z}$$

again has mean 0 and now also has variance 1.

NORMALIZED SUM 4.9.10: The *normalized sum* \overline{Y}_n of the sum Y_n of n independent copies of the variable Z is defined to be the random variable $\overline{Y}_n = \frac{Y_n - nC_{\mu_Z}}{\sqrt{n}\sigma_Z}$. The normalized sum \overline{Y}_n has expected value or mean $E(\overline{Y}_n) = \mu_{\overline{Y}_n} = 0$, variance $\text{Var}(\overline{Y}_n) = 1$ and standard deviation $\sigma_{\overline{Y}_n} = 1$.

EXAMPLE 4.9.11: For example, the normalization of the of the binomial variable K_n for total successes in a sequence of n Bernoulli trials with probability p —that is, where $Z = L$ — is $\overline{Y}_n = \frac{Y_n - nC_p}{\sqrt{n}\sqrt{pq}} = \frac{Y_n - nC_p}{\sqrt{n \cdot pq}}$.

To complete the comparison of G and Y_n , we assign probabilities to real intervals using \overline{Y}_n . We just define $\text{Pr}_{\overline{Y}_n}([a, b])$ to be the sum of the probabilities of all values y of Y_n whose *normalizations* lie between a and b . Formally,

$$\text{Pr}_{\overline{Y}_n}([a, b]) = \sum_{a \leq \overline{y} \leq b} \text{Pr}(E_{\overline{Y}_n, \overline{y}}).$$

4.9 Everything is back to normal

A better way to think of $\Pr_{Y_n}([a, b])$ —the way we’ll usually use—is to view it as the chance that the value of the original sum Y_n lies between a and b standard deviations away from its mean. We can do this by “unnnormalizing” a and b —that is scaling up by $\sqrt{n}\sigma_Z$ and shifting right by $n\mu_Z$ to get

$$\Pr_{Y_n}([a, b]) = \sum_{\sqrt{n}\sigma_Z a + n\mu_Z \leq y \leq \sqrt{n}\sigma_Z b + n\mu_Z} \Pr(E_{Y_n, y}).$$

In effect we measure deviations from the mean μ in *units* of σ and then we ask, “What is the chance the value of Y_n falls between a and b of these σ units from its mean?” It’s important to remember that a and b can be either positive (which means that we are looking σ units *above* the mean), or negative (which means that we are looking σ units *below* the mean).

CENTRAL LIMIT THEOREM 4.9.12: *For any basic random variable Z , let Y_n be the normalized sum of as above, of n independent trials of Z . Then, there is an integer $N := N_\epsilon$ for which:*

If $n \geq N_\epsilon$, then $|\Pr_G([a, b]) - \Pr_{Y_n}([a, b])| < \epsilon$ for any $[a, b]$.

More informally, if we combine enough trials in Y_n (that is, take N big enough), then we can make its probability distribution “as close as we like” to the universal Gaussian distribution G .

FIGURE 4.9.13 shows what this amounts to. In it, I have overlaid, the probabilities of the normalized variable \overline{K}_{100} for total number of successes in 100 Bernoulli trials with probability $p = 0.40$ and the graph of G . Of course, to make the pictures—especially, the areas that measure probabilities—match up, we need to “normalize” the binomial picture. This means shifting horizontally by $n\mu_Z$ and scaling horizontally by $\frac{1}{\sigma_{K_{100}}} = \frac{1}{\sqrt{n}\sigma_Z}$. This horizontal scaling reduces areas so I have multiplied all the *values* of \overline{K}_{100} by $\sigma_{K_{100}}$ to make the areas, and the graphs visibly match.

The probability of each number of successes from 20 to 60 is shown as a vertical line segment: the height of the segment gives the prob-

4.9 Everything is back to normal

ability of this number of successes (scaled up by $\sigma_{K_{100}}$) and the horizontal position gives the corresponding \bar{y} (shifted and scaled down).




FIGURE 4.9.13: Normal vs. binomial with $n = 100, p = 0.40$

For example, the mean of 40 successes corresponds to $\bar{y} = 0$; this has probability $C(100, 40) \cdot 0.4^{40} \cdot 0.6^{60} = 0.081219145$ which we scale up by $\sigma_{K_{100}} = \sqrt{100 \cdot 0.4 \cdot 0.6} = 4.8989795$ to produce a line of length 0.39789094 which is off by only 0.001 from the central value 0.39894228. The line segment just to the left of -2 corresponds to 30 successes and has height $\sigma_{K_{100}} \cdot C(100, 30) \cdot 0.4^{30} \cdot 0.6^{70} = 4.8989795 \cdot 0.010007504 = 0.049026561$.

But there's no need to check all these numbers. The agreement between the vertical lines and the curve is just too striking.

Let me emphasize the amazing part of this theorem, its universality: we don't need to know *anything* about the basic random variable Z , just that the trials in our series are *independent*. If we have independence, then *whatever* Z we take, the normalized sums Y_n will match the "National Bureau of Standards" Gaussian distribution G , more and more accurately as n gets bigger and bigger.

A warning is also in order here. Exactly *how* big you need to take n does depend on the original random variable Z used, as well as, of course, on how accurately you want the graphs to match. In our binomial examples, it turns out that when p (or q) is very close to 0, then you need a much larger n . [FIGURE 4.9.14](#) shows an example



4.9 Everything is back to normal

with, once again, 100 trials but with the much smaller value of $p = 0.03$. Notice how far the lines overshoot on the left and undershoot on the right. There are far fewer visible lines too because the mean number of successes is only 3—to the left of 0 for example, there are only three lines corresponding to 0, 1 and 2 successes. To the right, there are 97 but they are both mostly too short (even shorter than the little tick between 3 and 4) and too far too the right to be seen.




FIGURE 4.9.14: Normal vs binomial with $n = 100, p = 0.03$

How, then, *do* we know when n is big enough? Why is the [CENTRAL LIMIT THEOREM 4.9.12](#) any more effective a tool than the [LAW OF LARGE NUMBERS 4.9.3](#)? We still need to find that mysterious N_ε . Telling when n is big enough in general is what you’ll learn if you take a statistics course. For our binomial example, we’ll just use the rough [BIG-ENOUGH \$n\$ RULE-OF-THUMB 4.9.18](#) which I’ll explain when we come to applications.

What makes it possible to find a N that works in applications is the fact that the [CENTRAL LIMIT THEOREM 4.9.12](#) lets us compare the sum of random variables we are interested in to a fixed target, the standard normal distribution G . We simply tabulate probabilities that G has values in some standard intervals and then, by “de-normalizing”, translate these intervals back to intervals that tell us about the variable Y_N that we’re interested in.

[TABLE 4.9.15](#) contains the most commonly used values: to make the

table more readable I have used μ and σ for $\mu_{Y_n} = n\mu_Z$ and $\sigma_{Y_n} = \sqrt{n}\sigma_Z$. First, we list the probabilities of being “within” an interval about the mean (that is, close to the mean).

Informal name	Y_n interval	G interval	G - Probability
within σ	$[\mu - \sigma, \mu + \sigma]$	$[-1, 1]$	0.682689
within 2σ	$[\mu - 2\sigma, \mu + 2\sigma]$	$[-2, 2]$	0.954500
within 3σ	$[\mu - 3\sigma, \mu + 3\sigma]$	$[-3, 3]$	0.997300
within 4σ	$[\mu - 4\sigma, \mu + 4\sigma]$	$[-4, 4]$	0.999937

TABLE 4.9.15: “WITHIN” GAUSSIAN PROBABILITY VALUES

In our applications, we want to use this information to say that an observed value is unlikely, or very unlikely. This will be the case, not if the observation is close to the mean, but if it’s far away from it. So it is more convenient to display the information in [TABLE 4.9.15](#) using the complementary “beyond” or far from the mean ranges and the complementary probabilities as in [TABLE 4.9.16](#). For example the 0.317311 in the first row of this table is just 1 minus the 0.682689 from the first row of [TABLE 4.9.15](#). The middle two rows are the ones most commonly used in applications.

Informal name	Y_n interval	G interval	G - Probability
beyond σ	$(-\infty, \mu - \sigma) \cup (\mu + \sigma, \infty)$	$(-\infty, 1) \cup (1, \infty)$	0.317311
beyond 2σ	$(-\infty, \mu - 2\sigma) \cup (\mu + 2\sigma, \infty)$	$(-\infty, 2) \cup (2, \infty)$	0.045500
beyond 3σ	$(-\infty, \mu - 3\sigma) \cup (\mu + 3\sigma, \infty)$	$(-\infty, 3) \cup (3, \infty)$	0.002700
beyond 4σ	$(-\infty, \mu - 4\sigma) \cup (\mu + 4\sigma, \infty)$	$(-\infty, 4) \cup (4, \infty)$	0.000063

TABLE 4.9.16: “BEYOND” GAUSSIAN PROBABILITY VALUES

There’s one more refinement. In most applications, we only really care about observed values that are not only far from the expected ones (the mean), but to one side of it—that is, either far “above” or far “below” the mean.

For example, suppose we are trying to decide whether a new medical procedure is more effective than an old one. We compare the out-

4.9 Everything is back to normal

comes of two groups of patients, one getting the new treatment and the other the old one. If the results from the two groups are close to each other, we don't draw any conclusion, attributing the small difference to random effects. This is a "within" result and hence fairly likely to occur by chance.

We consider the results to be unlikely to be due to chance only if the two groups are show result that are far different—a "beyond" result. But there are two kinds of "beyond". Patients getting the new treatment can do a lot better than those who got the old one. This this kind of unlikely "above" observation is what would tend to make us think the new treatment is more effective than the old one. But patients getting the old treatment might very well do a lot better than those who got the new one. This "below" result is also a "beyond" result, but not one that argues for the new treatment. So we really want to ask how likely is the "above" observation.

Informal name	Y_n interval	G interval	G - Probability
more than σ above (or below)	$(\mu + \sigma, \infty)$	$(1, \infty)$	0.158655
more than 2σ above (or below)	$(\mu + 2\sigma, \infty)$	$(2, \infty)$	0.022750
more than 3σ above (or below)	$(\mu + 3\sigma, \infty)$	$(3, \infty)$	0.001350
more than 4σ above (or below)	$(\mu + 4\sigma, \infty)$	$(4, \infty)$	0.000032

TABLE 4.9.17: "ABOVE OR BELOW" GAUSSIAN PROBABILITY VALUES

Mathematically, this is again easy to derive from the numbers we have. Since the Gaussian distribution is symmmetric about the mean (the above side is the mirror image of the below), the probability of being either "above" or "below" is just half the probability of being "beyond". The fact that these numbers are smaller is an advantage: the same observation is more unlikely (half the probability). So when are only interested in one-sided differences, it's smart to test for them as we are more likely to be able to find good evidence. The table below gives the corresponding numbers. Again, the middle two rows are the ones most commonly used in applications.

4.9 Everything is back to normal

The numbers in these tables are often referred to rather loosely. For example, it's common to say that observed values lie "within 2σ " of expected ones is often loosely stated as "with 95% certainty" or that this is a "95% confidence interval". Likewise, the chance 0.001350 of an observation that is "more than 3σ above" an expectation is often loosely stated as "about one in a thousand".

All we need now is some way to tell when n is big enough that we can use these probabilities to analyze the total successes in n trials binomial random variable K_n . We'll just use the following rule-of-thumb which is based on simply calculating probabilities for K_n for many n and p and looking for those furthest from agreeing with normal probabilities in the table. It's actually rather conservative (usually the agreement between probabilities for K_n and for g is much closer than that claimed) but it will more than suffice for our purposes.

BIG-ENOUGH n RULE-OF-THUMB 4.9.18: If the number of trials in a binomial distribution is n and the probability of success in each trial is p , then n is big enough to apply the **CENTRAL LIMIT THEOREM 4.9.12** if $npq > 10$.

More precisely, for such n and p , the entries in the second rows of **TABLE 4.9.15**, **TABLE 4.9.16** and **TABLE 4.9.17** are within ± 0.015 of the true probability that K_n lies in the corresponding interval and those in the third rows are within ± 0.004 .

UNLIKELY SUCCESSES PERCENTAGES 4.9.19: When $npq > 10$, the chance of observing a value of K_n that is:

- i) less than $2\sigma_{K_n}$ from the expected value μ_{K_n} is more than 94%;
- ii) more than $2\sigma_{K_n}$ from μ_{K_n} is less than 6%;
- iii) more than $2\sigma_{K_n}$ above (or below) μ_{K_n} is less than 3%.

The chance of observing a value of K_n that is:

- i) less than $3\sigma_{K_n}$ from μ_{K_n} is more than 99%;
- ii) more than $3\sigma_{K_n}$ away from μ_{K_n} is less than 1%; and,
- iii) more than $3\sigma_{K_n}$ above (or below) μ_{K_n} is less than 0.5%.

4.9 Everything is back to normal

When analyzing deviations between observations and expectations in formal scientific situations, it's standard to give these kind of numerical measures of likelihood. But, in everyday life, we just want to distinguish between informal degrees of likelihood. The following rule-of-thumb tells us how to do this.

UNLIKELY SUCCESSES RULE-OF-THUMB 4.9.20: *Suppose $npq > 10$.*

- i) If an observed value of K_n is within $2\sigma_{K_n}$ of the expected value μ_{K_n} , then the difference is reasonably likely to be due to chance alone.*
- ii) If an observed value of K_n is further than $2\sigma_{K_n}$ from the expected value μ_{K_n} , then the difference is unlikely to be due to chance alone.*
- iii) If an observed value of K_n is further than $3\sigma_{K_n}$ from the expected value μ_{K_n} , then the difference is extremely unlikely to be due to chance alone.*

As a first example of how we can use this, let's see what it says about tossing a coin. Here $\mathcal{E}(L) = \mu_L = p = \frac{1}{2}$, $\text{Var}_L = \frac{1}{2} \cdot \frac{1}{2}$ and $\sigma_L = \sqrt{\text{Var}_L} = \frac{1}{2}$. Now we can finally quantify what it means to say that if we toss a coin 100 times (this, of course, is n) we expect “about” 50 heads. We expect about 50 heads because $\mu_{K_n} = n \cdot \mu_L = 100 \cdot \frac{1}{2} = 50$. And the standard deviation is $\sigma_{K_n} = \sqrt{n} \cdot \sigma_L = \sqrt{100} \cdot \frac{1}{2} = 5$.

We can now calculate the intervals referred to in the [TABLE 4.9.15](#), [TABLE 4.9.16](#) and [TABLE 4.9.17](#): we just start at the center μ_{K_n} and go up or down by a whole number multiple of σ_{K_n} . Warning: σ is almost never a whole number as above, so you need to work most of these examples using decimal bounds. For example, “within 2σ ” is the interval $50 \pm 2 \cdot 5$ or $[40, 60]$; and “more than 3σ above” means above $50 + 3 \cdot 5 = 65$ and gives the interval $(65, 100]$.

Now, $n \cdot p \cdot q = 25$ here, so the binomial probabilities we are asking about are close to Gaussian probabilities. We can therefore translate the statements above about K_n into statements about the coin. Using [UNLIKELY SUCCESSES PERCENTAGES 4.9.19](#), we expect to observe a number of heads between 40 and 60 at least 94% of the time. We

4.9 Everything is back to normal

expect to see more than 65 heads about 0.5% of the time—about one time in two-hundred. In these two cases, the actual binomial probabilities are 98% and 0.2%. We don't really mind these gaps because they're on the right side. We want to classify events as “likely” and “unlikely” and we're making the right calls if, in real life, “likely” events are even likelier (and unlikely ones less likely) than we predict.

But **UNLIKELY SUCCESSES RULE-OF-THUMB 4.9.20** lets us classify observations as “likely” and “unlikely” with no need to remember any numerical probabilities. Any number of heads between 40 and 60 is “reasonably likely” to be due to chance alone. More than 60 heads (or fewer than 40) constitute an observation “not likely” to be due to chance alone. More than 65 heads (or fewer than 45) constitute an observation “extremely unlikely” to be due to chance alone.

Now comes the last key element. As the number n of trials increases, the “reasonably likely” interval gets “tighter and tighter”. To see how this works, suppose we toss the coin 10,000 times—now n is 100 times bigger. The expected number of heads μ_{K_n} goes up by factor of n from 50 to 5,000 but σ_{K_n} only increases by a factor of $\sqrt{n} = 10$ from 5 to 50. Correspondingly, the expected width of the “within 2σ ” and “outside 3σ ” bands only increases by a factor of 10. For example, “within 2σ ” is the interval $5000 \pm 2 \cdot 50$ or $[4900, 5100]$; and “more than 3σ above” means above $5000 + 3 \cdot 50 = 5150$ and gives the interval $(5150, 10000]$.

What's more at this point n is big enough that for all practical purposes the Gaussian G and binomial \overline{K}_n coincide. So the chance is better than 95% that the observed number of heads will be between 4900 and 5100, and it's less than 0.001350 that we'll see more than 5150 heads.

PROBLEM 4.9.21: Use **TABLE 4.9.15** and **TABLE 4.9.17** to estimate the chance of seeing between 4950 and 5050 heads, and the chance of seeing more than 5200 heads.



4.9 Everything is back to normal

Why do I say that the intervals when $n = 10,000$ are “tighter” than those with $n = 100$? In fact, the former are 10 times as big as the latter, because σ_{K_n} goes up by 10. What I mean is that the *relative* size of these intervals—their length expressed as a *fraction or percentage of the expected value* μ_n is getting smaller and smaller.

With 100 tosses, $2\sigma = 10$ so we need to allow a deviation from the mean of ± 10 to catch 95% of observations; and 10 is 20% of the mean 50. With 10,000 tosses, $2\sigma = 100$ so we only need to allow a deviation of ± 100 from the mean of 5,000. But this deviation is now only 2% of the mean versus 20% when we had $n = 100$. When n goes up by a factor of 100, the *absolute* length goes up by a factor of 10 (from 10 to 100 in our example) but the *relative or percentage* length goes *down* by a factor of 10. And, as we’ll see in the examples that follow, it’s this relative length that we usually care about.

In general, if we increase n by a factor f (above, we had $f = 100$), then μ also increases by a factor of f . Of course, σ also increases, but only by a factor of \sqrt{f} . Thus ratios like $\frac{2\sigma}{\mu}$ and $\frac{3\sigma}{\mu}$ that measure the *relative or percentage* size of the intervals where “likely” observations get *smaller* by a factor of \sqrt{f} : $\frac{2\sqrt{f}\sigma}{f\mu} = \frac{1}{\sqrt{f}} \cdot \frac{2\sigma}{\mu}$. That’s the magic of the square root in σ .

EXAMPLE 4.9.22: Consider a series of n binomial trials by spinning a roulette wheel with S given by “seeing a black number”—so, as in [PROBLEM 4.8.21](#) $p = \frac{18}{38}$ —and the associated total number of successes variable K_n which is the sum of n independent Bernoulli success variables L .

- i) Find the μ_L , Var_L and σ_L and use them to give formulae for μ_{K_n} , Var_{K_n} and σ_{K_n} .
- ii) Find the “within 2σ ” and “more than 3σ above” intervals for:
 - a. $n = 400$.
 - b. $n = 40,000$.
 - c. $n = 4,000,000$.

4.9 Everything is back to normal

iii) Use the answers in the previous part and the **UNLIKELY SUCCESSES RULE-OF-THUMB 4.9.20** to justify the qualitative statements below.

- If you bet black at roulette 400 times tonight, you are reasonable likely to go home a winner. Hint: You expect to win about 189.47 spins and to go home a winner you need to observe at least 201 wins.
- If you go to the casino two nights a week—call it 100 times a year—and each night you play roulette 400 times, you're extremely unlikely to go home a winner on the year.
- Suppose a casino's roulette tables host 10,000 players a month and each player bets 400 times for a total of 4,000,000 bets. The casino expects to lose 1,894,736 of these bets and win the other 2,105,263 so it expects to make \$210,527. But the casino is extremely unlikely to make less than \$200,000.

Solution

Here, keeping 5 places, we have $p = \frac{18}{38} \simeq 0.47368$ and $q = \frac{20}{38} \simeq 0.52632$.

i) Using the values for p and q and **VARIANCE OF SUCCESSES IN BINOMIAL TRIALS 4.8.68**, we have $\mu_L = 0.47368$, $\text{Var}_L = 0.24931$ and $\sigma_L = 0.49930$ and use them to give formulae for $\mu_{K_n} = 0.47368n$, $\text{Var}_{K_n} = 0.24931n$ and $\sigma_{K_n} = 0.49930\sqrt{n}$. For example, for $n = 400$, $\mu_{K_n} = 0.47368 \cdot 400 = 189.47$, and $\sigma_{K_n} = 0.49930 \cdot 20 = 9.9861$.

ii) Thus “within 2σ ” and “more than 3σ above” intervals are:

- for $n = 400$, (169.50, 209.45) and (219.43, ∞);
- for $n = 40,000$, (18748, 19147) and (19247, ∞);
- for $n = 4,000,000$, $(1.8927 \cdot 10^6, 1.8967 \cdot 10^6)$ and $(1.8977 \cdot 10^6, \infty)$.

iii) We justify the statements as follows:

- The reasonably likely “within 2σ ” interval for K_{400} includes the range from 201 to 209 in which you are a winner.
- Here the range in which you are a winner starts at $K_{40000} = 20,000$ which is a lot “more than 3σ above” and hence extremely



4.9 Everything is back to normal

unlikely to occur.

- c. For the casino to make less than \$200,000, we'd have to observe a number $K_{4000000}$ of player wins above $1,900,000 - 2,100,000 - 1,900,000 = 200,000$. Again this is a lot “more than 3σ above” so is extremely unlikely.

This kind of qualitative assessment of likeliness is what we've been looking for, and it's all that needed in most cases. But, by going back to the tabulation of the Gaussian distribution we can get more quantitative assessments too. The next example illustrates how.

PROBLEM 4.9.23: Use [EXAMPLE 4.9.22](#) and the values in [TABLE 4.9.17](#) to justify the claims below.

- i) If you bet black at roulette 400 times tonight, you expect to lose about \$21 and you have bit less than a 15% chance of going home a winner.
- ii) If you go to the casino two nights a week—call it 100 times a year—and each night you play roulette 400 times, the chance you'll lose at least \$1700 on the year is over 97% and that you'll lost at least \$1500 is more than 99.86%.
- iii) Suppose a casino's roulette tables host 10,000 players a month and each player bets 400 times for a total of 4,000,000 bets. The casino expects to lose 1,894,736 of these bets and win the other 2,105,263 so it expects to make \$210,527. How likely are they to make more than \$219,000?

Partial Solution

For example, take the case when $n = 400$ and we have $\mu_{400} = 189.47$ and $\sigma_{400} = 9.9861$. So we expect to bet \$400 and $2 \cdot 189.47$ for a net loss of \$21.06. But to win, we just have to observe a K_{400} that's a bit more than “ σ above” (i.e. above 199.46), and the first line of [TABLE 4.9.17](#) says that this will happen about 15% of the time.

Likewise, $n = 40,000$, you expect to win 18,947 spins and lose about \$2106. To lose less than \$1700, you'd need to win at least

4.9 Everything is back to normal

19,150 spins. This is “more than 2σ above” so occurs, by the second line of TABLE 4.9.17, only about 2.3% of the time. You can handle the other parts similarly, but using the third and fourth lines.

That last problem puts the universal message of the LAW OF LARGE NUMBERS 4.9.3 and CENTRAL LIMIT THEOREM 4.9.12 in a nutshell. In a binomial experiment with rather few trials (small n), we can observe values quite far in relative or percentage terms from our expectation. Even though the house advantage on roulette is over 5% and we expect to lose \$21, we still have an almost 15 percent chance of beating the house over 400 spins.

As the number of trials increases, any observations we are likely to make become more and more tightly grouped—again, in a percentage sense—so that if we make 40,000 bets over a year, we expect to lose \$2,106 but we no longer have *any* practical chance of winning and we can be pretty sure our losses won’t be more than a few hundred dollars from this expectation. Finally, at the level of the casino, not merely are we certain that the players will lose, but we can accurately predict how much they will lose. If there’s a discrepancy of as little as 4% up or down in this handle, then the casino knows that someone is cheating! Yes, \$219,000 is more than they expected: too much more, and such a gap is so unlikely to happen by chance that we’re damn sure the binomial probabilities no longer apply—somebody’s cheating!

Here are a few more problems to give a feel for working with setting intervals to achieve desired levels of confidence.

PROBLEM 4.9.24: A math professor has decided to give a multiple choice final on which there will be 75 questions, each with 3 answers. He assumes that getting the right answers on any 2 questions are independent events.

i) The professor wants to set the minimum number of correctly answered questions for a pass high enough that a student who guesses



4.9 Everything is back to normal

all the answers will have little chance of passing. He decides that his cutoff should be a score “more than 3σ above” the expected number of correct answers of a student who did guess all the questions. Show that the passing grade should be 38. Hint: Each guessed question is a binomial trial with a probability of success—guessing the right answer—of $p = \frac{1}{3}$.

ii) How likely is it that a student who has a 50% chance of answering any question correctly will fail?

iii) Show that a student must be able to answer roughly 60% (or more) of the questions correctly for his or her chance of failing the test to be less than 5–6%? Hint: The passing grade needs to be “more than 2σ below” the student’s expected number of correct answers.

iv) Discuss the soundness of the professor’s assumptions of independence:

a. for the student who guesses all the answers.

b. for students who do their best to answer correctly and have a 50% or a 60% chance of getting any question correct.

PROBLEM 4.9.25: A math professor has decided to give a multiple choice final on which there will be 100 questions, each worth 1 point. He assumes that getting the right answers on any 2 questions are independent events.

i) The professor wants to set the minimum number of correctly answered questions for an A low enough that a student who has 95% chance of answering any question correctly will have earned an A with probability 99% or higher. What’s the highest he can set the A grade? Hint: The passing grade needs to be “more than 3σ below” the student’s expected number of correct answers.

ii) In the previous question, we considered a binomial distribution with $p = 0.95$ and $n = 100$. Show that this example falls outside the **BIG-ENOUGH n RULE-OF-THUMB 4.9.18** by finding npq .

iii) To check the cutoff for A’s, we can just use the **BINOMIAL DISTRIBUTION FORMULA 4.7.23** to directly compute probabilities for 100



4.9 Everything is back to normal

binomial trials with a probability p of success of 0.95. Show that the probability of

- a. at least 89 successes is about 99.57.
- b. at least 90 successes is about 98.85.

These answers illustrate that even when the **BIG-ENOUGH n RULE-OF-THUMB 4.9.18** does not hold, the agreement between a binomial distribution and its gaussian approximation is often—if not always—very good.

As an final example of what the **UNLIKELY SUCCESSES RULE-OF-THUMB 4.9.20** means, let's go back and look at the data from **CHUCK-A-LUCK**. This example is *not* binomial. However, we can still apply the **CENTRAL LIMIT THEOREM 4.9.12** to it. We will work with the random variable Z that gives our winnings in a single spin of the Chuck-a-luck cage. The expected value and standard deviation of Z were calculated in **PROBLEM 4.8.17** and **PROBLEM 4.8.6o** as $\mu_Z \simeq -\$0.078704$ and $\sigma_Z \simeq 1.1132$ —I've used more accurate values in the table below but rounded to 5 place accuracy. Now we want to turn to the variable Y_n we get by adding n independent copies of Z : in other words, by totaling the winnings (or losses) from playing Chuck-a-luck n times.

PROBLEM 4.9.26: Verify the bold entries in the table below which summarizes Y_n for various choices of n .

n	100	10,000	1,000,000
μ_{Y_n}	-7.8704	-787.04	-78,704
σ_{Y_n}	11.132	111.32	1113.2
$(\mu_{Y_n} - \sigma_{Y_n}, \mu_{Y_n} + \sigma_{Y_n})$	(-19.002, 3.2618)	(-898.32, -675.68)	(-79813, -77587)
$(\mu_{Y_n} - 2\sigma_{Y_n}, \mu_{Y_n} + 2\sigma_{Y_n})$	(-30.134, 14.394)	(-1019.6, -564.36)	(-80926, -76474)
$(\mu_{Y_n} - 3\sigma_{Y_n}, \mu_{Y_n} + 3\sigma_{Y_n})$	(-41.265, 25.525)	(-1121.0, -453.05)	(-82040, -75360)
$(\mu_{Y_n} - 4\sigma_{Y_n}, \mu_{Y_n} + 4\sigma_{Y_n})$	(-52.397, 36.657)	(-1232.3, -341.73)	(-83153, -74247)

TABLE 4.9.27: STANDARD INTERVALS FOR CHUCK-A-LUCK

Now we'd like to compare these answers to the observations in **CHUCK-A-LUCK**. First let's look at the case $n = 100$. Combining the

4.9 Everything is back to normal

intervals from TABLE 4.9.27 with the probabilities in TABLE 4.9.17, we get the prediction that about “15.8” of the 100 observations in the dataset in TABLE 2.1.23 should lie below -19.002 and about the same number should lie above 3.2618 . In both cases, we observe 15 such counts. Likewise, we’d expect to find 95 of these observations inside $(-30.134, 14.394)$ and we see 94. We expect that only about 0.27% of observations (or 1 in 400) will lie outside $(-41.265, 25.525)$ and none of our 100 do.

Our expectations are the same for the data in TABLE 2.1.24, but here the data agrees a bit less well. There are 21 observations below -19.002 but only 12 above 3.2618 , and only 3 outside $(-30.134, 14.394)$.

Does this mean that this second set of observations is “worse” than the first? Not at all! We’re once again dealing with the difficulty of making statements about theoretical probabilities on the basis of some observations. We *expect* to find about 95 of any 100 observations in the interval $(-30.134, 14.394)$ observing 94 or 97, as in the data sets, is likely to be due to random variation. Could we test our expectation? Yes, though I won’t do so here. Why? Because it would take us too far afield: we’d have to take a large sample of such data sets (not just 2) and look at the distribution of the number of observations outside $(-30.134, 14.394)$ across the sample, using the tools in this section.

Next let’s look at how the observations in TABLE 2.1.25 and TABLE 2.1.28. Both of these involve $n = 10,000$ spins. Here we expect to see about “15.8” observation below, and the same number above, the interval $(-898.32, -675.68)$: the actual counts are 17 and 19 below and 13 and 9 above. We expect to see about 95 entries in the interval $(-1019.6, -564.36)$ and observe 91 and 93. Recall that there is one really surprising entry in TABLE 2.1.28: the number -361 is almost 4σ above the expected value ($\frac{-361 - (-787.04)}{111.32} \simeq 3.8$). The chance of that an observation is this far above the mean is only about 1 in



4.9 Everything is back to normal

15,000. So our set of 200 observations contains one very rare bird. Such extreme observations are usually called **outliers**; our example is almost one-and-a-half σ 's above the next highest observation.

PROBLEM 4.9.28: Compare the observations in TABLE 2.1.28 with the intervals for $n = 1,000,000$ in PROBLEM 4.9.26 and discuss how well these observations agree with the expectations from TABLE 4.9.17.

Let me conclude this example by emphasizing the basic principle that as the number of trials n goes up, our observations become relatively more tightly grouped about the mean. When $n = 100$, the spread of the 100 observations is much bigger than the mean; for $n = 10,000$, the spread of the 100 observations is about the same size as the mean; for $n = 1,000,000$, the spread of the 100 observations is less than 8% of the mean.

Here's a final problem in which you can see how we might apply the ideas underlying the **CENTRAL LIMIT THEOREM 4.9.12** to a practical decision.

PROBLEM 4.9.29: You want to decide whether a prep course is an effective means of raising your expected score on the Verbal SAT. You know that the mean score on the Verbal SAT is 550 with a standard deviation of 130. For each set of prep course results described below, decide how likely it is that the results demonstrate that the course will raise your Verbal SAT score.

Hint: Here our basic random variable Z has mean $\mu_Z = 550$ and standard deviations $\sigma_Z = 130$, and we want to test a random variable Y_n which is a sum of n copies—assumed independent—of this Z .

- i) A sample of 25 graduates of course I scored an average of 575.
- ii) A sample of 25 graduates of course II scored an average of 625.
- iii) A sample of 225 graduates of course II scored an average of 575.
- iv) A sample of 225 graduates of course II scored an average of 625.

Chapter 5

Time is money

Everyone has heard the phrase, “Time is money”. If so, then just as there is a rate for converting between two different currencies like dollars and pesos, there ought to be a rate from converting between time and money. There is! Such a rate is called an **interest rate** and the use of such rates is the subject of this chapter.

Before we start, I'd like to mention that the material covered in this chapter is the one part of **MATH⁴LIFE** which just about every student finds useful in later life. If you ever use a credit card, buy a house, have a retirement plan at your job, take out an insurance policy, or save for your children's college education, you'll be able to make use of what we are going to learn here. And if you don't understand how interest affects you, you can be pretty sure that people who do will be taking advantage of your ignorance to their profit.

5.1 Simple interest

Why do we feel that “time is money”? The basic reason is that money can be used to take advantage of opportunities. Such opportunities



5.1 Simple interest

may be as small as the chance to buy a soda between classes, or as important as the chance to buy the house of your dreams but the principle is the same. If you don't have a buck in the first case or a bundle in the second, you can't take advantage of the opportunity. Ogden Nash put it very nicely: "*Money may not buy happiness, but have you ever tried to buy happiness without money*". So if I give you money now and you give the exact same amount of money back to me later I still lose out on the deal: while you are holding my money, I can't use it to take advantage of opportunities that interest me. This is a qualitative loss. The goal of the mathematics of interest is to put a numerical value on what I give up when I lend you the money. Then, we can strike a fair deal. You give me back what I loaned you *plus interest* in an amount which we agree roughly equals the cost of my missed opportunities and we're both happy. All that's left is to find an amount of interest we can agree on. To begin with, we can agree to denote the interest by I .

What should this amount depend on? Two principles are pretty clear. The more money I lend you and the longer you keep it the greater the opportunities I miss and the more I interest I will need to convince me to lend you the money. We're going to make two somewhat stronger assumptions. To state them we introduce two important definitions:

AMOUNT 5.1.1: *The amount A of a loan is the number of dollars loaned.*

TERM 5.1.2: *The term T of a loan is the length of time for which the money is loaned.*

So if I lend you a dollar until lunch time the amount A is \$1 and the time T is a few hours and if I lend you a \$100,000 for 10 years the amount A equals \$100,000 and the term T equals 10 years. While we almost always use dollars as the units for amounts, we will find it convenient to use many different units—all very familiar—to measure terms and call the unit used a *period*.



5.1 Simple interest

PERIOD 5.1.3: *A period is any unit of time used to measure the term of a loan.*

The most common periods are years, months and days but **quarters** (periods of three months) and **semesters** (periods of 6 months) will also come up. When we want to talk about the units of time used to measure a term without making a specific choice we will call them *periods*. Now we can state our basic principles:

EQUALITY OF DOLLARS 5.1.4: The interest I on a loan should be proportional the amount A loaned. In other words, every dollar earns the same interest during the period of the loan.

EQUALITY OF PERIODS 5.1.5: The interest I on a loan should be proportional to the term T of the loan. In other words, the interest earned during any two days of the loan should be the same, as should the interest earned in any two months, or any two years, and so on.

Remember that saying a quantity I is proportional to a quantity J means that $I = pJ$ where p does not depend on J at all. Mathematically, our two equality principles can thus be summarized by saying that interest I satisfies the equation

SIMPLE INTEREST FORMULA 5.1.6: $I = p \cdot A \cdot T$

where p does not depend on either the amount nor the term of the loan.

What does p depend on? Our answer is going to be: nothing. In other words,

PERIODIC RATE 5.1.7: *The quantity p is a constant called the periodic interest rate or periodic rate for short.*

Warning: although in any given interest calculation, the periodic interest rate p will be constant, we will definitely want to use different constants in different problems. Why do we call the constant of proportionality p a *periodic interest rate*? To answer this question, let's



5.1 Simple interest

look at the units of the [SIMPLE INTEREST FORMULA 5.1.6](#). On the left side, interest I is measured in dollars. So is the amount A on the right. The term T is measured in periods. So to make the units come out, we must have

$$\text{\$} = \text{units of } p \times \text{\$} \times \text{periods}$$

or

$$\text{units of } p = \frac{1}{\text{periods}} = \text{fraction-per-period}.$$

Remember that a *rate* generally stands for a ratio or fraction. Since the denominator of the rate p is measured in periods we see why p is called a periodic rate. It's the *fraction* of the amount I lend you which you must pay me as interest for each *period* you hold my money.

In the real world, people want to be able to discuss interest rates in general without worrying about the time units, or periods used to measure the term of a particular loan. This is generally done by quoting all interest rates as “per-years”. In most situations, the nominal interest fraction r is fairly small. It ranges from about .025 for money you lend to your bank by depositing it in your bank account, to about .19 for money the bank lends you by allowing you not to pay off your credit card bill every month. Since people prefer to work with whole numbers rather than small decimals, we usually talk about the interest rate not as an absolute fraction per year but as a *percentage* per year. In other words, we multiply the absolute fraction by 100 : so your bank pays you 2.5% per year on your deposits but charges you 19% per year on your credit card balance. To keep the difference in units between the fraction-per-period rate p we'll use in calculations and the informal percent-per-year rates clear we call the second kind *nominal interest rates* and use the letter r to denote them.

NOMINAL RATE 5.1.8: *A nominal rate (also called a nominal interest rate), denoted by the letter r , is one which expresses interest as a percentage-per year without reference to the periods used to measure the term or calculate the interest owed.*



5.1 Simple interest

Fine: in everyday life we speak of the nominal or r kind of interest rate: percent-per-year or often just a percentage with the per-year understood. But when we make any interest calculation in **MATH4LIFE**—for example, using the **SIMPLE INTEREST FORMULA 5.1.6**—we must work with the periodic or p kind of interest rate which is an absolute fraction per period. This means we need to:

- multiply by 0.01 to convert from percent to an absolute fraction; and,
- divide by the number of periods in a year to convert from years to periods.

PERIODS PER YEAR 5.1.9: *When a period or until of time is fixed in an interest problem, the corresponding number of periods per year will be denoted by m . In other words, one period is $(\frac{1}{m})^{\text{th}}$ of a year.*

For example, if periods are months, then $m = 12$; if they are quarters, then $m = 4$; and so on. We can summarize the conversion steps above with the equation:

INTEREST RATE CONVERSION FORMULA 5.1.10: $p = \frac{0.01 \cdot r}{m}$

There's one very important rule about using the **INTEREST RATE CONVERSION FORMULA 5.1.10** that is critical to getting the right—accurate to the nearest cent—answers in formulae that use it.

PERIODIC RATE RULE 5.1.11: *Never evaluate $\frac{0.01 \cdot r}{m}$ and plug this value in for p . Instead, always plug in the “raw” fraction $\frac{0.01 \cdot r}{m}$.*

Why? The reason is that m in the denominator. In the two most common cases, when the periods are months ($m = 12$) and days ($m = 365$), that m in the denominator means that the decimal value of p that your calculator returns is messy—the decimals go on as far as your calculator does.

All those messy decimals lead you into temptation: the temptation to break the **FIRST RULE OF ROUNDING 1.2.4** and to round p by writing down only *some* of those messy digits (in effect, rounding p). This

5.1 Simple interest

makes life easier for you two ways: fewer digits to copy down when recording the value of p and fewer to type into your calculator when plugging in the value of p in the final formula.

Unfortunately, there's very often a price to pay. The error you introduced by rounding p shows up in your final answer and it's off by a few cents or a few dollars. Experience teaches that it's very hard to decide in advance how many digits of p are "enough". And there's no way to tell from looking at the kind of messy final answers we'll be getting that those last few digits are wrong.

You may think that I am overdramatizing, but decades of teaching this material have shown me that rounding p is the single most common source of small errors in working financial math problems. Fortunately, the [PERIODIC RATE RULE 5.1.11](#) provides a simple and bulletproof solution. If you always plug in the formula $\frac{0.01 \cdot r}{m}$ for p , you'll never make any errors of this type.

What's more, it's actually *faster and easier* to just type in $\frac{0.01 \cdot r}{m}$ whenever you need to plug in p than it is to copy and retype the decimal value for p that it gives. So train yourself *now* to always follow this rule and you'll save yourself time and errors throughout this chapter.

In fact, the most important points to grasp in this introductory section are all of this type. You won't have much use for the [SIMPLE INTEREST FORMULA 5.1.6](#) later on. But if you train yourself to apply it as outlined below, you'll go a long way to being able to correctly apply all the more complex formulae that will come later. So pay careful attention to the [METHOD FOR FINDING SIMPLE INTEREST 5.1.15](#) and the [PERIODIC RATE RULE 5.1.11](#).

EXAMPLE 5.1.12: If the nominal interest rate is 6% a year and we measure periods in months, then $m = 12$ and the periodic interest rate is $p = \frac{0.01 \cdot r}{m} = \frac{0.01 \cdot 6}{12}$. This comes out to be 0.005 which is not such a messy decimal, but if you've absorbed the [PERIODIC RATE](#)

5.1 Simple interest

RULE 5.1.11 you'll never know this (and never need to know it) because you'll just plug in $\frac{0.01 \cdot 6}{12}$ when p is called for.

If we measure periods in days, then $m = 365$ and the periodic rate is $p = \frac{0.01 \cdot r}{m} = \frac{0.01 \cdot 6}{365}$. Now p comes out to be 0.0164383561643836 and you start to see why it's easier to stick with $\frac{0.01 \cdot 6}{365}$ and never have to worry about the decimal value.

Even if we measure periods in years so that $m = 1$, we *still* have to convert to get $p = \frac{0.01 \cdot r}{m} = \frac{0.01 \cdot 6}{1} = 0.06$.

Remember. The *first* step in almost *any* interest problem is to determine the units in which periods are measured and the corresponding m and to convert the nominal percent-per-year interest rate r in the problem to a fraction-per-period rate p using the **INTEREST RATE CONVERSION FORMULA 5.1.10**. This periodic rate p is the one you'll need in most formulae. You don't even need to evaluate $\frac{0.01 \cdot r}{m}$ —in fact, the **PERIODIC RATE RULE 5.1.11** tells you *not* to—but forgetting this simple conversion step is the *most common* source of *big* errors in interest calculations. So, don't forget it! Then remember the **PERIODIC RATE RULE 5.1.11** and plug in $\frac{0.01 \cdot r}{m}$ for p rather than precalculating it and you'll avoid the most common source of *small* errors too.

But wait: there's more. We need to use the *same* **period** or unit of time measure both the periodic rate p and the term T in the **SIMPLE INTEREST FORMULA 5.1.6**. In practice, the most common units are “months” (used for consumer loans like your credit card, mortgage, pension contributions, car loan) and “days” (used for bank accounts and most more commercial loans). But people being people, we like to speak about time in years if at all possible: after all it's a lot simpler to think about a 15 year mortgage than a 180 month mortgage, or a 5 year CD than a 1,826 day CD.

So, a conversion of time units is almost always needed. This almost always means converting a term stated in years, for which we'll use

5.1 Simple interest

the letter y , into a term T measured in whatever periods are the time units for the problem. The second step in working any interest problem is to make sure the term T in the problem is expressed in **periods**. Fortunately, this is very easy: if there are m periods in 1 year (this is the definition of m), then there will be my periods in y years, so

TERM CONVERSION FORMULA 5.1.13: $T = m \cdot y$

EXAMPLE 5.1.14: If term is given as 4 years and we measure periods in months, then $m = 12$ and the term in periods is $T = m \cdot y = 12 \cdot 4 = 48$. If we measure periods in days, then $m = 365$ and the term in periods is $T = m \cdot y = 365 \cdot 4 = 1460$.

What is the *second most common* source of errors in interest calculations? You guessed it: forgetting to convert from years to periods. So don't forget that either! We can formalize all this as a:

METHOD FOR FINDING SIMPLE INTEREST 5.1.15:

Step 1: Determine the periods in the problem (that is, the units in which the term is measured) and the value of m , the number of periods per year.

Step 2: Use the [INTEREST RATE CONVERSION FORMULA 5.1.10](#) to find the periodic interest rate p from the nominal interest rate r and the [TERM CONVERSION FORMULA 5.1.13](#) to find the term T in periods from the term in years y .

Step 3: Apply the [SIMPLE INTEREST FORMULA 5.1.6](#).

EXAMPLE 5.1.16: An amount of \$2,235.00 is loaned for a term of 3 years at a nominal rate of 9% a year. Find the simple interest due using *months* as periods.

Solution

Step 1: The periods are months so $m = 12$.

Step 2: $p = \frac{0.01 \cdot r}{m} = \frac{0.01 \cdot 9}{12}$ and $T = my = 12 \cdot 3 = 36$.

Step 3: $I = p \cdot A \cdot T = \left(\frac{0.01 \cdot 9}{12} \right) \cdot \$2,235.00 \cdot 36 = \$603.45$.



5.1 Simple interest

That's about it. When you think you have learned the basic ideas in this section try the Self Test. Once you can complete the self-test perfectly, you are ready to start on the problems.

Here are a few warm-up problems. You'll want your calculator. Remember the **FIRST RULE OF ROUNDING 1.2.4**: do *not* round until the very end of the calculation. I have provided a solution for one part of each to get you warmed up.

For each of the following loans, first find the periodic interest rate p and the term T in periods. Then determine the amount of interest in the loan using the **SIMPLE INTEREST FORMULA 5.1.6** to the nearest cent.

PROBLEM 5.1.17: In the following problem, measure the term T in periods of “years”.

- i) An amount of \$1,000.00 is loaned for a term of 2 years at a nominal rate of 6% a year.

Solution

Step 1: The periods are years so $m = 1$.

Step 2: $p = \frac{0.01 \cdot r}{m} = \frac{0.01 \cdot 6}{1}$ and $T = my = 1 \cdot 2 = 2$.

Step 3: $I = p \cdot A \cdot T = \left(\frac{0.01 \cdot 6}{1} \right) \cdot \$1,000.00 \cdot 2 = \$120.00$.

- ii) An amount of \$1,100,000.00 is loaned for a term of 5 years at a nominal rate of 4.9% a year.
- iii) An amount of \$16,235.00 is loaned for a term of four and a half years at a nominal rate of 8.22% a year.

PROBLEM 5.1.18: In the following problem, measure the term T in periods of “months”.

- i) An amount of \$1,000.00 is loaned for a term of 2 years at a nominal rate of 6% a year.
- ii) An amount of \$1,100,000.00 is loaned for a term of 5 years at a nominal rate of 4.9% a year.

Solution

Step 1: The periods are months so $m = 12$.



5.1 Simple interest

Step 2: $p = \frac{0.01 \cdot r}{m} = \frac{0.01 \cdot 4.9}{12}$ and $T = my = 12 \cdot 5 = 60$.

Step 3: $I = p \cdot A \cdot T = \left(\frac{0.01 \cdot 4.9}{12} \right) \cdot \$1,100,000.00 \cdot 60 = \$269,500.00$.

iii) An amount of \$16,235.00 is loaned for a term of four and a half years at a nominal rate of 8.22% a year.

EXAMPLE 5.1.19: Before we go on, let's see what could have happened in ii) above, if I'd ignored the **PERIODIC RATE RULE 5.1.11**. In this example $p = \left(\frac{0.01 \cdot 4.9}{12} \right) = 0.0040833333333333$. Who needs *all* those 3s, especially as the amount $A = \$1,100,000.00$ is such a round number?

Let's just write down $p = 0.004008333$. Nine decimals has *got* to be enough. Now I get

$$I = p \cdot A \cdot T = 0.004008333 \cdot \$1,100,000.00 \cdot 60 = \$269,499.97800.$$

which rounds to $I = \$269,499.98$. Close, but still off by 2¢.

It turns out in this problem, *ten* decimals were needed. If I had written down write down $p = 0.0040083333$, I'd have gotten \$269,499.997800 and, after rounding, \$269,500.00. But the only way to know that 10 is enough and 9 isn't, is by having the right answer to compare to. So the only smart plan is *not* to round and the easy way to do this is never to write p down, as in the solution above.

PROBLEM 5.1.20: In the following problem, measure the term T in periods of "quarters".

i) An amount of \$1,000.00 is loaned for a term of 2 years at a nominal rate of 6% a year.

ii) An amount of \$1,100,000.00 is loaned for a term of 5 years at a nominal rate of 4.9% a year.

iii) An amount of \$16,235.00 is loaned for a term of four and a half years at a nominal rate of 8.22% a year.

Solution

Step 1: The periods are quarters so $m = 4$.

Step 2: $p = \frac{0.01 \cdot r}{m} = \frac{0.01 \cdot 8.22}{4}$ and $T = my = 4 \cdot 4.5 = 22$.

5.1 Simple interest

Step 3: $I = p \cdot A \cdot T = \left(\frac{0.01 \cdot 8.22}{4}\right) \cdot \$16,235.00 \cdot 22 = \$7,339.8435 = \$7,339.84$ to the nearest cent.

PROBLEM 5.1.21: In the following problem, measure the term T in periods of “days”. Don’t worry about leap years.

i) An amount of \$1,000.00 is loaned for a term of 2 years at a nominal rate of 6% a year.

Solution

Step 1: The periods are days so $m = 365$.

Step 2: $p = \frac{0.01 \cdot r}{m} = \frac{0.01 \cdot 6}{365}$ and $T = my = 365 \cdot 2 = 730$.

Step 3: $I = p \cdot A \cdot T = \left(\frac{0.01 \cdot 6}{365}\right) \cdot \$1,000.00 \cdot 730 = \$119.999999 = \120.00 .

Here’s another example which illustrates that it’s actually easier *not* to write down p . In this example, $p \approx 0.000164383561643836$. Which is easier, $\frac{0.01 \cdot 6}{365}$ or 0.000164383561643836?

ii) An amount of \$1,100,000.00 is loaned a term of for 5 years at a nominal rate of 4.8% a year.

iii) An amount of \$16,235.00 is loaned for a term of four and a half years at a nominal rate of 8.22% a year.

If you made all the conversions and used the formula correctly, you should find that the answers to the corresponding parts of the three problems above are equal *to the penny!* So why did I make all that noise about converting from nominal to periodic rates and from years to periods? It really doesn’t seem to matter. The answer is that for calculations which only use the [SIMPLE INTEREST FORMULA 5.1.6](#) it doesn’t matter—I’ll show why in a moment—but the interest calculations which come up in real life almost always involve the [COMPOUND INTEREST FORMULA 5.2.4](#) which we’ll study in the next section. And, in compound interest calculations the units always *do* affect the answer. I think it’s just easier to learn how to make the conversions at the start before the formulas get complicated (which they will) and to get in the habit of *always* making the conversions so you never forget to do so.



5.1 Simple interest

Why *do* the answers above agree even though the intermediate calculations are very different? Answer: the m 's cancel! (But they won't in the future). For example, suppose we measure the term T in months. Then the formulas above give

$$p = \frac{0.01 \cdot r}{12} \quad \text{and} \quad T = 12 \cdot y$$

hence

$$I = pAT = \frac{0.01 \cdot r}{12} \cdot A \cdot (12y) = (0.01r)Ay.$$

Note that the final answer is what we'd get right away if we used years as periods. To really convince ourselves that the answer will not involve the period, we can just replace 12 months-per-year by a general number m of periods-per-year to get

$$p = \frac{0.01 \cdot r}{m} \quad \text{and} \quad T = m \cdot y$$

hence

$$I = pAT = \frac{0.01 \cdot r}{m} \cdot A \cdot (my) = (0.01r)Ay.$$

After cancelling the m 's, the final answers agree: Note, however, that even if we measure the term in years ($m = 1$), we need to convert rates from per-cent to fractions: although T equals y in this case, p equals not r but $0.01 \cdot r$.

Although we think of the [SIMPLE INTEREST FORMULA 5.1.6](#) as telling us the interest I given the amount A , periodic rate p and term T , we can, as usual, turn it around and solve for any one of these four quantities given the other three. Here are some problems to practice this. The [METHOD FOR FINDING SIMPLE INTEREST 5.1.15](#) still applies but with one small difference. If a problem asks us for an interest rate or term, the corresponding conversion in Step 2 of the method has to be postponed until *after* Step 3—before Step 3, we won't know the value to be converted—and we have to carry out a fourth step and convert in the *opposite* direction. The [SIMPLE INTEREST FORMULA 5.1.6](#) will give us the **periodic rate** or the term in periods and the conversion will give us the corresponding **nominal rate** or term in years. Once again I have provided a few sample solutions as models.



5.1 Simple interest

Solve each of problems 5-7 below three ways, measuring the term T in periods of years, months and days.

PROBLEM 5.1.22: Your three answers to each part of this problem should agree to the nearest penny.

- i) Find the amount which will earn interest of \$22.75 when loaned at a nominal rate of 5% for a term of 2 years.

Solution (using months for periods)

Step 1: The periods are months so $m = 12$.

Step 2: $p = \frac{0.01 \cdot r}{m} = \frac{0.01 \cdot 5}{12}$ and $T = my = 12 \cdot 2 = 24$.

Step 3: Since $I = p \cdot A \cdot T$,

$$A = \frac{I}{p \cdot T} = \frac{22.75}{\left(\frac{0.01 \cdot 5}{12}\right) \cdot 24} = \$227.50000004 = \$227.50$$

to the nearest cent.

- ii) Find the amount which will earn interest of \$121.00 when loaned at a nominal rate of 8% for a term of 3 years.

PROBLEM 5.1.23: Your three answers to each part of this problem should be different numbers of periods. However, if you convert these periods into years, the three numbers of years should be identical.

- i) Find the term (as a number of periods) over which an amount of \$1,200.00 will earn interest of \$128.00 when loaned at a nominal rate of 8%.

- ii) Find the term (as a the number of periods) over which an amount of \$500.00 will earn interest of \$146.00 when loaned at a nominal rate of 7.3%.

Solution (using days for periods)

Step 1: The periods are days so $m = 365$.

Step 2: $p = \frac{0.01 \cdot r}{m} = \frac{0.01 \cdot 7.3}{365}$. Since the answer we are seeking is the term T , we wait until we have found T to convert the term.

Step 3: Since $I = p \cdot A \cdot T$, $T = \frac{I}{p \cdot A} = \frac{146.00}{\left(\frac{0.01 \cdot 7.3}{365}\right) \cdot 500} = 1460$ days.



5.1 Simple interest

Step 4: Now we convert the term T in days back to years y . Since

$$T = m \cdot y, y = \frac{T}{m} = \frac{1460}{365} = 4 \text{ years.}$$

PROBLEM 5.1.24: Your three answers to each part of this problem should be different periodic interest rates. However, if you convert these periodic rates into nominal rates the three nominal rates you obtain should be identical.

i) Find the periodic interest rate at which which an amount of \$900.00 will earn interest of \$135.00 when loaned for a term of 2 years.

Solution (using quarters for periods)

Step 1: The periods are quarters so $m = 4$.

Step 2: Since the answer we are seeking is the periodic rate p , we wait until we have found p to convert to a nominal rate r .

But we do want to convert the term: $T = m \cdot y = 4 \cdot 2 = 8$.

Step 3: Since $I = p \cdot A \cdot T$, $p = \frac{I}{A \cdot T} = \frac{135.00}{900 \cdot 8} = 0.01875$.

Step 4: Now we convert the periodic rate p back to a nominal percent-per-year rate r . Since $p = \frac{0.01 \cdot r}{m}$, $r = 100 \cdot m \cdot p = 100 \cdot 4 \cdot 0.01875 = 7.5\%$.

Let's just check that we get the same nominal rate when we use years as periods. The calculation proceeds: $m = 1$, $T = m \cdot y = 1 \cdot 2 = 2$, $p = \frac{I}{A \cdot T} = \frac{135.00}{900 \cdot 2} = 0.075$ and hence we get a nominal rate $r = 100 \cdot m \cdot p = 100 \cdot 1 \cdot 0.75 = 7.5\%$ exactly as before.

ii) Find the periodic interest rate at which which an amount of \$850.00 will earn interest of \$170.00 when loaned for a term of 4 years.

Here are a few more problems to practice with. Solve each of problems 8-11 below three ways, measuring the term in periods of years, months and days and check that your three answers agree (converting periodic rates to nominal rates and periods to years if needed).

PROBLEM 5.1.25: Find the term over which an amount of \$600.00 will earn interest of \$64.80 when loaned at a nominal rate of 7.2%.

5.1 Simple interest

PROBLEM 5.1.26: Find the amount which will earn interest of \$70.00 when loaned at a nominal rate of 9% for a term of 1.5 years.

PROBLEM 5.1.27: Find the interest due when an amount of \$21,000.00 is loaned for a term of 15 years at a nominal rate of 3.22% a year.

PROBLEM 5.1.28: Find the periodic interest rate at which which an amount of \$1,125.00 will earn interest of \$46.00 when loaned for a term of 8 years.

Here is one that involves a bit more calculation but will prepare you for the ideas in the next section on compound interest.

PROBLEM 5.1.29:

i) A sum of \$1,000.00 is loaned at 6% interest using periods of years. Find the total (amount loaned plus interest) which will be due if the loan term is:

- a. 1 year.
- b. 2 years.
- c. 3 years.

ii) Suppose that, at the end of one year, the loan amount plus the interest due is paid and the total is loaned for a second year at 6% interest. How much will the second loan plus interest come to at the end of the second year? Your answer should be \$1,123.60 whereas the answer to part i)b is \$1,120.00 Can you explain the difference?

iii) If the total (including interest) at the end of the second year is loaned for a third year at 6% interest, how much will the total due at the end of the third year be? Not only is this answer different from that to i)c but the difference is bigger than the difference in ii). Can you explain why this is?

PROJECT 5.1.30: The **SIMPLE INTEREST FORMULA 5.1.6** is based on three principles: **equality of dollars**, **equality of periods**, and constancy of the **interest rate**. Are our three principles correct in the real world? As is so often the case, the answer is “Not quite!”.



5.1 Simple interest

- i) Look at the mortgage rates quoted in your local Sunday newspapers real estate section. Are the rates the same for all mortgage amounts? If not, how do they vary? What might the explanation be for such a variation?
- ii) Look at the interest rates paid for Treasury Bills or various maturities (a fancy word for the term during which the government gets to keep the money). Are the rates the same for all maturities? If not, how do they vary? What might the explanation be for such a variation?
- iii) We have already seen that there is no such thing as *the* interest rate: instead there are many different rates for many different kinds of loans. Make a list of a variety of different rates and discuss what factors might explain the differences. For example, why is a bank able to attract deposits when it pays only a few percent a year in interest while you have to pay rates close to 20% a year to borrow via your credit card?
- iv) Interest rates not only vary by the type of loan but change over time. How much have interest rates on bank deposits, or mortgages, or treasury bills varied over the last thirty years. How can we try to explain this variation? What is the role of the Federal Reserve Bank in influencing this process? You might find it helpful to speak to a friend who is majoring in economics or finance (or to a professor in your economics or finance department).

If you looked into the research project, you will see that our interest formulas are based on some imprecise assumptions. We assume that, in any given transaction, there's a single fixed interest rate that applies. While this is often true—examples are bank CDs and fixed rate mortgages—it's more often false. We'll just ignore the possibility that the interest rate charged on a transaction often changes over time. The reason is the usual one in math. If we do not ignore this, then all our interest calculations become much more complicated—so complicated that a substantial component of many jobs is keep-

5.2 Compound interest

ing track how interest rates change and taking suitable actions to respond to these changes. People who devote a lot of effort to this include: bankers, accountants, stockbrokers, actuaries, ...

5.2 Compound interest

Let's start by recalling the basic argument for interest on loans: if you loan me money, you receive interest to compensate you for losing various opportunities to make attractive purchases with the money. In this section, we refine this argument.

Let's imagine that I want to borrow an amount $A = \$100,000.00$ from the bank for a term of 30 years to purchase a house at a nominal interest rate of 8% and calculate the simple interest I would owe at the end of the thirty year term using years as periods. Here $p = \frac{0.01 \cdot r}{m} = \frac{0.01 \cdot 8}{1} = .08$ —I break my rule of never evaluating p in the [INTEREST RATE CONVERSION FORMULA 5.1.10](#) here to make the calculations easier to follow—and $T = y = 30$). We find $I = pAT = 0.08 \times \$100,000.00 \times 30 = \$240,000.00$. The interest outweighs the amount when I go to pay back.

Now let's ask what the situation would be if the bank asked me to pay them back after 15 years, but offered to give me a second 15 year loan once I did. We now have $T = y = 15$ so

$$I = pAT = 0.08 \times \$100,000.00 \times 15 = \$120,000.00$$

so I have to pay the amount of \$100,000.00 plus interest of \$120,000.00. Now, I take out a second \$100,000.00 loan and after the second 15 years pay it back with another \$120,000.00 in interest. At first it looks like nothing has changed: in both cases the bank has its original loan money back along with a total of \$240,000.00 in interest at the end of 30 years. However, there is one big difference. In the second scenario, the bank is holding an extra \$120,000.00 for the

5.2 Compound interest

second 15 years. This affords the bank many opportunities it would not have in the first scenario. In effect, if I can wait the full 30 years to pay, then for the second 15 years the bank is making me an interest free loan of \$120,00.00. Nice work if you can get it.

One thing the bank might be willing to let me do is to borrow the \$120,000.00 interest for the second 15 years but make me pay interest on the second loan too: I'd owe an extra

$$I = 0.08 \times \$120,000.00 \times 15 = \$144,000.00.$$

Instead of receiving its \$100,000.00 loan plus \$240,000.00 in interest at the end of the thirty years the bank now gets the \$100,000.00 loan plus the \$120,000.00 in interest on that loan plus the \$120,000.00 it loaned me back after 15 years *plus* \$144,000.00 in interest on the amount of the reloaned interest. By making me pay the interest I owe twice rather than just once, the bank comes out ahead \$144,000.00 at the end of the thirty years. This \$144,000.00 is the interest due in the *second* 15 years, not on the amount of the loan but on the *interest* due in the first 15. Such interest on the interest of a loan is called *compound interest*. More generally,

COMPOUND INTEREST 5.2.1: *An amount of money is at compound interest when, at regular intervals called compounding periods, the accrued interest is added to the amount current at the start of the period and this sum is used as the current amount for the next compounding period. A compounding period is no different from a period like those we used in the previous section. Unless we really want to emphasize that we are dealing with compound interest, we'll generally just speak of periods.*

I suspect you are already asking the same question the bank ought to. If it was smart for the bank to ask to get paid off twice wouldn't it be smarter still to ask to be paid off three times? The answer is yes: I'll let you fill in the details and just summarize the numbers this time. After 10 years, I'd owe \$80,000.00 in interest. If I borrowed

5.2 Compound interest

this back, then after 20 years, I'd owe another \$80,000.00 interest on the original loan plus the rolled over interest loan amount of \$80,000.00 plus \$64,000.00 in interest on the rollover loan for a total of \$224,000.00 in various interests. Finally, if I rollover this second amount of interest and borrow it back back for another 10 years, then at the very end, in addition to repaying the original loan of \$100,000.00 I'd owe another \$80,000.00 in interest on the original loan plus the \$224,000.00 amount of the second rollover loan plus \$179,200.00 in interest on this *second* rollover. The upshot is that in addition to repaying the bank its \$100,000.00 this time I'd owe \$483,200.00 so the bank would be up \$243,200.00 compared to the scenario in which I only paid interest once and \$99,200.00 compared to the scenario in which I paid interest twice.

Before we go any further, let's ask what happens if I borrow the money for a term of three years instead of thirty. I claim that if interest is collected once, I will owe \$24,000.00 in interest at the end of the three years; if it is collected twice, I will owe \$25,440.00; and, if it is collected three times, I will owe \$25,971.20. Before you read any further, test whether you have followed the discussion above by checking these numbers. If you have not worked [PROBLEM 5.1.29](#) which gets you warmed up for these calculations, you should do so first.

PROBLEM 5.2.2:

- i) Check the figures given above for a three year term by imitating the the calculations in the preceding paragraphs.
- ii) What are the corresponding numbers if I borrow the money for a term of 12 years?

Several ideas seems clear from the examples above. Qualitatively, we can draw two basic conclusions. First, the more often I have to pay off the outstanding interest on my loan the better off the bank is going to be. Put in more technical language the more frequently my loan is compounded, the more interest I will owe. Second, the

5.2 Compound interest

longer my loan is outstanding the greater the differences in interest owed caused by more frequent compounding, or put another way, the more important interest on interest becomes.

We can draw some quantitative conclusions too. First, the simple interest formula contains all the information we need to calculate exactly how much better off the bank will be in any of these cases *in principle*. What if the bank wants me to pay interest 4 times, or 5 times, or 30 times (once a year) or 360 times (once a month)? We could probably slog our way through all the calculations for 4 or even 5 payments. But, in principle be damned, there is no way we are going to extend these calculations to 30 or 360 payments. To make such calculations practical, we need a better way to think about compound interest. Fortunately, there is a very effective point of view.

To describe it, we begin with a simple question: If I borrow an amount A at a periodic interest rate p for a term T of 1 period how much will I owe at the end of the period *including accrued interest*? The answer is easy: the [SIMPLE INTEREST FORMULA 5.1.6](#) tells me that I owe interest of $I = pAT = pA$ since $T = 1$. To this I have to add the amount A which I borrowed so the total is

$$A + I = A + pA = A(1 + p).$$

In words,

ONE PERIOD PRINCIPLE 5.2.3: *The total amount I owe at the end of any period is the amount I owed at the beginning of the period times the magic factor $(1 + p)$ where p is the periodic interest rate. Alternatively, to get from the starting amount to the ending amount, you add the starting amount times the periodic rate.*

The key point about this statement is that we only have to multiply by the magic factor $(1 + p)$ *regardless of the amount of the loan or the units used to measure periods*. Every \$1 dollar owed at the start becomes $\$(1+p)$ owed at the end: this is a consequence of [EQUALITY OF DOLLARS 5.1.4](#).



5.2 Compound interest

Let's restate this yet another way. We'll write A_0 for the amount owed now: the "0" is called a subscript and is usually used, as we will use it here, to distinguish a series of related quantities, in this case related amounts. We'll also write A_1 for the outstanding balance (that is, amount owed with interest) after 1 period, A_2 for the outstanding balance at the end of 2 periods, A_3 for the outstanding balance at the end of 3 periods and so on with A_T standing for the outstanding balance at the end of a general number T of periods. The formula above can then be rewritten

$$A_1 = A_0 (1 + p).$$

But we immediately get many similar equations:

$$A_2 = A_1 (1 + p)$$

$$A_3 = A_2 (1 + p)$$

and more generally

$$A_T = A_{T-1} (1 + p).$$

In each case, the amounts on the left and right sides of each equality side are the totals owed exactly one period apart and the [ONE PERIOD PRINCIPLE 5.2.3](#) tells us these *always* differ by a factor of $(1 + p)$. But, things are even nicer: we can write all the amounts A_n very simply in terms of the original amount A_0 . I claim:

COMPOUND INTEREST FORMULA 5.2.4: $A_T = A_0(1 + p)^T$

Conceptually, this follows directly from the [ONE PERIOD PRINCIPLE 5.2.3](#). Each period the loan is outstanding causes the amount outstanding to get multiplied by a factor of $(1 + p)$. So, over a term of T periods the amount picks up T factors of $(1 + p)$ which is the same as a factor of $(1 + p)^T$. We can check this argument algebraically for

5.2 Compound interest

small values of the number T of periods.

$$A_1 = A_0(1 + p) \qquad \qquad \qquad = A_0(1 + p)^1$$

$$A_2 = A_1(1 + p) = (A_0(1 + p)^1)(1 + p) = A_0(1 + p)^2$$

$$A_3 = A_2(1 + p) = (A_0(1 + p)^2)(1 + p) = A_0(1 + p)^3$$

$$A_4 = A_3(1 + p) = (A_0(1 + p)^3)(1 + p) = A_0(1 + p)^4$$

and so on. The first equality on each line is one of those listed above, the second comes from substituting the equality of the previous line, and the third comes by combining the powers of $(1 + p)$.

Finally, let's check that the calculations we made above by hand also agree with this formula. This will also, help us to start coming to grips with a point that confuses many students. In the compound interest formula, the term T is given in units of the *compounding* period, the period of time between interest calculations (so, as we will see again below, T is *always* equal to the number of compoundings). In the previous section, you were forced to get used to using different units for T but the answer came out the same regardless of what units we used.

This is no longer true: we will *only* get the right answer in a compound interest calculation if we use compounding periods as our units of time. In other words, if we pay interest monthly, then we *must* express the term of the loan in months. If we pay compound annually, we *must* express the term of the loan in years. If we pay compound daily, we *must* express the term of the loan in days.

To start with, we consider the \$100,000.00 loan at a nominal rate of 8% with a thirty year term with which we began the section. In first calculation we made, interest was paid only once at the end of the thirty years. This means that the compounding period was 30 years, so that the quantity m which represents the number of compounding periods per year is a fraction!: $m = \frac{1}{30}$. (This problem will be one of the very few cases in which m is not a whole number.) Once we swallow this, however, everything becomes very simple. The initial



5.2 Compound interest

loan amount A_0 is \$100,00.000.00. The term $T = m \cdot y = \frac{1}{30} \cdot 30 = 1$ corresponding to the fact that 30 years is 1 compounding period and the periodic interest rate is

$$p = \frac{0.01 \cdot r}{m} = \left(\frac{0.01 \cdot 8}{\frac{1}{30}} \right) = 0.08 \cdot 30 = 2.4.$$

This problem will also be one of the very few cases in which p is not a small decimal and I have again broken my rule and evaluated the [INTEREST RATE CONVERSION FORMULA 5.1.10](#) to make it easier to see what's happening. So, the [COMPOUND INTEREST FORMULA 5.2.4](#) $A_T = A_0(1 + p)^T$ becomes $A_1 = \$100,000.00(1 + 2.4)^1 = \$340,000.00$ as hoped: this is the sum of the \$100,000.00 needed to repay the principal of the loan plus the \$240,000.00 in simple interest we computed above.

What happens when we compound twice? Now the compounding period is 15 years so $m = \frac{1}{15}$, the 30 year term $T = m \cdot y = \frac{1}{15} \cdot 30 = 2$, the periodic interest rate is $p = \frac{0.01 \cdot r}{m} = \frac{0.01 \cdot 8}{\frac{1}{15}} = 0.08 \cdot 15 = 1.2$, and the compound interest equation gives a final outstanding amount of

$$A_2 = \$100,000.00(1 + 1.2)^2 = \$100,000.00 \cdot 4.84 = \$484,000.00.$$

This is exactly \$144,000.00 more than the we got when we compounded once exactly as above. By now, I hope you are getting the hang of things. When we compound three times, the compounding period is 10 years, $m = \frac{1}{10}$, $T = 3$ and $p = 0.8$ and the compound interest equation gives a final outstanding amount of

$$A_3 = \$100,000.00(1 + 0.8)^3 = \$583,200.00$$

which is exactly \$99,200.00 more than the we got when we compounded twice and \$243,200.00 more than the we got when we compounded once. Once again, we recover exactly the hand calculation above.

We can formalize these examples with a method. It looks a lot like the [METHOD FOR FINDING SIMPLE INTEREST 5.1.15](#). The only differ-



5.2 Compound interest

ence is the third step where we use the COMPOUND INTEREST FORMULA 5.2.4 instead of the SIMPLE INTEREST FORMULA 5.1.6. We call this a provisional method because we'll only apply it to some easy problems. Later in this section, you'll find a final version which you'll be able to apply to a larger range of problems.

PROVISIONAL METHOD FOR FINDING COMPOUND INTEREST 5.2.5:

Step 1: Determine the periods in the problem (that is, the units in which the term is measured) and the value of m , the number of periods per year.

Step 2: Use the INTEREST RATE CONVERSION FORMULA 5.1.10 to find the periodic interest rate p from the nominal interest rate r and the TERM CONVERSION FORMULA 5.1.13 to find the term T in periods from the term in years y .

Step 3: Apply the COMPOUND INTEREST FORMULA 5.2.4.

EXAMPLE 5.2.6: Let's try redoing PROBLEM 5.2.2 using our new method. Since this problem tells us how many compoundings we are to make—that is, the number T of periods—we first need to determine the length of each compounding period in years and then, by inverting this, the number m of periods per year. We then use these to find the periodic rate and compounded amount. The tables below summarize the results.

First, we switch the term to three years. I have once again evaluated the periodic rates p in these examples to let you warm up.

T	y	m	p	A_T
1	3	$\frac{1}{3}$	0.24	$\$100,000.00(1 + 0.24)^1 = \$124,000.00$
2	2	$\frac{1}{2}$	0.16	$\$100,000.00(1 + 0.16)^2 = \$125,440.00$
3	1	$\frac{1}{1}$	0.08	$\$100,000.00(1 + 0.08)^3 = \$125,971.20$

TABLE 5.2.7: COMPOUNDED AMOUNTS FOR A 3 YEAR TERM

5.2 Compound interest

Finally, we switch the term to 12 years. Now, it's time to *start* listening to the **PERIODIC RATE RULE 5.1.11** and *stop* evaluating p .

T	y	m	p	A_T
1	12	$\frac{1}{12}$	$\left(\frac{0.01 \cdot 8}{\frac{1}{12}}\right)$	$\$100,000.00 \left(1 + \left(\frac{0.01 \cdot 8}{\frac{1}{12}}\right)\right)^1 = \$196,000.00$
2	6	$\frac{1}{6}$	$\left(\frac{0.01 \cdot 8}{\frac{1}{6}}\right)$	$\$100,000.00 \left(1 + \left(\frac{0.01 \cdot 8}{\frac{1}{6}}\right)\right)^2 = \$219,040.00$
3	4	$\frac{1}{4}$	$\left(\frac{0.01 \cdot 8}{\frac{1}{4}}\right)$	$\$100,000.00 \left(1 + \left(\frac{0.01 \cdot 8}{\frac{1}{4}}\right)\right)^3 = \$229,996.80$

TABLE 5.2.8: COMPOUNDED AMOUNTS FOR A 12 YEAR TERM

You might notice that the columns labelled T are identical in both tables. They'd better be: remember anytime you use the compound interest formula, you *must* measure the term T in compounding periods, so this is just another way of saying that T equals the number of compoundings.

Now you should be ready to try a few easy problems. While you are working these, try to get a feel for how the final amount A_T changes as we change the length of the compounding period. I have inserted a few solutions as models. As usual, we'd like to get amounts correct to the nearest penny so we want to plug in the "raw" **INTEREST RATE CONVERSION FORMULA 5.1.10** for p rather than evaluating it first.

PROBLEM 5.2.9: Find the final amount which will accrue if an amount of \$13,253.44 earns nominal interest of 4% for a term of 7 years if:

- i) interest is compounded annually.

Solution

Step 1: The periods are years so $m = 1$.

Step 2: $p = \frac{0.01 \cdot r}{m} = \frac{0.01 \cdot 4}{1}$ and $T = my = 1 \cdot 7 = 7$.

Step 3: $A_T = A_0(1+p)^T = \$13,253.44 \left(1 + \left(\frac{0.01 \cdot 4}{1}\right)\right)^7 = \$17,440.62$.

- ii) interest is compounded quarterly.
- iii) interest is compounded monthly.

5.2 Compound interest

PROBLEM 5.2.10: Find the final amount which will accrue if an amount of \$2,600.00 earns nominal interest of 9% for a term of 3 years if:

- i) interest is compounded annually.
- ii) interest is compounded quarterly.

Solution

Step 1: The periods are quarters so $m = 4$.

Step 2: $p = \frac{0.01 \cdot r}{m} = \frac{0.01 \cdot 9}{4}$ and $T = my = 4 \cdot 3 = 12$.

Step 3: $A_T = A_0(1 + p)^T = \$2,600.00 \left(1 + \left(\frac{0.01 \cdot 9}{4}\right)\right)^{12} = \$3,395.73$.

- iii) interest is compounded monthly.

PROBLEM 5.2.11: Find the final amount which will accrue if an amount of \$1,255,000.00 earns nominal interest of 6.73% for a term of 5 years if:

- i) interest is compounded annually.
- ii) interest is compounded quarterly.
- iii) interest is compounded monthly.

Solution

Step 1: The periods are months so $m = 12$.

Step 2: $p = \frac{0.01 \cdot r}{m} = \frac{0.01 \cdot 6.73}{12}$ and $T = my = 12 \cdot 5 = 60$.

Step 3: Using $A_T = A_0(1 + p)^T = \$1,255,000.00 \left(1 + \left(\frac{0.01 \cdot 6.73}{12}\right)\right)^{60} = \$1,755,397.74$.

So far all the examples and problems have had one common feature: we know how much we want to borrow (or invest) and would like to find out how much we will owe (or collect) *after* a certain term at interest. In other words, the amount A_T that we are seeking only exists at a time which is in the *future* relative to the time at which the amount A_0 which we know exists. For this reason, problems like those we have been working are often called *future value* problems.

FUTURE VALUE 5.2.12: *A future value problem is one in which an amount or value is sought which lives in the future relative to a known amount or value.*

5.2 Compound interest

However, there are many situations in which there is an amount which we'd like to be able to collect in the future and what we'd like to figure out is the amount we'd need to set aside at interest *today* to do so. For example, on my daughter's eighth birthday, it might occur to me that I'll need to have \$120,000.00 or so available for her college education in 10 years (perhaps, half of what sending her to a good private college will cost) and I might like to know how much I'd need to put away now in order to have \$120,000.00 on hand 10 years from today. Or, I might want to plan ahead in similar ways for retirement or other goals. Business often face a similar need to make provisions today for anticipated future expenses. Since the amount we are trying to determine usually lives in the *present* such problems are generally called *present value* problems.

PRESENT VALUE 5.2.13: *A present value problem is one an amount or value is sought which lives in the past relative to a known amount or value.*

If you compare this definition to the previous one, you'll see that the terminology is not very good. Unfortunately, it is completely standard. So let me emphasize that the important factor is *not* whether the known sum lives in the past present or future and it is *not* whether the unknown sum lives in the past, present or future. What's important is whether the unknown sum lives after the known sum (in the future *relative to* it) in which case we have a future value problem or whether it lives before the known sum (in the past *relative to* it) in which case we have a present value problem.

Perhaps you're also wondering, "What's the difference between a value and an amount?" None. They both stand for a quantity of money at a point in time. The word "value" has many other meanings in mathematics so I'd prefer to avoid it. The word "amount" is the natural one to introduce in connection with interest, as we did [SECTION 5.1](#). And I'd like to use the letter *A* for all amounts wherever I can. Then, if you see an *A*, you know it's an amount and to find

5.2 Compound interest

the amount, you just look for the A . However, the terms present and future *value* are completely standard in dealing with compound interest. As a practical matter, you need to learn to recognize and work with them too. The solution I've adopted is a compromise. I'll try to keep an A in the notation for all amounts, but I'm going to feel free to use the word value instead of amount at times. In particular, *future value* and *present value* are such standard terms that you'll almost never see terms like "future amount" or "present amount problem".

EXAMPLE 5.2.14: We can use the **COMPOUND INTEREST FORMULA** 5.2.4 $A_T = A_0(1+p)^T$ to solve present value problems too. Just divide both sides by $(1+p)^T$ to get the formula

$$A_0 = \frac{A_T}{(1+p)^T}$$

which allows us to find the present value A_0 given the future value A_T (and of course p and T which we'd figure out as usual from r , y and m). Suppose, in the example of my daughter's college fund that my bank will sell 10 year CDs (certificates of deposit with a 10 year term) which earn 3.6% interest compounded monthly. What I want to know is how much money I need to put into these CD's now to have $A_T = \$120,000.00$ worth of them in 10 years. Fine: the bank compounds monthly so $m = 12$. Thus $p = \frac{0.01 \cdot r}{m} = \frac{0.01 \cdot 3.6}{12}$ and $T = m \cdot y = 12 \cdot 10 = 120$. Finally,

$$A_0 = \frac{A_T}{(1+p)^T} = \frac{\$120,000.00}{\left(1 + \left(\frac{0.01 \cdot 3.6}{12}\right)\right)^{120}} = \$83,766.29.$$

Gleep! I don't have that kind of money. I should have started planning for this sooner. What if I had invested when my daughter was born? All that changes is that now $y = 18$ so $T = 216$ and when I calculate A_0 , I now get \$62,831.83. I could never have afforded that much eight years ago. Yow! What should I do? We'll find out the answer to this question in the section on **SECTION 5.6**.

For now, let's just draw a few simple conclusions. First, the principle that a sum of money at one point in time changes in value

5.2 Compound interest

when we try to move it to a second point in time—interest is earned to compensate for the loss of use of the money—continues to apply. However, there is now one significant difference. Amounts which we know about today become larger when we move them to the future because interest is paid on them. Turning this around, amounts which we know about in the future (like that \$120,000.00 college fund) become *smaller* when we move them to the present (or past) because the interest they could have been earning must be deducted. Second, solving present value problems involves the same ideas and formulas we use in solving future value problems. With the approach in the previous paragraph, the only difference is whether we multiply or divide by $(1 + p)^T$.

However, I'd like things to be even simpler. I'd like to *always* use A_0 to denote the sum of money which is *known* and I'd like to *always* use A_T to denote the sum of money which is to be *determined*. Moreover, I'd like to *always* be able to find A_T by *multiplying* A_0 by $(1 + p)^T$: I just don't want to worry about whether to multiply or divide. I *always* want to be able to use the using the **COMPOUND INTEREST FORMULA 5.2.4** as it stands: $A_T = A_0(1 + p)^T$.

"Yeah, well if pigs had wings, they could fly. Get real, Dr. Morrison." Not so fast: if we are willing to stretch our brains just a bit, we can have everything I ask for. The key idea is simply to allow the number of periods T to be *either positive or negative*. Just think of T as telling us the number of periods between the moment in time when the amount A_0 whose value we know lives and the moment in time when the amount A_T whose value we are seeking lives. A *positive* T means that we have to move *forward* in time—towards the *future*—to get from the amount A_0 to the amount A_T . A *negative* T means that we have to move *backward* in time—towards the *past*—to get from the amount A_0 to the amount A_T . In all the *future* value problems we have worked to this point, the the amount we knew lived *before* the amount we were seeking, we moved *forward* in time, and so T was

5.2 Compound interest

positive.

But, in the example of my daughter's college fund, the amount I know I'll need—the \$120,000.00—lives in the future when she turns 18. The amount I want to determine—the “unknown” \$83,766.29 I'd need to invest—lives in the present: to get to this time I have to travel 10 years backwards in time. I now want to think of this as (-10) years. Correspondingly, I have to consider a term of $T = m \cdot y = 12 \cdot (-10) = -120$ monthly periods. A bit strange, but once I do so, I can just apply the [COMPOUND INTEREST FORMULA 5.2.4](#) unchanged using as A_0 the *known* amount \$120,000.00 and $p = \left(\frac{0.01 \cdot 3.6}{12}\right)$ as above to find the “unknown” amount A_T . I get

$$A_T = A_0(1+p)^T = \$120,000.00 \left(1 + \left(\frac{0.01 \cdot 3.6}{12}\right)\right)^{-120} = \$83,766.29$$

exactly as before.

This problem is pretty typical. Even though we are mentally moving backwards in time towards the past, we never actually enter the past. We stop at the present because that is when the amount we are seeking—the \$83,766.29 I'd need to invest today—lives. That's why problems where we move towards the past and use a negative number of periods T are called *present* rather than past value problems. But, things work just as well if I really do travel into the past. We can check this by thinking about the investment I might have made when my daughter was born. Now, the known amount $A_0 = \$120,000.00$ lives 10 years in the future when my daughter turns 18 and the unknown amount lives 8 years in the past when she was born. To get between them, I have to travel backwards 18 years in time so $y = -18$ and $T = -216$. When I compute the “unknown” amount this time, I find that $A_T = A_0(1+p)^T = \$120,000.00 \left(1 + \left(\frac{0.01 \cdot 3.6}{12}\right)\right)^{-216} = \$62,831.83$ exactly as before.

If you remember your [RULES OF EXPONENTS 1.4.10](#), you've probably already figured out why both approaches give the same answer. When I did the problem the first time, I *divided* the known amount



5.2 Compound interest

of \$120,000.00 by $\left(1 + \left(\frac{0.01 \cdot 3.6}{12}\right)\right)^{120}$. When I did it the second time, I wound up multiplying the \$120,000.00 by $\left(1 + \left(\frac{0.01 \cdot 3.6}{12}\right)\right)^{-120}$. But, by [PROBLEM 1.4.57](#),

$$\left(1 + \left(\frac{0.01 \cdot 3.6}{12}\right)\right)^{-120} = \frac{1}{\left(1 + \left(\frac{0.01 \cdot 3.6}{12}\right)\right)^{120}}$$

so multiplying by $\left(1 + \left(\frac{0.01 \cdot 3.6}{12}\right)\right)^{-120}$ has exactly the same effect as dividing by $\left(1 + \left(\frac{0.01 \cdot 3.6}{12}\right)\right)^{120}$ and the answer comes out the same. More generally, dividing by $(1 + p)^T$ and multiplying by $(1 + p)^{-T}$ are the same.

Many of you are probably saying to yourselves: “He may think it’s simpler to multiply by a power with a negative exponent and than to divide by one with a positive exponent, but I sure don’t. If you get the same answer both ways, why do I need to worry about positive or negative numbers of periods?” Fair question. Most books try to avoid the negative periods and negative exponents. There are two good reasons why it’s better to swallow and use them.

First, doing so allows us to solve all compound interest problems—present or future value—with a *single formula* and a *single method*. If there is one key to the power of mathematics, it’s the ability to simplify by **seeing patterns** and generalizing. It is much easier and less error prone to understand one idea well enough to be able to use it to solve several different kinds of problems than to learn several ideas each of which can only be applied to one particular kind of problem. To take advantage of this power, you must continually strive to apply methods you know to new problems rather than seeking new methods for every different problem. Here is a concrete case where you can try to do this. Rather than having to learn two formulas—one for future value problems and one for present value problems—and then to learn how to apply each and which problems each applies to, we will make the [COMPOUND INTEREST FORMULA 5.2.4](#) do *all* the work.

5.2 Compound interest

The second reason for using negative exponents here, where we do not *absolutely* have to, is like the reason for using periodic interest rates in [SECTION 5.1](#). By always converting terms to periods and rates to periodic rates even when it was still optional, we avoided having to train ourselves to do this in compound interest calculations when it these conversions *are* obligatory and learned to avoid the most common error in interest problems.

We'll get a similar benefit from training ourselves *now* to ask whether terms and periods are to be counted positively or negatively—that is, whether we have to find present value problem or a future value—even though we could avoid the negative exponents in this section. That's because, when we reach the sections which discuss amortization (things like mortgages and retirement savings), there *will* be formulae where negative terms and negative exponents are *unavoidable*.

So I want you to train yourselves now to ask the question, “Does the unknown sum of money live *after* the known sum or does it live *before* the known sum?” and to use a *positive* term if the answer is *after* (future value) and a *negative* term if the answer is *before* (present value).

The payoff will come when we work with amortizations because once again we won't need to learn any new concepts—just a new formula. We'll be able to classify all our amortization problems by asking exactly the question above. Finally, can you guess what the most common error in working amortization problems is? Right, it's getting positive and negative terms mixed up! Once again, we can train ourselves to avoid these errors by learning to use negative exponents in this section.

So, let's get to it. First, let's formalize the examples above into a method we can use on all the remaining problems in this section. The only difference from the [PROVISIONAL METHOD FOR FINDING](#)



5.2 Compound interest

COMPOUND INTEREST 5.2.5 will be that we'll need to determine the sign of the term T .

METHOD FOR FINDING COMPOUND INTEREST 5.2.15:

Step 1: Determine the periods in the problem (that is, the units in which the term is measured) and the value of m , the number of periods per year.

Step 2: Use the **INTEREST RATE CONVERSION FORMULA 5.1.10** to find the periodic interest rate p from the nominal interest rate r .

Step 3: Subtract the point in time at which the known amount in the problem lives *from* the point in time at which the amount unknown in the problem lives to obtain the term as a *signed* number of years y or other periods: the sign is *positive* if the known amount lives *before* the unknown amount and *negative* if the known amount lives *after* the unknown amount. If the *signed* term in you obtain is in years y , use the **TERM CONVERSION FORMULA 5.1.13** to convert to a *signed* term T in *periods*.

Step 4: Apply the **COMPOUND INTEREST FORMULA 5.2.4**

Experience shows that lots of students have a hard time with signed periods and that you can't repeat the basic idea enough, so ... If the known amount lives *before* the unknown amount then the difference in step 3 gives a positive number of years and the unknown amount will be larger than the known amount. If the known amount lives *after* the unknown amount, then the difference gives a negative number of years and the unknown amount will be smaller than the known amount.

Here are some present value problems so you can get a feel for working with negative terms.

PROBLEM 5.2.16: Find the present value of a amount of \$4,200.00 which will be due in 4 years if the nominal interest rate is 12% and:

i) interest is compounded annually.



5.2 Compound interest

Solution

Step 1: The periods are years so $m = 1$.

Step 2: The known amount of \$4,200.00 lives 4 years from now; the unknown present value lives today. The difference (time of unknown amount) subtract (time of known amount) equals (today) subtract (today +4) years or (-4) years.

Step 3: $p = \frac{0.01 \cdot r}{m} = \frac{0.01 \cdot 12}{1}$ and $T = my = 1 \cdot (-4) = -4$.

Step 4: $A_T = A_0(1 + p)^T = \$4,200.00 \left(1 + \left(\frac{0.01 \cdot 12}{1}\right)\right)^{-4} = \$2,669.18$.

Two comments. Usually, we'll be more concise in step 2 and just say the known amount lives -4 years from the known amount. Also, note that, as in any present value problem, the present value is smaller than the known amount.

ii) interest is compounded quarterly.

iii) interest is compounded monthly.

PROBLEM 5.2.17: Find the amount you'd have to put in the bank at age 25 at an interest rate of 4.8% to have \$100,000.00 when you retire at age 65 if:

i) interest is compounded annually.

ii) interest is compounded quarterly.

iii) interest is compounded monthly.

Solution

Step 1: The periods are months so $m = 12$.

Step 2: The known amount of \$100,000.00 lives $65 - 25 = 40$ years *after* the unknown amount so the term is -40 years.

Step 3: $p = \frac{0.01 \cdot r}{m} = \frac{0.01 \cdot 4.8}{12}$ and $T = my = 12 \cdot (-40) = -480$.

Step 4: $A_T = A_0(1 + p)^T = \$100,000.00 \left(1 + \left(\frac{0.01 \cdot 4.8}{12}\right)\right)^{-480} = \$14,716.95$.

PROBLEM 5.2.18: Find the amount you'd have to put in the bank at age 45 at an interest rate of 7.2% to have \$100,000.00 when you retire at age 65 if:

i) interest is compounded annually.

ii) interest is compounded quarterly.



5.2 Compound interest

Solution

Step 1: The periods are quarters so $m = 4$.

Step 2: The known amount of \$100,000.00 lives $65 - 45 = 20$ years *after* the unknown amount so the term is -20 years.

Step 3: $p = \frac{0.01 \cdot r}{m} = \frac{0.01 \cdot 7.2}{4}$ and $T = my = 4 \cdot (-20) = -80$.

Step 4: $A_T = A_0(1+p)^T = \$100,000.00 \left(1 + \left(\frac{0.01 \cdot 7.2}{4}\right)\right)^{-80} = \$23,998.15$.

iii) interest is compounded monthly.

You may find one point about your answers to these exercises a bit puzzling. When we were computing future values—sending money forward in time—it grew in value. Moreover, if we kept the nominal rate and term fixed, then the more frequently we compounded the greater the growth. In these present value problems—where we send money backward in time—it shrinks in value. But, if we keep the nominal rate and term fixed as in the problems above, we find that the more frequently we compound the greater the shrinkage. Aren't these two conclusions about how compounding more frequently affects money travelling forward and backward in time contradictory? The answer is no, but it can be a bit confusing. The way to remove the confusion is to forget which amount is known and which is unknown for a moment.

Let's just think about the present and future values. It'll help to name them. I'll use B for the present value—the amount at the *start* of the term—and S for the future value—the amount at the *end* of the term. (Why the apparently random letters B and S ? In later sections, we'll want to discuss buying and selling investments. The buying price B will be the present value of the investment and the selling price S will be its future value. We'll need similar notations when we discuss savings and loan amortizations, where B will usually be a loan Balance and S a Savings goal. It's easier to just use B all present values and S for all future values than to introduce new letters in each new situation.)

Now let's ask, "How does the frequency of compounding affect the



5.2 Compound interest

ratio $\frac{S}{B}$ of the future value S to the present value B ?” We can restate the fact that more frequent compounding makes future values like S *larger* by saying that more frequent compounding makes the ratio *bigger*. But, we can *also* restate the fact that more frequent compounding makes present values like B *smaller* by saying that more frequent compounding makes this the *bigger*. The conclusions are the identical. This is worth noting for future reference.

EFFECT OF MORE FREQUENT COMPOUNDING 5.2.19: *If the nominal interest rate and term are fixed, then the more frequently interest is compounded, the greater the ratio $\frac{S}{B}$ of the future value S to the present value B . If the nominal interest rate, term and present value B are fixed, then the more frequently interest is compounded, the greater the future value S . If the nominal interest rate, term and future value S are fixed, then the more frequently interest is compounded, the smaller the present value B .*

Now, some exercises in which present and future values problems are intermixed and you must *think* about the sign of the term. I have worked the first one for you.

PROBLEM 5.2.20: A U.S. bond which earns 5.2% interest compounded quarterly will be redeemable for \$1,000.00 in 15 years. What does it sell for?

Solution

Step 1: The periods are quarters so $m = 4$.

Step 2: The known amount of \$1,000.00 lives 15 years *after* the unknown amount—what the bond sells for today—so the term is -15 years.

Step 3: $p = \frac{0.01 \cdot 5.2}{12}$ and $T = my = 12 \cdot (-15) = -300$.

Step 4: $A_T = A_0(1 + p)^T = \$1,000.00 \left(1 + \frac{0.01 \cdot 5.2}{12}\right)^{-300} = \459.18 .

PROBLEM 5.2.21: A money market account containing \$3,200.00 earns 3% interest compounded monthly. What will its balance be in 5 years?



5.2 Compound interest

PROBLEM 5.2.22: I put \$32,000.00 in my son’s college fund when he was born. If the fund pays 6% interest compounded monthly, how much money will he have for college when he is 18?

PROBLEM 5.2.23: Five years ago, the University trustees put money in an account which earns 4% a year interest compounded daily to be used to retire a loan of \$250,000.00 next year. How much did they set aside?

PROBLEM 5.2.24: Loan sharks typically charge interest of 2% a week, compounded weekly. This is one of the few cases where interest is stated in terms of the compounding period and the periodic rate (i.e., here $p = 0.02$ —we’re given the *periodic* rate rather than, as usual, the nominal one—and $m = 52$). What would a debt of \$1,000.00 amount to in a year at these rates if no payments were made to the loan shark? What would the debt amount to in 5 years?

PROJECT 5.2.25: In 1636, Peter Stuyvesant bought the island of Manhattan for \$24.00. Was this good deal? In other words, if Peter had deposited the \$24.00 at interest for the intervening 370-some years, would his heirs have enough in their account to buy Manhattan today? To answer this question, you will need to find out two types of information. First, about how much would the land in Manhattan be worth today? (Just ignore the cost of buildings roads etc.) Second, what kind of interest could the Stuyvesants have expected to earn over the past four centuries? One thing you’ll find out is that interest rates have varied considerably over this period. The [COMPOUND INTEREST FORMULA 5.2.4](#) only deals with accruals at a fixed rate of interest. So another question you’ll need to answer is: how should we to deal with the variations in interest rates in a problem like this? The interest rate on common accounts—like money market accounts—also changes frequently. Ask people in your business school how financial institutions handle the complications this introduces. You’ll see how much *simpler* math is than everyday life!



5.3 Approximating compound interest

The [SIMPLE INTEREST FORMULA 5.1.6](#) is just that: simple. As long as you remember [PERIODIC RATE RULE 5.1.11](#) and you just multiply and it is pretty hard to go wrong. What's more there is an easy way to check any answer: since changing the periods used to measure the term has no effect on the final value of the interest, you can just recalculate with different periods. The [COMPOUND INTEREST FORMULA 5.2.4](#) is not really more complicated. However, since the choice of the compounding period now has a substantial effect on the answer, you can't just blindly recalculate with different periods to check.

In this section, we'll learn how to keep our eyes open while changing periods. This leads to easy approximations for compound interest which can be used to check that a compound interest answer at least "looks right". While doing this, we'll find out that calculators don't always compute the compound interest formula correctly and learn how to work around their limitations. In later sections ([SECTION 5.4](#) and [SECTION 5.5](#), for example), what we'll learn will turn out to be handy in many other ways.

The simple interest approximation

The first approximation we'll use is truly simple: just ignore compounding and use the [SIMPLE INTEREST FORMULA 5.1.6](#). If we use years as periods—so $p = \frac{0.01 \cdot r}{m} = 0.01r$ and $T = y$ —then the interest on an amount A_0 is $I = p \cdot A_0 \cdot T = 0.01r \cdot A_0 \cdot y$ and the future value A_T —the total of principal plus interest at the end of the period—is *roughly* $A_0 + 0.01r \cdot A_0 \cdot y = A_0(1 + 0.01r \cdot y)$. We can use this to approximate the compound interest formula. Here's an example.

5.3 Approximating compound interest

EXAMPLE 5.3.1: Let's try checking my answer to one of the Problems in SECTION 5.2. In part ii) of PROBLEM 5.2.10 which asked for the future value S of \$2,600.00 at nominal interest of 9% for a term of 3 years compounded quarterly, my answer was \$3,395.73. The SIMPLE INTEREST APPROXIMATION 5.3.3 gives $A_0(1 + 0.01r \cdot y) = \$2,600(1 + 0.09 \cdot 3) = \$3,302.00$ which is just a bit smaller than my answer.

If we think of A_0 as an initial amount or present value B and think of A_T as a final amount or future value S , we can rewrite the approximation

$$A_T \simeq A_0(1 + 0.01r \cdot y) \quad \text{as} \quad S \simeq B(1 + 0.01r \cdot |y|)$$

This new formula on the right can be used to check both future and present value problems. Just remember that S is to be the final amount and B the starting amount, so in a present value problem we'd have to reverse the A 's: $S = A_0$ and $B = A_T$. The bars around the y are absolute value signs and remind us to *always* make y positive in this formula—just as we have always done when working with simple interest.

EXAMPLE 5.3.2: In part i) of PROBLEM 5.2.16 which asked for the present value of a sum of $S = \$4,200.00$ which will be due in 4 years at a nominal interest rate of 12% compounded annually, I got $B = \$2,669.18$. The SIMPLE INTEREST APPROXIMATION 5.3.3 gives $B(1 + 0.01r \cdot |y|) = \$2,669.18(1 + 0.12 \cdot 4) = \$3,950.39$. There are two points to note. First, this time we plugged in the *answer* as the present value or starting amount—because we were checking a present value problem. Second, even though I used a negative period of -4 years in PROBLEM 5.2.16 I plugged in $|y| = 4$ in the check.

Unfortunately, this “approximation” is only close to the correct value when the term is quite short. If you look over the example at the start of section SECTION 5.2, you'll see that over terms of many years, the interest on the interest in a compound interest calculation can become much larger than the original amount or the interest on the

5.3 Approximating compound interest

original amount. However, we can still make some use of this formula by thinking of it somewhat differently. Using the [SIMPLE INTEREST FORMULA 5.1.6](#) amounts to making the entire term a single period so that, in effect, there is no compounding. In other words we want $T = 1$. We can check this idea by working backwards. Since $T = m \cdot y$ by the [TERM CONVERSION FORMULA 5.1.13](#), $m \cdot y = 1$ so $m = \frac{1}{y}$. Then, by the [INTEREST RATE CONVERSION FORMULA 5.1.10](#)

$$p = \frac{0.01 \cdot r}{m} = \frac{0.01r}{\frac{1}{y}} = 0.01r \cdot y.$$

Finally, applying the [COMPOUND INTEREST FORMULA 5.2.4](#)

$$A_T = A_0(1 + p)^T = A_0(1 + 0.01r \cdot y)^1 = A_0(1 + 0.01r \cdot y).$$

By combining this with what we know about the [EFFECT OF MORE FREQUENT COMPOUNDING 5.2.19](#), we can squeeze out a bit of information even when the approximation is way off. Remember that more frequent compounding increases future values and decrease present values. Since using simple interest amount to doing the least compounding possible—none at all—it should definitely underestimate future values. Note that both examples above confirm this: the approximations are both slightly smaller than the answers being checked.

SIMPLE INTEREST APPROXIMATION 5.3.3: *The ending or future amount S in a compound interest calculation should be greater than the approximation $B(1 + 0.01r \cdot |y|)$ (where B is the starting or present value) and the approximation should be fairly good if the term is not too long.*

Both [EXAMPLE 5.3.2](#) and [EXAMPLE 5.3.2](#) illustrate reasonable uses of the [SIMPLE INTEREST APPROXIMATION 5.3.3](#). I often do something even cruder in my head when working problems in class. In the first case, I'll say: "The interest is $0.09 \cdot 3 = 0.27$ which is about $\frac{1}{4}$ and $\frac{1}{4}$ of \$2,600.00 is \$650.00 so I should expect my answer to be a bit bigger than \$3,250.00". Similarly, in the second problem, I'd say, "Here

5.3 Approximating compound interest

the interest is $0.12 \cdot 4 = 0.48$ which is about $\frac{1}{2}$ so I expect \$2,669.18 times $(1 + \frac{1}{2}) = \frac{3}{2}$ to be a bit smaller than \$4,200.00; it's easier to turn this around $\frac{2}{3}$ of \$4,200.00 or \$2,800.00 should be somewhat bigger than \$2,669.18". The point is that since we are only approximating anyway we can afford to ignore the difference between 0.27 and $\frac{1}{4}$ or between 0.48 and $\frac{1}{2}$. I know I am much more likely to perform a quick mental check than one where I have to get out my calculator and a check you don't perform is not much of a check.

EXAMPLE 5.3.4: To see the limitations of the approximation, let's look at part iii) of [PROBLEM 5.2.17](#) which asked for the present value of an amount of \$100,000.00.00 due in 40 years at an interest rate of 4.8% compounded monthly. My answer was \$14,716.95. The [SIMPLE INTEREST APPROXIMATION 5.3.3](#) gives $B(1 + 0.01r \cdot y) = \$14,716.95(1 + 0.048 \cdot 40) = \$42,973.49$. Again this is less than the correct future value of \$100,000.00, .00 but now it is so much less—barely two-fifths—that is not much use as a check. The period here was just too long for the approximation to be useful.

PROBLEM 5.3.5: Use the [SIMPLE INTEREST APPROXIMATION 5.3.3](#) to check your answers to [PROBLEM 5.2.11](#), [PROBLEM 5.2.21](#) and [PROBLEM 5.2.23](#). Are there any other problems in [SECTION 5.2](#), for which you'd expect the [SIMPLE INTEREST APPROXIMATION 5.3.3](#) to be fairly accurate?

The continuous approximation

We have already seen in [EFFECT OF MORE FREQUENT COMPOUNDING 5.2.19](#) that keeping the nominal rate and the term in years fixed, the more often we compound the larger the amount owed at the end of the term. Let's go back to the [PROBLEM 5.2.2](#) where we borrowed \$100,000.00.00 at 8% interest and ask: What happens if the bank compounds ever more frequently? Let's try compounding monthly,



5.3 Approximating compound interest

daily, hourly and once a second in the two problems above. The corresponding values of m are : 12, 365, 8,760 and 31,536,000. (In other words, there are 8,760 hours and 31,536,000 seconds in a year. So while a million seconds might seem like forever it is actually only about 11.57 days.) Now we can just apply the [METHOD FOR FINDING COMPOUND INTEREST 5.2.15](#).

I have not given the details of the calculations, but it will be important to note that I carried them out on a calculator with 10 place accuracy. The results are shown in [TABLE 5.3.6](#) below for a term of 3 years:

m	12	365	8,760	31,536,000
T	36	1,095	26,280	94,608,000
p	$0.08 \cdot 12$	$\frac{0.01 \cdot 8}{365}$	$\frac{0.01 \cdot 8}{8760}$	$\frac{0.01 \cdot 8}{31536000}$
A_T	\$127,023.71	\$127,121.56	\$127,123.37	\$132,819.91

TABLE 5.3.6: COMPOUNDING TABLE FOR A 3 YEAR TERM

and [TABLE 5.3.7](#) for a term of 12 years.

m	12	365	8,760	31,536,000
T	144	4,380	105,120	378,432,000
p	$\frac{0.01 \cdot 8}{12}$	$\frac{0.01 \cdot 8}{365}$	$\frac{0.01 \cdot 8}{8760}$	$\frac{0.01 \cdot 8}{31536000}$
A_T	\$260,338.92	\$261,142.08	\$261,156.97	\$311,209.46

TABLE 5.3.7: COMPOUNDING TABLE FOR A 12 YEAR TERM

PROBLEM 5.3.8: Make your own calculation of each of the amounts in the tables above. You should get the answers in the table to the penny when you compound monthly. Some of your other answers may be somewhat different for reasons I'll explain in a moment. If

5.3 Approximating compound interest

so, don't worry.

There are many interesting things to note about these tables. Let's look at the p rows to start. First notice that these rows in the two tables contain identical values: we should expect this since p does not depend on the *term* of the loan—which is what differs between the two tables—but only on the *frequency* with which we compound.

Second, notice how tiny the periodic rates have become: this is because we are compounding many times a year with very short compounding periods and hence we get very little interest in each period. The values are so small that if you try to evaluate p , your calculator will use scientific notation to show you as many decimals as possible. For example, in my 10 place calculator, $\frac{0.01 \cdot 8}{31536000} = 2.536783358\text{e-}9$ which we can also write $2.536783358 \times 10^{-9}$ or 0.000000002536783358. My calculator just does not have room for all those 0's.

Next, let's look at the final amount of A_T rows. Does anything strike you about these? All the amounts are getting closer and closer together—in the three year table they seem to be settling down around \$127,120.00 or so and in the 12 year table around \$261,150.00—and then suddenly the last amount where we compound in seconds is much bigger.

What's going on? Two things. First, the final amounts are *wrong*! The problem is that to find them I asked my calculator to compute

$$A_T = \$100,000.00 \cdot \left(1 + \left(\frac{0.01 \cdot 8}{31536000}\right)\right)^{94,608,000}$$

and

$$A_T = \$100,000.00 \cdot \left(1 + \left(\frac{0.01 \cdot 8}{31536000}\right)\right)^{378,432,000}$$

it choked. It just can't compute that value accurately. The 10 digits of precision it uses is just not enough to get the final amount even to 2 places!! If I use a much better calculator (which carries 20 places) and make the same calculation, I get the amounts in [TABLE 5.3.9](#).



5.3 Approximating compound interest

A_T , 3 years	\$127,023.71	\$127,121.56	\$127,124.78	\$127,124.92
A_T , 12 years	\$260,338.92	\$261,142.18	\$261,168.50	\$261,169.65

TABLE 5.3.9: AMOUNTS USING A 20 PLACE CALCULATOR

Notice that it turns out it wasn't just the final “seconds” amounts that were wrong. They were just the only wrong answers that were so far off that it was clear to the naked eye that something was fishy. Both the amounts for hourly compounding were off by dollars and the 12 year final amount with daily compounding—which is what your bank uses—was off by 10 cents. But all these other answers were close enough that the only way we would ever know that they were wrong was by making a second more accurate calculation. The moral here is:

MURPHY’S LAW OF CALCULATORS 5.3.10: *Never trust a calculator’s answers unless you have some other way to check them.*

What went wrong? Looking at the formulae for A_T above, you’d at first guess that the calculator has trouble handling those enormous exponents. That’s partly right. But, the main source of the error comes inside the parentheses when we *add the 1 to p !!* Let’s write out what’s involved:

$$\left(1 + \left(\frac{0.01 \cdot 8}{31536000}\right)\right) = 1 + 2.536783358 \times 10^{-9} = 1.00000000253678335.$$

So far so good. But my calculator only works to 10 place accuracy so this number has to get rounded to 1.000000003. That’s the number that get raised to that huge power and this number looks like we’d used a 1-digit value of 0.000000003 for p ! The effect is to raise the periodic rate by 20%— $1.2 \cdot 2.5 = 3$ —so no wonder the final answer is a whole lot bigger.

You may be wondering how you are going to do problems like the one where we got the wrong answer if you are not able to use your calculator. Relax: I’m just not going to ask *you* to work any problem

5.3 Approximating compound interest

where the exponent in the [COMPOUND INTEREST FORMULA 5.2.4](#) is dangerously large. But, in real life where daily compounding is common, you might be asked to: if you are, remember to watch out.

Did some of the amounts you computed in [PROBLEM 5.3.8](#) differed from mine as I suggested in the problem that they might? Perhaps now you can guess why I told you not to worry about this. I knew my answers were wrong and (unless you have a very good calculator) that yours probably would be wrong too. But there's no chance we'd get the *same* wrong answers. Why? Because every calculator is a bit different inside. While they'll all give the same answer when they can get the right one, when things go wrong each calculator goes wrong in its own way. If, like the largest number of students, you have a TI-8x calculator and you want to check your work, you should have obtained the correct answers in all but the last column.

So much for the bad news. Let's get back to those amounts. The correct answers give a striking confirmation of our initial impression that as you compound more and more frequently, the final amounts get larger and larger but do so ever more slowly. Eventually, these amounts appear to settle down. In fact, no matter how often you compound—even if you compound a trillion times a second—the final amount you'll wind up with in these two problems will never grow by another *cent*: after 3 years, you'll have \$127,124.92 and after 12 you'll have \$261,169.65. (You need a calculator that keeps 25 places to check these answers so you'll just have to trust me on this. If this worries you a bit, take a gold star: you're catching on.)

What would be very nice is to have some “easy way” to get this magic maximum amount. It's not that we'd ever want to compound interest every second in real life and so have a direct need for a way around the limitations of our calculators. But, if we did we'd have a good way to make a rough check of any compound interest calculation. Our answer should be close to, but somewhat smaller than this magic amount: the more frequently we compound, the closer the two



5.3 Approximating compound interest

should be. In particular when we compound monthly or daily as in the majority of real world loans we should see a few matching digits, as in the tables above. If we do not, we'll know right away that something is wrong.

Well kids, life is good! Finding formulas for limits—which is fancy way to say, for how things “settle down”—is one of the main applications of calculus. And, a standard formula from calculus computes the magic amount which appears in the tables above. Moreover, you don't need to know any calculus to understand and use this formula.

CONTINUOUS APPROXIMATION 5.3.11: *The quantity $(1 + p)^T$ lies between 1 and the $e^{(p \cdot T)} = e^{(0.01r \cdot y)}$ and is close to the latter. Therefore, A_T lies between A_0 and the continuous approximation*

$$A_0 \cdot e^{(p \cdot T)} = A_0 \cdot e^{(0.01r \cdot y)}$$

and, if the compounding period is short, A_T is close to the continuous approximation.

THE NUMBER e 5.3.12: *The base e in this formula is a very important number: $e \simeq 2.71828182845904523536$. But you don't need to try to memorize any of these decimals: e is so important it's got its very own key on your calculator. Moreover, exponentials with base e occur so often in so many places that there is also a key usually called **exp** for taking the exponential base e of the current value.*

So to use the **CONTINUOUS APPROXIMATION 5.3.11**, you just calculate the product $0.01r \cdot y$, hit 2nd LN or exp on your calculator, and multiply by the amount A and you've got the magic amount. In the two problems above, it gives

$$A \cdot e^{(0.01r \cdot y)} = \$100,000.00 \cdot e^{(0.018 \cdot 3)} = \$127,124.92$$

and

$$A \cdot e^{(0.01r \cdot y)} = \$100,000.00 \cdot e^{(0.018 \cdot 12)} = \$261,169.65.$$

Try it with your calculator! You might wonder why I preferred to calculate the exponent in the form $0.01r \cdot y$ rather than in the apparently simpler form $p \cdot T$. First, let's remark that they really are equal.



5.3 Approximating compound interest

Using the [INTEREST RATE CONVERSION FORMULA 5.1.10](#) and [TERM CONVERSION FORMULA 5.1.13](#) gives $p \cdot T = \frac{0.01 \cdot r}{m} \cdot my = 0.01r \cdot y$. Forgetting to use one of these conversion formulas is the most common error in working interest problems. Thus, the fact that I can use the nominal interest rate and the term in years in the formula with no need for converting to periodic rates and periods, makes the continuous approximation perfect for catching such errors. It's one of the nicest features of the formula that it lets us work with the real life quantities we like to think in terms of—nominal rates and years.

But the formula has other amazing properties. First, even with a standard calculator you can use it to find the magic number to the penny. Using the [COMPOUND INTEREST FORMULA 5.2.4](#), I needed a supercalculator to get these numbers. You'd have no way to compute them. Makes you think there might be something to this [calculus](#) after all, and that it might not be as hard as it's cracked up to be. (Both guesses are correct and I hope this section will inspire a few of you to take a calculus course. If you are planning to do serious work in any of the mathematical, computational, physical, biological or social sciences, you will have to do so eventually and the sooner you start the easier a time you'll have with the math and the further ahead you'll be in your major. If you don't believe me, ask your major's undergraduate advisor.)

The final remarkable feature of the [CONTINUOUS APPROXIMATION 5.3.11](#) is that the signs of the quantities T and y which measure time are not mentioned anywhere. So far we have only used the formula in future value problems where these quantities are positive but everything works just as well in present value problems when they are negative. The future value version says that when T and y are positive that $A_0 < A_T < A_0 e^{(0.01r \cdot y)}$ or, since $A_T = A_0 (1 + p)^T$ that

$$A_0 < A_0 (1 + p)^T < A_0 e^{(0.01r \cdot y)}.$$

The present value version says that if we replace T by $-T$ and y by

5.3 Approximating compound interest

—y we should have

$$A_0 > A_0 (1 + p)^{-T} > A_0 e^{(0.01r \cdot (-y))}.$$

Note that, although the directions of the inequalities are reversed, the exact value A_T is still in the between A_0 and the continuous approximation as claimed. I leave this to you: it is good practice in playing with exponents and inequalities: you need no special knowledge about the number e .

PROBLEM 5.3.13: Show that if $A_0 < A_0 (1 + p)^T < A_0 e^{(0.01r \cdot y)}$, then $A_0 > A_0 (1 + p)^{-T} > A_0 e^{(0.01r \cdot (-y))}$.

If you are hoping that I will now explain where this approximation comes from, bless you. First off, there's not much doubt that the [CONTINUOUS APPROXIMATION 5.3.11](#) is correct. The fact that it computes the two amounts above to the penny is pretty convincing. And, as we've already noted you don't need to understand where it comes from to use it. So if you could care less, you can skip to [EXAMPLE 5.3.16](#).

Everything comes down to [BERNOULLI'S LIMIT FOR exp 1.4.56](#). As I remarked when we were studying exponentials, Bernoulli stumbled on his limit when he was trying answer exactly the same question about compound interest that we've been asking in this section: How much interest can be earned if we compound ever more frequently? So the connection of e with the mathematics of finance is not only an intimate, but a very early one. I'll write down the formula again here since we'll use it so much.

BERNOULLI'S LIMIT FOR exp 5.3.14: *If n is a large positive number, then $(1 + \frac{1}{n})^n$ is slightly smaller than e . The bigger n you take, the closer the approximation. In fancier notation, as $n \rightarrow \infty$, of $(1 + \frac{1}{n})^n \rightarrow e$.*

Remember that this formula depends on the logarithmic version [BERNOULLI'S LIMIT FOR exp 1.4.56](#) which was what we actually



5.3 Approximating compound interest

proved. But you can convince yourself that it's right, without going back there, by calculating with a few big values of n . For example,

$$\left(1 + \frac{1}{1000}\right)^{1000} = 2.716923932$$

and

$$\left(1 + \frac{1}{10000}\right)^{10000} = 2.718145927.$$

The first less than e by about 0.013 and the second by about 0.0014.

PROBLEM 5.3.15: You shouldn't take n too big, however. Why? What will happen if you do?

Moreover, it's easy to see how it [BERNOULLI'S LIMIT FOR exp 5.3.14](#) leads to the continuous approximation. If we take the $(p \cdot T)^{\text{th}}$ power of both sides, we find that

$$\left(\left(1 + \frac{1}{N}\right)^N\right)^{(p \cdot T)} = \left(1 + \frac{1}{N}\right)^{(N \cdot p \cdot T)}$$

is a bit smaller than $e^{(p \cdot T)}$.

The [CONTINUOUS APPROXIMATION 5.3.11](#) just says we can always find a value of N for which the exponential $\left(1 + \frac{1}{N}\right)^{(N \cdot p \cdot T)}$ equals the exponential $(1 + p)^T$ in the [COMPOUND INTEREST FORMULA 5.2.4](#). Equating bases we need $\frac{1}{N} = p$ and equating exponents we need $N \cdot p \cdot T = T$: there are two equations for the single unknown N which would usually be impossible to satisfy. However, here is where a small miracle happens. Solving the first equation tells that we must have $N = \frac{1}{p}$. If we plug this into the left side of the second equation it becomes $N \cdot p \cdot T = \frac{1}{p} \cdot p \cdot T = T$ so the exponents automatically match up too! The final point to note is that to get a good approximation we need to have N large. But

$$N = \frac{1}{p} = \frac{1}{\frac{0.01 \cdot r}{m}} = \frac{m}{0.01r}.$$

Thus N is big when m is: in other words, we get a good approximation when we are compounding frequently.

Any way, using this approximation is a cinch. Let's use it to check a few problems from the last section.



5.4 Yields

EXAMPLE 5.3.16: In [PROBLEM 5.2.11](#), we worked out future values of an amount of \$1,255,000.00 earning nominal interest of 6.73% for a term of 5 years. In part [iii](#)), where we compounded monthly I got a future value of \$1,755,397.74. The [CONTINUOUS APPROXIMATION 5.3.11](#) gives $\$1,255,000.00e^{(0.016.73 \cdot 5)} = \$1,757,048.78$ which is, as predicted, slightly higher but matches the exact answer to 3 places.

EXAMPLE 5.3.17: In part [ii](#)) of [PROBLEM 5.2.18](#), we computed the present value of \$100,000.00.00 due in 20 years at interest of 9.6% compounded quarterly to be \$14,996.97. The [CONTINUOUS APPROXIMATION 5.3.11](#) gives $\$100,000.00.00e^{(0.019.6 \cdot (-20))} = \$14,660.70$. Note that since this was a present value problem where we were moving money backwards in time, we used a negative value $y = -20$ and that this time, as expected, the continuous approximation is slightly lower than the exact answer. Because we are compounding less frequently here, we get a less accurate approximation—only 2 places match. But, the approximation is good enough that we'd be sure to catch any silly errors like forgetting to convert from rate or term or miskeying one of the numbers into our calculator.

PROBLEM 5.3.18: Use the continuous approximation to check your answers to the 4 problems [PROBLEM 5.2.20](#) to [PROBLEM 5.2.23](#).

5.4 Yields

If you're very alert, you may have noticed that while [SECTION 5.1](#) ended with an assortment of problems which asked you to find terms and rates as well as interest and amounts, all the problems in [SECTION 5.2](#) asked about amounts. Don't we ever want to use the compound interest formula $A_T = A_0 (1 + p)^T$ to find an interest rate r by solving for p and using the [INTEREST RATE CONVERSION FORMULA 5.1.10](#) to find or a term in years y by solving for T and using



5.4 Yields

the **TERM CONVERSION FORMULA 5.1.13**? Yes and no. When money is borrowed or loaned at fixed interest in everyday life, all the variables which affect the calculation of interest—amount, rate, term, compounding frequency—are almost invariably known by both parties. For all consumer loans, the lender is obliged to state these in writing to the borrower. This is the no side of the answer.

The yes side of the answer involves, paradoxically, evaluating financial transactions which *do not involve fixed interest rates*. There are many situations in which we give up the use of a sum of money at one point in time and at a later time receive the use of a different sum. As common simple examples, you might buy a house, live in it for 10 years and sell it, or you might buy a stock keep it for 15 months and sell it. As a more complicated example, you might put money into a pension plan through regular contributions throughout your working life and then receive pension checks from the plan after you retire.

INVESTMENT 5.4.1: *We'll use **investment** as a general term for any transaction in which a sum of money is surrendered at one point in time (we speak of making or buying the investment) and a second sum of money is received in return at a later point in time (we speak of liquidating or selling the investment).*

How can we compare investments? If I really knew the answer to that question, I wouldn't be teaching math. The reason the question is so hard is that when we make an investment we know how much money we are surrendering but we can only guess how much we will realize when we sell the investment. However, to the extent that we know or can guess a selling price, we can compare investments using the ideas of compound interest. In particular, we can compare the *past* performance of investments since in this case we know both the before and after prices.

So let's suppose that we are given a buying price B (the present value of the investment), a selling price S (the future value) and the term

5.4 Yields

y in years between the date of purchase and the date of sale. What we'd like to do is to use the [COMPOUND INTEREST FORMULA 5.2.4](#) $A_T = A_0 (1 + p)^T$ to come up with an interest rate r which is generally called the *yield* of the investment. When we invest money at fixed interest, we would rather get more interest than less—at least if we are willing to ignore other factors like the risk of losing our investment, taxes ...). So, with the same provisos, we can compare any two investments by simply computing their yields and seeing which is greater.

The first steps in computing the yield are clear: we set A_0 equal to the buying price B and A_T equal to the selling price S . However, we hit a snag when we try to evaluate the term T . It is measured in periods and to convert to these from the term in years y we need to know the compounding frequency m . On the one hand, we now know that changing m can substantially affect our answer so the choice matters. On the other, no actual compounding was going on so there is no “right” choice for m . This dilemma is usually resolved by picking m somewhat arbitrarily. After all, if our main interest is in coming up with a *fair* way of comparing investments, then all that matters is that we use the *same* m for all the investments we compare. Changing m will affect each yield individually but won't affect which investments have higher and lower yields. Thus there is not one yield but *many*.

YIELD 5.4.2: *A yield on an investment is a nominal interest rate r which would have allowed an amount equal to the buying price B of the investment to accrue to the selling price S over the term of the investment assuming a specified compounding frequency or equivalently assuming a specified compounding period.*

Annualized yields

By far the most common choice for m is to take $m = 1$: in other words, to “compound” annually.

ANNUALIZED YIELD 5.4.3: *A yield computed by using annual compounding is called **annualized yield** and is denoted by r_a .*

In an annualized yield r_a , r reminds us that we have an annual rate of some sort and the subscript a reminds us that this rate is a yield rather than a nominal rate. The difference between nominal rates and yields is mainly one of viewpoint. Both play the same roles in formulas. A nominal rate is one we know at the start of the investment while a yield is a rate we have to compute. Annualized yields are almost the only ones seen in everyday life. In fact, if you see a “yield” in an ad or financial report, you can pretty much assume that it is an annualized yield. The compounding frequency will never be mentioned.

Using annualized yields—taking $m = 1$ —has two big advantages. It greatly simplifies conversion: $T = 1 \cdot y = y$ by the [TERM CONVERSION FORMULA 5.1.13](#) and $p = \frac{0.01 \cdot r_a}{1} = 0.01r_a$ by the [INTEREST RATE CONVERSION FORMULA 5.1.10](#). The nominal interest rate r in this last formula is just the annualized yield r_a that we are looking for. Setting $A_0 = B$ and $A_T = S$, the [COMPOUND INTEREST FORMULA 5.2.4](#) $A_T = A_0 (1 + p)^T$ becomes

$$S = B (1 + 0.01r_a)^y$$

and the question is how can we solve for r_a . The first step is clear: divide both sides by B to get

$$\frac{S}{B} = (1 + 0.01r_a)^y .$$

Next we have to eliminate the y^{th} power so we can get at the yield r_a we are after. As always when solving equations, we have to ask what operation will *undo* a y^{th} power? There are two ways to state

5.4 Yields

the answer. In terms of roots, a y^{th} root undoes a y^{th} power. Taking the y^{th} root of both sides, we get

$$\sqrt[y]{\frac{S}{B}} = \sqrt[y]{(1 + 0.01r_a)^y} = 1 + 0.01r_a.$$

Now we just subtract 1 from both sides and multiply by 100 to get

$$r_a = 100 \cdot \left(\sqrt[y]{\frac{S}{B}} - 1 \right).$$

The y^{th} root operation that undoes a y^{th} power is the same as a power operation with fractional exponent $\left(\frac{1}{y}\right)$. Using this notation we can make the same calculation:

$$\left(\frac{S}{B}\right)^{\frac{1}{y}} = \left((1 + 0.01r_a)^y\right)^{\frac{1}{y}} = 1 + 0.01r_a$$

and hence

$$r_a = 100 \cdot \left(\left(\frac{S}{B}\right)^{\frac{1}{y}} - 1 \right).$$

ANNUALIZED YIELD FORMULA 5.4.4: *If an investment is bought for \$B and sold after y years for \$S, then the annualized yield r_a on the investment is*

$$r_a = 100 \cdot \left(\sqrt[y]{\frac{S}{B}} - 1 \right) = 100 \cdot \left(\left(\frac{S}{B}\right)^{\frac{1}{y}} - 1 \right).$$

For emphasis, I repeat that the two forms are completely equivalent: I only give them both because some calculators have a $\sqrt[y]{}$ key and some have a power key (usually marked x^y). Use whichever is more convenient. I will use roots henceforth.

EXAMPLE 5.4.5: Here's an example of how this formula is used. Suppose you are offered the choice of two investments. The first costs \$5,000 and returns \$10,000 in 10 years. The second costs \$4,000 and returns \$5,000 in 3 years. We ask: which is better investment? "No contest", says the guy selling the first investment. "My way your

5.4 Yields

double you money in 10 years. With hers you only make 25% in 3 years. It'd take you 12 years to double your money.” To check, we simply compute both annualized yields and see which is higher. We find that the two yields are

$$100 \cdot \left(\sqrt[10]{\frac{10000}{5000}} - 1 \right) = 7.18\%$$

and

$$100 \cdot \left(\sqrt[3]{\frac{5000}{4000}} - 1 \right) = 7.72\%.$$

So it is no contest: the second investment has a yield more than half a percent *higher*! What was wrong with the first salesman's argument? He ignored the fact that after the first three year term you continue to make 25% every three years but it's 25% of the new larger sum. We can check this by hand: after 3 years you have $1.25 \cdot \$4,000 = \$5,000$; after 6, $1.25 \cdot \$5,000 = \$6,250$ —note that extra \$250; after 9, $1.25 \cdot \$6,250 = \$7,812.50$. You have almost doubled your money in 9 years, you'll more than double it in 10.

PROBLEM 5.4.6: Show that \$4,000 at 7.72% interest compounded annually for 10 years amounts to \$8,414.41.

PROBLEM 5.4.7: How much will a bank CD which costs \$1,000 today and earns 4.5% interest compounded daily be worth at the end of 5 years?

PROBLEM 5.4.8: Which has the highest and which the lowest yield of the following three investments: a Treasury bill which costs \$24,850 today and returns \$100,000 in 30 years; a CD which costs \$1,000 today and earns 4.5% interest compounded daily over the next 5 years; and a house which costs \$120,000 12 years ago and is worth \$205,000 today? Warning: the annualized yield on the CD is *not* 4.5% as you might at first think! To find out what it is, you'll need to use the answer to the preceding problem.



5.4 Yields

What's going on with that CD? If the nominal interest rate r is 4.5%, why is the annualized yield r_a *not* 4.5%? We have already seen the reason. When you compound more frequently, interest piles up faster. You earn more at interest of 4.5% compounded *daily* than you do at the same *nominal* rate of v compounded *annually*. The annualized yield on the CD tells what higher rate—you should have found that it's 4.60%—we'd have to earn *compounded annually* to wind up with the same amount of interest as we do at a nominal rate of 4.5% *compounded daily*.

We can use the annualized yield to compare any two fixed interest propositions—even if the compounding frequencies are different and neither is annual. In this special case, the annualized yield r_a is generally called an *effective yield* or an *effective interest rate*: we'll use r_a for effective rates to emphasize that they are just particular annualized yields. The effective rate is the actual interest we earn in a year with the compounding frequency already factored in: it's what we'd like to compare. The nominal rate gets its name from the fact that it while it *names* a rate it does not really say what we'll earn until we also specify a compounding frequency.

It's straightforward to get a formula for computing the effective rate r_a from the nominal rate r and the compounding frequency m . We invest an initial amount $B = A_0 = \$1$ for one year and see what amount $S = A_T$ we have at the end of the year. The [TERM CONVERSION FORMULA 5.1.13](#) says that if $y = 1$ then $T = m \cdot y = m$ and the [INTEREST RATE CONVERSION FORMULA 5.1.10](#) says that $p = \frac{0.01 \cdot r}{m}$ so

$$S = A_T = (1 + p)^T = \left(1 + \frac{0.01 \cdot r}{m}\right)^m.$$

Then we just compute an annualized yield with $y = 1$. In fact, since $y = 1$ there is no need to take the y^{th} root and since $B = 1$ too, the formula for r_a simplifies to $r_a = 100 \cdot (S - 1)$. Plugging in for S , we find,

EFFECTIVE YIELD FORMULA 5.4.9: *The effective yield or effective rate r_a corresponding to a nominal interest rate r compounded m*



5.4 Yields

times a year is the interest we'd actually earn in one year and is given by the formula

$$r_a = 100 \cdot \left(\left(1 + \frac{0.01 \cdot r}{m} \right)^m - 1 \right).$$

EXAMPLE 5.4.10: Let's check that the effective rate for that CD as computed by the formula really is just the annualized yield of 4.60% we found in [EXAMPLE 5.4.5](#). Plugging in $r = 4.5\%$ and $m = 365$ we find that

$$r_a = 100 \cdot \left(\left(1 + \frac{0.01 \cdot 4.5}{365} \right)^{365} - 1 \right) = 100(1.046025084 - 1) = 4.60\%.$$

PROBLEM 5.4.11: Which offers a higher yield? A CD with a nominal rate of 5.25% compounded quarterly or one with a nominal rate of 5.20% compounded daily?

PROBLEM 5.4.12: Which is a better investment? A CD with a nominal rate of 4.45% compounded daily or a bond which sells for \$8,000 today and returns \$10,000 in 5 years.

PROBLEM 5.4.13: In 1972, I bought a bottle of 1967 Chateau d'Yquem for \$17.35. In 1997, a bottle of the same wine sold at auction for \$625.00. What yield would I have realized had I auctioned my bottle instead of drinking it with friends for my wife's 50th birthday?

PROBLEM 5.4.14: Loan sharks charge compound interest of 2% a week compounded weekly: see [PROBLEM 5.2.24](#). What is their annualized yield?

The last problem illustrates an important point. You do not actually need to liquidate an investment to calculate its yield: it's enough to know what price you *could* get if you *did* want to sell. This point is more commonly applied in evaluating investments in stock, bonds, mutual funds and other financial instruments for which a public market with published prices exists.

5.4 Yields

PROJECT 5.4.15: I collect wine as a hobby. You might be interested in classic cars, Van Goghs, old editions of the Bible, baseball cards, Choose some good which you are interested in and for which you can find historical price data and investigate what the yields have been on investments in this good over the past century (or as far back as you can find data). How much have yields varied from time to time? Look up the terms “risk” and “liquidity”: these are aspects of evaluating an investment which do not simply involve its price. Discuss the risk of investing in your chosen good and its liquidity.

Continuous yields

Recall that at the beginning of the section on yields we settled on comparing investments on the basis of annualized compounding with $m = 1$. While the annualized yield which this produces is pretty much the only one used in daily life, there is a second choice for m which is used in more theoretical applications: $m = \infty$! Of course, we cannot really use $m = \infty$ in any of our formulas. What we mean by this is that we try to understand what would happen to the yield if we took larger and larger finite values of m . We already know the answer: when we do this the [COMPOUND INTEREST FORMULA 5.2.4](#) gets replaced by the [CONTINUOUS APPROXIMATION 5.3.11](#).

CONTINUOUS YIELD 5.4.16: *A yield calculated from the [CONTINUOUS APPROXIMATION 5.3.11](#) is called a continuous yield and is denoted by r_c .*

The formulas for continuous yields r_c are actually *simpler* than those for annualized yields r_a . Moreover, they lead to a classic rule-of-thumb for estimating compound interest.

To get a formula for continuous yield, we begin as above by equating our buying price B with the the initial amount A_0 in the [CONTINUOUS APPROXIMATION 5.3.11](#) and our selling price S with the final amount



A_T so that

$$S = A_T = A_0 \cdot e^{(0.01r \cdot y)} = B \cdot e^{(0.01r \cdot y)}.$$

Then we just plug in the number of years y the investment was held and solve for the rate r . The main novelty comes, after dividing both sides of the yield equation by B when we try to solve

$$\frac{S}{B} = e^{(0.01r_c \cdot y)}$$

for r_c . Instead of hiding out inside the base of an exponential, the way r_a did above, the continuous yield r_c is in the exponent itself. Thus taking a root or power won't help: it just moves the r_c to a different exponent. For example if I raise both sides to the power $(0.01r_c \cdot y)$, I get

$$\left(\frac{S}{B}\right)^{\left(\frac{1}{\frac{r_c}{100} \cdot y}\right)} = \left(e^{(0.01r_c \cdot y)}\right)^{\left(\frac{1}{(0.01r_c \cdot y)}\right)} = e.$$

What we need is an operation which moves exponents down onto the main level of our expressions. This is exactly what a *logarithm* does.

The key property of logarithms we'll need is that $\log_b(x^a) = a \log_b(x)$. There are, however, many logarithm functions—one for each positive base so we need to choose a base before proceeding. Basically, any choice will do the job but, as we saw in [SECTION 1.4](#) choosing b to be the number e has many advantages. As I noted there, almost everyone who works with logarithms does *all* their work in this base because it makes many formulas simpler—including our formula for continuous yield. That's why we think of \ln as God's logarithm and call it the *natural* logarithm. Recall from [exp AND ln ARE INVERSES 1.4.53](#) that the key property then becomes: $\ln(e^a) = \log_e(e^a) = a \log_e(e) = a$.

Once again, recall that your calculator has a special \ln key for calculating natural logarithms (which incidentally confirms how widely used \ln is). Using it, we can solve for the **continuous yield**. We just take the natural logarithms of both sides of

$$\frac{S}{B} = e^{(0.01r_c \cdot y)}$$

5.4 Yields

and apply the key property $\ln(e^a) = a$ with $a = (0.01r_c \cdot y)$ to find

$$\ln\left(\frac{S}{B}\right) = \ln\left(e^{(0.01r_c \cdot y)}\right) = (0.01r_c \cdot y).$$

Finally, we multiply both sides by $\left(\frac{100}{y}\right)$ to get

CONTINUOUS YIELD FORMULA 5.4.17: *If an investment is bought for \$B and sold after y years for \$S, then the continuous yield r_c on the investment is*

$$r_c = \frac{100}{y} \cdot \ln\left(\frac{S}{B}\right).$$

A natural question is: how do continuous and annualized yields compare to each other? The answer is that continuous yields are always slightly lower but the two are usually fairly close. You might at first think that continuous yields should be slightly higher. After all, we know that the more frequently you compound—the bigger m —the more interest you earn at a fixed nominal interest rate and the [CONTINUOUS APPROXIMATION 5.3.11](#) informally means letting $m = \infty$. But, here we have turned the question around: we want to achieve a fixed amount of growth in our investment from B to S in a fixed time of y years. The annualized yield r_a is the nominal rate which would achieve this with annual compounding. If we used this rate r_a in the [CONTINUOUS APPROXIMATION 5.3.11](#) we'd wind up with a sum *larger* than S because we're compounding more frequently. So to end up at S while compounding continuously we have to use a rate r_c a bit *lower* than r_a . Let's check this with some examples.

EXAMPLE 5.4.18: What are the continuous yields of the two investments in [EXAMPLE 5.4.5](#)—recall that the first costs \$5,000, returns \$10,000 in 10 years and had an annualized yield $r_{a1} = 7.17\%$ while the second costs \$4,000, returns \$5,000 in 3 years had an annualized yield $r_{a2} = 7.72\%$. The corresponding continuous yields are

$$r_{c1} = \frac{100}{10} \cdot \ln\left(\frac{10000}{5000}\right) = 6.93$$

and

$$r_{c2} = \frac{100}{3} \cdot \ln\left(\frac{5000}{4000}\right) = 7.44\%.$$



5.4 Yields

In both cases, the continuous yield is about a quarter of a percent lower than the annualized yield as we predicted. Note, however, that difference between the two remains about a half a percent.

Here are a few problems for you to try.

PROBLEM 5.4.19: Which has the highest and which the lowest continuous yield of the following three investments: a Treasury bill which costs \$24,850 today and returns \$100,000 in 30 years; a CD which costs \$1,000 today and earns 4.5% interest compounded daily over the next 5 years; and a house which cost \$120,000 12 years ago and is worth \$205,000 today?

PROBLEM 5.4.20: Find the continuous yields on a CD with a nominal rate of 5.25% compounded quarterly and on one with a nominal rate of 5.20% compounded daily? Which offers a higher continuous yield?

PROBLEM 5.4.21: Use continuous yields to decide, which is a better investment? A CD with a nominal rate of 4.45% compounded daily or a bond which sells for \$8,000 today and returns \$10,000 in 5 years.

PROBLEM 5.4.22: In 1972, I bought a bottle of 1967 Chateau d'Yquem for \$17.35. In 1997, a bottle of the same wine sold at auction for \$625.00. What continuous yield would I have realized had I auctioned my bottle instead of drinking it with friends for my wife's 50th birthday?

PROBLEM 5.4.23: Loan sharks charge compound interest of 2% a week compounded weekly: see [PROBLEM 5.2.24](#). What is their continuous yield?

If you compare these problems with [PROBLEM 5.4.20](#) to [PROBLEM 5.4.22](#), you'll see that all the continuous yields are indeed a bit lower than the corresponding annualized yields. However, in all cases, the ranking of investments by annualized yield is the same as the ranking by continuous yield. This is no accident.

5.4 Yields

CHALLENGE 5.4.24: Show that if one investment has a higher annualized yield than a second, then it also has a higher continuous yield and vice versa.

In other words, if we only want to compare investments it does not matter which of the two yields we use. The rates we get will be a bit higher if we use annualized yield and a bit lower if we use continuous yield. This explains why what you see are annualized yields. The people publishing these yields are generally trying to make their investments look attractive. So they prefer the yield which results in higher percentages: they'd use centennialized yields if they could get away with it because these would appear higher still. However, as a practical matter the difference is small and either yield is equally good for comparing investments.

CHALLENGE 5.4.25: Just as stating yields in terms of very short compounding periods—in the extreme, continuous compounding—shrinks them a bit, stating yields in terms of very long compounding period can really make them appear to grow. Show that an investment with a 5% annualized yield has a decennialized yield of 6.2% and a centennialized yield of 130.5% (yes, that's 130.5%). Hint: First, compute the amount A_T to which \$1 would accrue at interest of 5% compounded annually over term of 10 and 100 years. Then, find the interest in each amount and divide by the corresponding term. You should get the numbers above. Finally, explain why this procedure gives the yields asked for.

There's one possibility we haven't considered yet. What if your investment decreases in value? Then, you must have a *negative* yield. You can still use the formulas above to compute either an **annualized yield** or **continuous yield** with no change as the next exercise shows. I've done the first part to get you started.

PROBLEM 5.4.26: Suppose that you invest \$1,000. Find the (negative) annualized and continuous yields you obtain if:



5.4 Yields

- i) After 3 years your investment returns \$800.

Solution

To find the annualized yield, we just plug in for r_a :

$r_a = 100 \cdot \left(\sqrt[3]{\frac{S}{B}} - 1 \right) = 100 \cdot \left(\sqrt[3]{\frac{800}{1000}} - 1 \right) = -7.16822333$ so the annualized yield is about -7.2% .

To find the continuous yield, just plug in for r_c :

$r_c = \frac{100}{y} \cdot \ln\left(\frac{S}{B}\right) = \frac{100}{3} \cdot \ln\left(\frac{800}{1000}\right) = -7.438118376$ so that the continuous yield is about -7.4% .

- ii) After 5 years your investment returns \$700.

- iii) After 8 years your investment returns \$300.

Are you bothered by anything about the answers above? I hope so. Ignoring the sign, the annualized yield is also *less* than the continuous yield. I just convinced you that continuous yields would be smaller—see the comments following the [CONTINUOUS YIELD FORMULA 5.4.17](#). What gives? The one difference in the problems above is that they involved negative yields rather than positive ones so this must be what gives.

PROJECT 5.4.27: Convince yourself that *when yields are negative continuous yields will always be larger in size than annualized ones*.

- i) First, compare the future value of \$1,000 which earns a nominal rate of $(-6)\%$ for 5 years compounded annually with the future value at the same rate when compounded continuously.

- ii) Next, compare the present values of \$3,500 due in 8 years if the nominal rate is $(-7.5)\%$ when compounded annually and when compounded continuously.

- iii) Explain why the answers to i) and ii)) are both what we'd expect from the italicized claim.

- iv) Examples are a good first step in verifying a claim but mathematicians are professional skeptics. Can you come up with an argument to convince me that the claim above is *always* right?

The rules of 69.3 and 72

There is one feature of the [CONTINUOUS YIELD FORMULA 5.4.17](#) which is considerably nicer than the [ANNUALIZED YIELD FORMULA 5.4.4](#). The term in years y is out in the open in the [CONTINUOUS YIELD FORMULA 5.4.17](#) while it is hidden in the radical or exponent in the [ANNUALIZED YIELD FORMULA 5.4.4](#). This means we can easily solve for y in the former formula:

TERM EQUATION 5.4.28: $y = \frac{100}{r_c} \cdot \ln\left(\frac{S}{B}\right).$

It's tempting to ask whether we can make any use of this formula. The kind of "How long?" question which we'd be able to answer is one which seldom arises in everyday financial transactions. In these, the term of the transaction is almost always fixed. However, when we want to think informally about investments, worrying about the big picture rather than the last few pennies, this question can be quite helpful. The typical example of how we'll phrase it is the question "How long will it take my money to double?" which we can use this to get an idea of how long we will have to wait to achieve financial goals without a lot of decimals.

What does "doubling my money" mean? Just that the selling price of my investment is twice the buying price: $S = 2B$ or $\frac{S}{B} = 2$. The corresponding period of time is usually called a *doubling time* and we'll denote it by y_2 . Plugging into our equation for y , we find $y_2 = \frac{100}{r_c} \cdot \ln\left(\frac{S}{B}\right) = \frac{100}{r_c} \cdot \ln(2)$. For future reference, we'll record this.

DOUBLING TIME FORMULA 5.4.29: *The doubling time y_2 and the continuous rate r_c are related by the formula*

$$y_2 = \frac{100}{r_c} \cdot \ln(2).$$

However, we will not use this formula until a later section. The reason is that the $\ln(2)$ is not suitable for thinking informally. We'd at least need a calculator to use the formula. But we *can* get an easy rule-of-thumb formula if we just plug in and simplify. Since

5.4 Yields

$\ln(2) \simeq 0.6931471806 \simeq 0.693$, we find that $\frac{100}{r_c} \cdot \ln(2) \simeq \frac{69.3}{r_c}$. This gives the

RULE OF 69.3 5.4.30: *The time in years y_2 which it will take money to double at a continuous yield of $r_c\%$ is roughly $\frac{69.3}{r_c}$.*

The main virtue of this formula is its arithmetic simplicity. You just divide 69.3 by the rate in percent to estimate the doubling time in years. As a rule-of-thumb for making rough or “ballpark” estimates, it has two defects. First, what you’ll usually have in mind when you are making such estimates are annualized rather than continuous yields. We certainly don’t want a rule-of-thumb that requires us to convert between the two kinds of yields. This is easily resolved: just use the annualized yield if that’s easier. In fact, if all you have is a nominal rate, use that. As we’ve seen all these rates differ, but in the practical range of 0-20% they don’t differ by very much. Since a rule-of-thumb is only asked to give rough estimates, we can afford to ignore these small differences.

The second problem is that decimal, 69.3. So what if we can use the first rate we hit on. If we have to divide into 69.3, many of us will have to get out our calculators. The point of a rule of thumb is to avoid this. How can we do so? The first idea to simply round to 69. Still kind of an awkward number. Why not round up to 70 which is a bit “rounder”? Making the numerator a bit bigger also helps compensate for the fact that the nominal or annualized rates we likely use are a bit bigger than the r_c in the **RULE OF 69.3 5.4.30**. You sometimes see this rule (called the “rule of 70”) in books. But having decided to play fast and loose with the numerator, there is a much cleverer replacement: 72. This is cleverer for two reasons. First, the increase does a better job of compensating for fact that the rates we’ll be using are higher than the continuous yield in the **RULE OF 69.3 5.4.30**. Second, 72 is in some ways “rounder” than 70: it’s evenly divisible by 2, 3, 4, 6, 8, 9, 12. Better still, we can just have our cake and eat it too by using the:



5.4 Yields

RULE OF 69.3, 70 AND 72 5.4.31: *The time in years y_2 which it will take money to double at interest of $r\%$ is roughly equal to any of the three fractions $\frac{69.3}{r}$, $\frac{70}{r}$ or $\frac{72}{r}$. So as long as we only want a rough estimate of the doubling time, we can use whichever of the 3 numerators is easiest, and we can use any nominal rate for the denominator r .*

Amazingly, this rule was first stated in 1494(!), almost a hundred years before the invention of logarithms, in the book *Summa de Arithmetica* of Fra Luca Pacioli:

“If you want to know for every percentage interest per year, in how many years your capital will come back doubled, hold the rule of 72 in mind, which you always divide by the interest, whatever the quotient, in that many years it will be doubled. Example: When the interest is 6 percent per year, I say that you divide 72 by 6; you get 12, so in 12 years your capital will be doubled.”

EXAMPLE 5.4.32: Here’s an example of the kind of rough and ready estimation for which this rule is suited extending [EXAMPLE 5.2.14](#). Suppose my parents make an initial contribution of \$10,000 to my daughter’s college fund when she is born. How much can I anticipate having in the fund when she is 18? The answer clearly depends on what interest rate I can realize, but how. With a very conservative investment like a CD, I can expect to earn about 4%. This means my money will double in $y_2 = \frac{72}{4} = 18$ years so I’ll have \$20,000.

Investing in bonds, which are somewhat riskier—there’s a chance of default—I can expect to earn 6%. Now my money will double every $y_2 = \frac{72}{6} = 12$ years. So, the \$10,000 would double once to \$20,000 when she’s 12 and a second time to \$40,000 when she’s 24. When she’s eighteen it would be somewhere in between, say about \$30,000. Actually, it’s closer to \$28,000 but the whole virtue of this kind of rough estimation is that such differences are small enough to ignore.



5.4 Yields

If I add some stocks to the mix, increasing my risk again, maybe I can expect to earn 8%: now my doubling time is $y_2 = \frac{72}{8} = 9$ years. In 18 years, my money doubles twice to about \$40,000.

Hmm, perhaps I should really gamble and put my money into emerging markets, hoping for a yield of 10%. Now I'll switch to a numerator of 70 (why?) to find that my doubling time is about $y_2 = \frac{70}{10} = 7$ years. So, I'll have \$20,000 when she's 7, \$40,000 when she's 14 and \$80,000 when she's 21. (Note that every time I wait one doubling-time, I double the sum at the start of the period, not my initial \$10,000. The second doubling from age 7 to 14 doubles \$20,000 and the third from ages 14 to 21 doubles \$40,000.) At age 18, I'd expect to be roughly halfway between \$40,000 and \$80,000, at around \$60,000.

What if I could figure out some way to earn 12%? Then my doubling time would drop to $y_2 = \frac{72}{12} = 6$ years, and by the time my daughter is 18 my \$10,000 would have doubled 3 times to \$80,000.

What conclusions can I draw from this analysis? First, I am clearly going to need to make some contributions of my own to the fund. There's no way the initial gift can grow to the \$120,000 I think I'll need. In fact, with prudent investment strategies and realistic yield assumptions, I'll need to provide the bulk of the fund. Better start working on this now. The right strategy is discussed in section [SECTION 5.6](#).

Here are some problems for you to try. Do them using the [RULE OF 69.3, 70 AND 72 5.4.31](#).

PROBLEM 5.4.33: How long does it take money to double at interest of 5%? 7%? 9%? Express your answers to the nearest year.

PROBLEM 5.4.34: About how much will \$20,000 in your retirement account at age 25 amount to by the time you are 65, if the account has a yield of roughly: 2%, 5%, 7%, 9%, 12%? About how much will \$20,000 in your retirement account at age 45 amount to by the time



5.4 Yields

you are 65, if the account has a yield of roughly: 3%, 6%, 7%, 10%, 14%?

PROBLEM 5.4.35: How long does it take a debt carried on a credit card with an 18% interest charge to double?

Finally, we should remark that we can turn the [RULE OF 69.3, 70 AND 72 5.4.31](#) around. If we are given the doubling time y_2 for an investment, we can use these to determine its *approximate* yield r . Since $y_2 \simeq \frac{69.3/70/72}{r}$, we have $r \simeq \frac{69.3/70/72}{y_2}$. I emphasize the word approximate here because we'll usually only know the doubling time approximately. Here are a few problems for you to try.

PROBLEM 5.4.36: Determine the approximate annualized yield which an investment would provide if it has a doubling time of:

i) 10 years?

Solution

Just plug in: $r \simeq \frac{72}{y_2} = \frac{72}{10} = 7.2$ so the yield is about 7%. Note that I do not expect the rule of thumb to give me the tenths of percent so I rounded to the nearest percent.

ii) 6 years?

iii) 18 months?

What if your investment decreases in value? An investment which decreases in value is never going to double so you might think that the [RULE OF 69.3, 70 AND 72 5.4.31](#) could never apply. Actually, only a small change is needed to make it useful for such situations. The key is to make the correct analogy: *doubling* is for an investment which grows, as *halving* is for an investment which shrinks. So we should ask for a halving time, $y_{\frac{1}{2}}$: that is, ask when will $\frac{S}{B} = \frac{1}{2}$. We can answer this using formula [TERM EQUATION 5.4.28](#) just as for the doubling time. We find

$$y_{\frac{1}{2}} = \frac{100}{r_c} \cdot \ln\left(\frac{S}{B}\right) = \frac{100}{r_c} \cdot \ln\left(\frac{1}{2}\right) = \frac{100}{r_c} \cdot -0.693147 \simeq \frac{-69.3}{r_c}.$$

The only thing that's new is the minus sign which reflects the fact that $\ln\left(\frac{1}{2}\right) = -\ln(2)$. (This is a general property of logarithms—it still holds if we replace $\frac{1}{2}$ by $\frac{1}{x}$ and 2 by x —which we won't go

5.4 Yields

into.) We saw in [CONTINUOUS YIELDS](#) that we can still talk about the continuous yield for an investment decreasing in value: r_c is simply *negative*. The minus in r_c will cancel the minus above and we'll wind up with a positive number of years for the halving time, as we'd hope. In practice, the formula is so close to the formula for a doubling time that there's no point in having two. Instead, we'll just use the existing [RULE OF 69.3, 70 AND 72](#) [5.4.31](#) but remember that, when the yield is negative—and hence we get a negative value for y_2 —we should just drop the minus and interpret the time as a *halving* time. We'll see several applications for this in [SECTION 5.5](#).

PROBLEM 5.4.37: Find the approximate halving time of an investment whose annualized yield is -6% , -9% , -12% .

Finally, we can also turn a halving time into an approximate yield just as for a doubling time above.

PROBLEM 5.4.38: Find the approximate yield (which will be negative) on an investment whose value halves

- i) every 6 years.
- ii) every 8 years.
- iii) every 10 years.

Here are a few final general practice problems.

PROBLEM 5.4.39: We can combine yields and doubling times and pass from buying and selling prices to doubling times. Use continuous yields to find the doubling time of an investment which

- i) was purchased for \$1,200 and after 3 years is worth \$1,500?
- ii) was purchased for \$800 and after 6 years is worth \$1,200?
- iii) was purchased for \$23,000 and after 18 months is worth \$26,500?

Solution

We proceed in two steps: first find the yield and second convert this to a doubling time. In step one, we have $B = \$23,000$, $S =$

5.4 Yields

\$26,500 and $y = 1.5$ —note that we needed to convert the time to years because we want an *annualized* yield—so that

$$r_c = \frac{100}{y} \cdot \ln\left(\frac{S}{B}\right) = \frac{100}{1.5} \cdot \ln\left(\frac{26500}{23000}\right) = 9.443367800$$

for a yield of about 9.4%. (Here I have hard buying and selling prices so I can compute a fairly accurate yield. But, if I really wanted accuracy I'd have to use the [ANNUALIZED YIELD FORMULA 5.4.4](#). Check, if you like, that the annualized yield here is 9.90%.)

Step two is even easier. I'll use the numerator 69.3 since I am already using decimals in my yield: $y_2 \approx \frac{69.3}{9.4} = 7.37234042553$ so the doubling time is about 7 and a third years. Let's check this by computing the future value of \$23,000 at 9.90% interest compounded annually for 7.37 years. What answer should we expect? We have $p = 0.099$ and $T = 7.37$ so we actually obtain $\$23,000(1 + \frac{0.01 \cdot 9.9}{1})^{7.37} = \46119.1005943 or a bit more than the expected \$46,000.

PROBLEM 5.4.40: The same steps apply to investments that shrink. In [PROBLEM 5.4.26](#), we carried out the first half—finding the yields—for three such investments. What is the halving time of each of these investments?

Solution to i)

We just plug the yield of -7.4% into the [RULE OF 69.3 5.4.30](#) to find: $y_2 = \frac{69.3}{-7.4} = -9.364864865$. Since this is negative, we realize that what we have is a halving time, and since we are using a rule of thumb we round the answer. The halving time is about 9 years.

PROBLEM 5.4.41: Use the [RULE OF 69.3 5.4.30](#) to find the continuous yield on an investment whose value doubles in 18 months? We'll use this answer in discussing [MOORE'S LAW 5.5.18](#).

5.5 Applications of compounding

The formulas which we have developed in the last few sections can be used to study a number of problems having nothing to do with interest at all. It's amazing how often this happens in mathematics. You isolate some essential pattern in one problem and then you notice the identical pattern in a completely unrelated subject. You can then apply the techniques you developed to study the first problem to the others without doing any additional work. It's a kind of "something for nothing" which is one of the things which makes mathematics so useful. So, in this section, we'll pause to look at some of these other applications before going back and looking at the more complex topic of amortization.

Inflation

When we introduced the topic of the college fund as an example of a [FUTURE VALUE 5.2.12](#) problem, I computed the \$120,000 balance as 4 times the *current* \$30,000 cost per year of sending a child to a good private university. There's a major problem with that computation. My daughter isn't going to college today, she'll be going in 18 years. So the amount I really need to provide for is the 4 times the annual cost of sending a child to a good private university *in 18 years*. I don't (and can't) know what that cost is going to be exactly. However, we can estimate it using the [COMPOUND INTEREST FORMULA 5.2.4](#).

INFLATION 5.5.1: The general term used by economists to describe rising prices or costs is **inflation**. In the rare cases where prices fall, they speak of **deflation**.

The metaphor is supposed to make you think of a tire or balloon—the good or service—which can be filled with more or less air. The

5.5 Applications of compounding

amount of air corresponds to the price so when we raise the price we need to inflate.

The standard measure of inflation is the *rate of inflation* which is most often stated in the news as a percentage—"The Department of Commerce announced to day that in March the inflation rate stood at 3.4%"—but as the name rate suggests this is a concise way of referring to what is really a percentage-per-year. Informally, the inflation rate is the percentage by which prices will rise in a year. In other words,

INFLATION RATE 5.5.2: The inflation rate is the annualized or nominal annual rate at which prices are rising.

I'm already starting to use terminology from our study of compound interest to describe inflation but I can hear some of you protesting. "Just a minute, Dr. Morrison." You're thinking that prices behave very differently from the balance in your bank account and you're correct. Prices don't rise periodically (by predictable jumps which are timed at regular one period intervals) but sporadically (by jumps which are unpredictable in both size and timing). Moreover, the prices of different goods and services change somewhat independently of each other. Milk doesn't go up 2 cents a quart just because gas went up 8 cents a gallon. An inflation rate is some kind of average of the changes in prices of all good and services. It's more like an average of a large number of individual bank accounts earning interest at different rates and periods but even that's a simplification. In fact, just deciding how to make up this average is a subject of some controversy amongst economists. If they can't even agree on how to compute it with all the prices in front of them, how can we deal with it *without* the prices?

This kind of question comes up again and again when you want to use math to study a real-life problem. Real life is complex, disorderly and full of wrinkles and exceptions. Math is hard enough when you try to use it to study simple problems. If you incorporate all the com-

5.5 Applications of compounding

plexity of the real world into a mathematical problem, you generally get a problem you can't solve. People who use math to study the real world generally confront such difficulties in three stages.

In step 1, they try to find ways to simplify the questions they are interested in—by ignoring certain features or assuming that properties which are only roughly true actually hold exactly—until they can state their questions as mathematical problems which they can solve. You wish you could paint *The Last Supper* but you can't: instead of just giving up, you draw a cartoon instead. The goal is to capture enough of the essential features of the problem that you can still recognize it—we'd like to still be able to pick out Christ and the Apostles in the cartoon. The polite term for this process is *mathematical modelling* but people also speak of making a *caricature* and this term captures the flavor of what goes on somewhat better.

Step 2 is to do the math. This can range from very easy to very hard depending on the particular problem. In **MATH⁴LIFE**, we'll stick to the easy range. If things work out, you now have an answer to a math problem. But you were interested in a real-life problem.

Step 3 is to ask what the mathematical answer tells you about that problem. This involves more than just interpreting numbers. Usually, to check that the answer we have applies in the real world we need to see if it makes predictions which can be checked against known facts. The danger is that in step 1 we simplified too much and while we were able to get an answer it no longer accurately reflects reality.

All three steps require considerable care and effort. But, there is almost no field in which careful mathematical modelling has not proved useful. And I do not just mean the sciences. Mathematics has been used to determine the authorship of ancient manuscripts, to reconstruct the migrations of stone-age peoples from the genetic makeup of their descendants, to plan the U.S. economy in World War Two, to understand how our mind converts the stream of photons



5.5 Applications of compounding

hitting the retina into the images we see ...—I could go on indefinitely.

With this philosophy of trying to use mathematics to model the real world, let's get back to inflation. OK, we speak of inflation as a percentage rate. But, why does that make it have anything to do with compound interest. When we say that the inflation rate is 5% a year, all we mean is that something that costs \$100 last year will cost \$105 today. The extra \$5 is the inflation or rise in price. In general, the new price is the old price plus the inflation rate times the old price. Another way of putting this—the one we'll prefer—is to say the new price is the old price *times* 1 plus the inflation rate. The key point here is that *what counts is not the current price but the time interval*. If we wait a year, *any* price will rise 5%. A good that cost \$500 will rise 5% to \$525, one that cost \$20 will rise 5% to \$21.

This certainly looks familiar and so we should at least ask: “Does inflation work like either simple or compound interest?”. We can rule out one answer: it's not simple interest. If it were, then a year from now that \$105 item would cost \$110. But, to get next year's price we add 5% of *this* year's price of \$105. That's \$5.25 and so the price next year will be \$110.25. I hope you recognize that quarter: it's like the interest on interest—5% of \$5.00 is 25 cents—and tells us that inflation works like *compound* interest. In fact, the [COMPOUND INTEREST FORMULA 5.2.4](#) applies directly to inflation because the [ONE PERIOD PRINCIPLE 5.2.3](#) on which it is based does. Recall that this principle says that “to get from the amount at the start of a period to the amount at the end of that period you multiply by $(1 + p)$ where p is the periodic rate.” That's exactly how the previous paragraph says inflating prices behave. In other words, inflation is like interest on prices and the inflation rate really is a nominalized or annualized rate at which prices are rising.

So, can we just apply the [COMPOUND INTEREST FORMULA 5.2.4](#) to some questions about inflation? Not quite; one key ingredient is

5.5 Applications of compounding

missing. Before going on, see if you can spot what it is.

That's right: to use [COMPOUND INTEREST FORMULA 5.2.4](#) we need to know the *compounding frequency* m . As I noted above: “Prices don't rise periodically (by predictable jumps which are timed at regular one period intervals) but sporadically (by jumps which are unpredictable in both size and timing).” In other words, the [ONE PERIOD PRINCIPLE 5.2.3](#) seems to capture the basic mechanism of inflation nicely except that we have no idea what a period is. What we'd like is some way to calculate something like compound interest which does depends only on the annualized rate—the rate of inflation—but not on any compounding frequency. Put this way, a moment's recollection shows that we're in luck: this is exactly what the [CONTINUOUS APPROXIMATION 5.3.11](#) does

Maybe, you're worried because the [CONTINUOUS APPROXIMATION 5.3.11](#) is just what it says—only an approximation to compound interest. How do we know that such an approximation will give sensible answers when applied to a quantity like inflation which itself is only approximately like compound interest? Good question. There are several ways to convince yourself that we'll get useful answers. On a general level:

TRIANGLE RULE 5.5.3: If two quantities (like inflation and the [CONTINUOUS APPROXIMATION 5.3.11](#)) are both close to a third (compound interest), they must be close to each other.

This idea is a very basic one which appears all over the place in mathematics. The name triangle rule comes from drawing a picture with the three quantities as the vertices of a triangle.

PROBLEM 5.5.4: Draw a picture of a circle of radius 1 (and hence diameter 2) with the center labelled C . Now put place two points A and B anywhere you like in the circle. The points A and B are both close to C —within 1 unit. Now ask: what's the furthest apart that A and B can be from each other?

Hint: In the worst case, the triangle degenerates to a line segment.



5.5 Applications of compounding

A more direct argument for applying the [CONTINUOUS APPROXIMATION 5.3.11](#) to inflation depends on recalling how we obtained the [CONTINUOUS APPROXIMATION 5.3.11](#). We simply let the compounding frequency get very large, or equivalently, let the compounding periods get very short. We can think of inflation somewhat similarly. We already noted that the rise in the general level of prices is really made up of many independent jumps in the prices of individual goods and services. We can think of each such jump as a tiny rise in overall prices. When gas goes from \$3.32 to \$3.42 we think of this not as a 3% rise in the price of gas but as a much smaller rise in overall prices—smaller because only a small fraction of our money is spent on gas. The next day when milk goes up 2 cents a quart, we view this too as a small rise in overall prices (much smaller since we spend less on milk than gas). The inflation rate tells us what annualized increase in *overall* prices all these many tiny changes caused by rises in the prices of *individual* goods amount to. Instead of thinking of prices as compounding in discrete chunks, we view them as rising smoothly and continuously like the level of a liquid in a jar.

Later in this section, we are going to consider several other quantities which, like price levels, grow in accordance with the [ONE PERIOD PRINCIPLE 5.2.3](#) except that there is no regular compounding period. Instead, the overall change in the quantity is the result of an accumulation of many small changes in its parts. All these quantities can be modelled (that is, approximated and studied) using a form of the [CONTINUOUS APPROXIMATION 5.3.11](#). Rather than repeat the discussion above several times, we now give a general statement which applies to them all.

GENERAL COMPOUNDED QUANTITY 5.5.5: *We will say that a quantity is a general compounded quantity if it varies according to the [ONE PERIOD PRINCIPLE 5.2.3](#). That is, there is a period of time and a fixed periodic rate p such that regardless of the value B of the quantity at the start of a period, you can find its value S at end of that period by*



5.5 Applications of compounding

multiplying the starting value B by $(1 + p)$. The crucial point which makes for a compounded quantity is that the same magic factor of $(1 + p)$ works whatever starting amount B we choose or, what is the same thing, for whatever time we choose to start the period.

In discussing inflation above, we've worked with years as periods so that the periodic rate is an annualized rate (like the 5% rate of inflation used above). This is purely to keep things simple for you. With a bit more arithmetic, any period can be used to replace the year. Of course, as with interest, if we change the period, we will change the periodic rate p and hence the magic factor $(1 + p)$ but the same basic principle will apply: *over equal periods, any two starting amounts get multiplied by the same magic factor.* We'll use this principle often in the rest of section.

For discussing inflation, I chose years because that's the period the media commonly use. But, unlike interest calculations, where there is a single right "to-the-penny" answer and we have to use the right periods to get it, for a quantity like inflation, there is no such canonical choice for the periods and no way to get accurate answers. Once we accept this, we're free to choose our periods to make is easiest to get approximate answers. In working with yields in [SECTION 5.4](#), we have already seen that annual compounding is not best choice if such approximate answers are our goal. Instead, we'll choose to compound continuously. One immediate advantage is that we can always measure time in years.

GENERAL CONTINUOUS APPROXIMATION 5.5.6: *If B and S are the values of a [GENERAL COMPOUNDED QUANTITY 5.5.5](#) at two times that are y years apart, then $S_y \simeq B e^{(0.01r \cdot y)}$.*

The notation reflects the fact that we'll usually be thinking more in terms of a starting value B to an ending value S . I hope the reason for the name is clear: if we replace B by A_0 and S by A_T , the [GENERAL CONTINUOUS APPROXIMATION 5.5.6](#) turns into the [CONTINUOUS APPROXIMATION 5.3.11](#). There's one more point to note about this for-

5.5 Applications of compounding

mula. Although we can easily use it to get lots of decimals, we are usually only trying to get a rough approximation from it. The reason is that, unlike an interest calculation where we know the exact constant interest rate and compounding period the bank will use, quantities like inflation involve fluctuating rates when we often can only guess at and don't compound at regular periods. They are enough like compound interest to make it worth applying the same formulas but are different enough that we don't expect the formulas to give accurate answers.

Now some problems to warm up. I have done a few parts worked as models.

PROBLEM 5-5.7:

i) Suppose the rate of inflation is 5% a year. If a coat costs \$120 today, about how much is it likely to cost in

a. 3 years.

Solution

Here $B = \$120$, $r = 5\%$ and $y = 3$. Plugging in we find that $S \simeq \$120e^{(0.05 \cdot 3)} = \139.42 . Thus the coat is likely to cost about \$140 in 3 years. (Note how I claim only an approximate idea of the future cost.)

b. 6 years.

c. 12 years.

ii) Suppose the rate of inflation is 9% a year. If a car costs \$23,000 today, how much is it likely to cost in

a. 4 years.

b. 8 years.

c. 12 years.

Solution

Here $B = \$23,000$, $r = 9\%$ and $y = 12$. Plugging in we find that $S \simeq \$23,000e^{(0.09 \cdot 12)} = \$67,727.63$. Thus the car is likely to cost a bit less than \$70,000 in 12 years.

There is another way to try to track inflation. Instead of focussing on prices, we can focus on the buying power of the dollar. This shrinks



5.5 Applications of compounding

due to inflation: as prices rise, we need more dollars to buy the same goods, hence each dollar is worth less. Economists like to express sums of money that will exist in the future in terms of “constant dollars”—the dollars that exist today. How do we make the conversion? Simple! Suppose I set $B = \$1$ in [GENERAL CONTINUOUS APPROXIMATION 5.5.6](#). Then the value of S will tell me how many dollars I need in y years to buy something which cost \$1 today—like a dollar. In other words, S tells me how many future dollars my present dollar is worth. This is close to but not quite what the economists want. They’d like to know how much I could buy today with \$1 from y years in the future. I could just set S to \$1 and solve for B . However, I can avoid solving by recognizing that what the economists are asking for is the present value (the value today) of \$1 from y years in the future and recalling that the [CONTINUOUS APPROXIMATION 5.3.11](#) can be used to get present values by simply using negative numbers of years. The next problem let’s you practice this.

PROBLEM 5.5.8:

i) What will the value, in terms of today’s dollar, of a dollar 4 years in the future be if the inflation rate is

a. 3%?

Solution

Here $y = -4$ (note the minus so I can get a present value), $r = 3\%$ and $B = \$1$. I find that $S \simeq \$1e^{(0.013 \cdot (-4))} = \0.8869204367 so the present value is about 89 cents. In other words, a dollar 4 years from now will only buy what 89 cents buys today. You can check that $\$1 \simeq \$0.8869e^{(0.013 \cdot 4)}$ if you are suspicious of that negative y .

b. 5%?

c. 7%?

ii) What will the value in terms of today’s dollar of a dollar 12 years in the future be if the inflation rate is

a. 3%?

b. 6%?



5.5 Applications of compounding

c. 9%?

Solution

Here $y = -12$, $r = 9\%$ and $B = \$1$. I find that $S \approx \$1e^{(0.019 \cdot (-12))} = \0.3395955256 so the present value is about 34 cents. We can check this: the car from the [PROBLEM 5.5.7](#) which will cost \$67,727.63 is worth $\$67,727.63 \times 0.3395955256 = \$23,000$ today.

The most interesting questions about inflation—what causes it? how can we try to control it?—are not mathematical but economic in nature. All the math does is help economists formulate and test models about inflation to try to come to grips with these basic questions. I will say (oversimplifying somewhat) that the basic mechanism is too much money chasing too few goods and services. The excess demand leads to generally higher prices which is another way of saying inflation.

Hyperinflation is an extreme case which sometimes results when governments attempt to compensate for the falling buying power of their currency by simply printing more. This of course means that even more money is chasing the limited supply of goods and causes even further inflation of prices. When inflation reaches very high levels, sellers try to get a jump on it by raising prices even further. They fear that before they can sell their goods and spend the proceeds, prices will have risen even further. To protect against this, they add a sort of “margin of safety” to prices. The government then has to start printing money even faster to service its debts and so on.

PROBLEM 5.5.9: In the 1970's, Argentina underwent a period of hyperinflation in which prices were rising an average of 30% a month.

i) What *annual* rate of inflation does this correspond to? HINT: This is like asking for the annualized yield of an investment given its monthly yield. Since we view inflation as compounding continuously, it is most appropriate to use the [CONTINUOUS YIELD FORMULA 5.4.17](#).

5.5 Applications of compounding

Now ask yourself: if I bought something for $B = \$100$, what price S could I sell it for in 1 month?

ii) At this rate of inflation, how much will a car which costs \$15,000 today sell for in 1 year? in 5 years?

PROJECT 5.5.10: Look up the history of the hyperinflation in the Weimar Republic in Germany in the 1920's which is often blamed for the rise of the Nazi Party. What was the highest *daily* rate of inflation? To what *annual* rate of inflation would this correspond? Suppose that a loaf of bread costs \$1 and that a house costs \$100,000 today. If you have enough just money to buy a house today, how long would it have taken in the Weimar Republic before you did not have enough money to buy a loaf of bread?

Population

Populations—of people, animals, plants ...—are the next topic we want to look at. If we can convince ourselves that these behave like a [GENERAL COMPOUNDED QUANTITY 5.5.5](#), then we can model them using the [GENERAL CONTINUOUS APPROXIMATION 5.5.6](#), something which was first done in the 19th century by the economist Malthus but which has important applications today in geography, biology and other areas.

If populations are to satisfy the [ONE PERIOD PRINCIPLE 5.2.3](#), there should be an annual rate at which they grow. Let's ask: what is this rate? A moment's thought suggests a first answer: populations grow because of new births and the rate we want to consider is thus the *birth rate*. This is half right. Can you see what's missing? We'd be done if we were immortal. Unfortunately, we have to account for deaths too. The net growth of a population is the difference of these two factors: if, for example, a population has a birth rate of 3.5% a year and a death rate of 1.5% a year, we'll simply say that the net *growth* rate of the population is the difference, 2%, of these two rates.



5.5 Applications of compounding

For example, if we had 10,000 people at the start of a year, we'd expect about 350 births and 150 deaths during the year for an net increase in population of 200 to 10,200. The growth during the year was 200 or 2% of the original 10,000 population.

Just being able to name a rate does not automatically guarantee that populations behave like a [GENERAL COMPOUNDED QUANTITY 5.5.5](#). We need to convince ourselves of the key fact that the percentage change over a given period—say a year—only depends on the period and *not* on the starting population. This certainly seems so be the case for birth and death rates. We intuitively view the birth rate as expressing a general propensity to have children (which may depend on the percentage of the population of childbearing age, on cultural factors like the desire to be able to provide a good education for every child and so on) which is uniform across the whole population and not affected by the size of the group we choose to measure. Similarly the death rate depends on things like number of elderly people, quality of medical care etc. but is again the same for all sizes of groups. Hence, the net growth rate is also independent of the size of the population we measure. The [ONE PERIOD PRINCIPLE 5.2.3](#) applies and we can use the [GENERAL CONTINUOUS APPROXIMATION 5.5.6](#) to model populations.

PROJECT 5.5.11:

- i) What's wrong with what I just said even if we assume that populations have uniform birth and death rates? More precisely, briefly criticize the argument above. What factors which might be important have I left out? (Hint: three you might want to consider are migration, war and plagues). Give some examples of ways in which my assumptions that populations are uniform might be inaccurate. How does that fact that birth and death rates change—both are much lower now than they were a hundred years ago—limit the validity of my argument?
- ii) One property of any [GENERAL COMPOUNDED QUANTITY 5.5.5](#) is



5.5 Applications of compounding

that it always grows. (Or, as we'll see is also possible when we look at radioactivity, always shrinks). Find statistics on the population of Europe (or of the country of your choice in Europe) over the period from 1000 AD to 1600 AD. You'll see both rises and declines in the population so the model of a [GENERAL COMPOUNDED QUANTITY 5.5.5](#) definitely does *not* apply. What factors were responsible? What is it about these factors that makes our model invalid for these populations?

iii) Another major objection my argument that populations of people behave like a [GENERAL COMPOUNDED QUANTITY 5.5.5](#) is that they do *not* have uniform birth and death rates. What kinds of factors identify groups with higher or lower birth and death rates? Give some examples of groups in the United States whose populations are growing more or less rapidly due to such internal factors.

The point of this project is to let you try your hand at the hard part of mathematical modelling, which is, remember, not usually the math but the analysis of how far the model corresponds to reality, what its uses and limitations are, and the work needed to overcome these limitations. The moral of parts [i\)](#) and [ii\)](#) seems to be that a population will behave like [GENERAL COMPOUNDED QUANTITY 5.5.5](#) only if there are no variable external factors which affect it.

The moral of part [iii\)](#) is that like inflation which is made up of lots of different changes in the prices of individual goods, population growth is made up of lots of different rates of growth in sub-populations. As with inflation, this does not eliminate but merely limits our ability to use the [GENERAL COMPOUNDED QUANTITY 5.5.5](#) model. We can't apply conclusions about general price changes to a single good like gasoline and we can't apply conclusions about general populations changes to a sub-population unless we *know* it is representative of a larger population.

This poses the question, "How can we tell when a group is representative?". It's a fascinating one but *very* hard to answer. The whole



5.5 Applications of compounding

aim of statistics is to find good ways to answer questions like this which come up almost anytime we try to model human phenomena (and in many other contexts).

The discussion above suggests that we look for populations which are highly uniform and whose growth is not influenced by external factors. The classic examples are populations of small organisms grown under controlled conditions.

PROBLEM 5.5.12: Let's consider the growth of bacteria in a laboratory. We'll assume that the scientists who run the lab provide the bacteria with everything they need to grow. Bacteria can reproduce very quickly by cell division. Suppose that a bacterium can divide every 3 hours. In other words, starting with $1 = 2^0$ bacteria after $0 = 0 \cdot 3$ hours, we get $2 = 2^1$ after $3 = 3 \cdot 1$ hours, then $4 = 2^2$ after $6 = 3 \cdot 2$ hours, then $8 = 2^3$ after $9 = 3 \cdot 3$ hours and so on.

What we're saying is populations of this bacterium have a *doubling time* of 3 hours. In this problem, we'll investigate the consequences of this statement in various ways.

- i) How many bacteria will there be after one day (i.e., 24 hours?)
Hint: The pattern above is that there are 2^k bacteria after $3 \cdot k$ hours. What does the pattern tell us will happen after one day?
- ii) How many bacteria will there be after 4 days?

You should find that there will be over 4 billion—4,294,967,296 !—bacteria after 4 days. If you think, that's a lot, after a week there will be 72,057,594,037,927,936 bacteria. Your calculator won't be able to give you this figure exactly as it has too many digits but it should tell you something like “.7205759404e17”.

To get an idea how big this is suppose that a billion bacteria weigh a gram (about one twenty-eighth of an ounce): then after a week you'd have over 72,000 kilograms (about 150,000 pounds!) of bacteria. This seems awfully big. Can we check our answers to be sure there is no mistake? Here are two ways to do so.



5.5 Applications of compounding

PROBLEM 5.5.13:

i) In [PROBLEM 5.4.36](#) we noticed that the [RULE OF 69.3](#) [5.4.30](#)— $y_2 \simeq \frac{69.3}{r_c}$ —can be turned around. Given a doubling time we can deduce a rate of growth: $r_c \simeq \frac{69.3}{y_2}$. We can do this here. We just have to be a bit careful about units. Since our doubling time is given in *hours*, our continuous rate r_c will be in percent-*per-hour* not percent-per-year. Show that the rate $r_c = 21.3\%$ -per-hour. Now that we have r_c , we can use the [GENERAL CONTINUOUS APPROXIMATION 5.5.6](#) to check the populations above. Once again, we need to be consistent about units and use hours rather than years. What populations do you obtain after 4 days and after 7 days starting with a population $B = 1$?

ii) Your answers in i) should have the same size as those in [PROBLEM 5.5.12 ii\)](#): after 4 days you should have about $4,000,000,000 \simeq 0.4e10$ bacteria and after 7 days about $0.7e17$. But, the two sets of answers will *not* match those to many digits. Why should we not expect them to?

iii) Here's an exact check which also illustrates the key property of a [GENERAL COMPOUNDED QUANTITY 5.5.5](#), namely, that we can use any period we like. In [PROBLEM 5.5.12 i\)](#), we found that a single bacteria grows to a population of 256 in one day. In other words, in one day a population of 1 gets multiplied by the magic factor of 256. But the key property of a [GENERAL COMPOUNDED QUANTITY 5.5.5](#) then says that, over a period of 1 day, *any* population should get multiplied by 256. A population of 256 should grow to $256 \cdot 256 = 65,536$ in a day. In other words, a population of 1 should grow to 256^2 in 2 days, to 256^3 in 3 days and so on. Use this to check the answer to [PROBLEM 5.5.12 b\)](#) exactly.

iv) OK, so there was nothing wrong with the math in [PROBLEM 5.5.12](#). But, the idea that after a week we'd have 72,057,594,037,927,936 bacteria—over 150,000 pounds—is clearly wrong. If the math is right, why is this answer wrong?

The social economist, Malthus, mentioned above, was the first to ob-

5.5 Applications of compounding

serve that human populations behave much like a **GENERAL COM-POUNDED QUANTITY 5.5.5** in the absence of external factors. However, his conclusions about our future—that population growth would eventually lead to a catastrophe have so far proven false: the world's population has long ago passed the figure he was sure could not be sustained. Some species do suffer such catastrophic growth and decline cycles. Have you ever seen photos of millions of lemmings rushing madly off the side of cliff in Scandinavia? This happens when the lemming population grows too rapidly and exhausts the supply of food.

PROBLEM 5.5.14:

- i) Suppose that: lemmings have litters of eight young of whom 6 are female and 2 are male; a female lemming can have her first litter when she is 3 months old and can have another litter every three months thereafter; no lemmings ever die. How long will it be before a starting group of 3 females and 1 male grows to a population of a million?
- ii) How does the answer to i) change if we change the assumptions and suppose that:
 - a. lemmings can reproduce every 2 months instead of every 3 months?
 - b. each generation of lemmings dies when the next generation is born?
 - c. each litter contains 7 females and 1 male?

PROBLEM 5.5.15:

- i) Suppose that every fruit fly lays 100 eggs and then dies and that all of these are ready to lay their eggs in a week. (I'm ignoring the male flies to simplify). Without using any of our formulas, can you say how many fruit flies would there be in
 - a. 2 weeks?
 - b. 3 weeks?
 - c. 4 weeks?



5.5 Applications of compounding

d. a general number T of weeks? Hint: Just try to see the pattern in the answers after 2, 3 and 4 weeks.

ii) How many weeks would it be before the offspring of a single fruit fly cover the entire surface of the earth to a depth of 100 meters (over the length of a football field) with flies?

Hint: First estimate how many flies it would take to cover the earth to this depth. You may assume that the area of the earth is about 10^{15} square meters and that fruit fly is a cube of side equal to 0.001 meters. (This means that over 300,000 would fit into a standard can of soda.)

iii) Why can we touch ground?

PROBLEM 5.5.16:

i) The informal numbers in [PROBLEM 5.5.15](#) can be obtained—with more effort—by considering fruit fly populations as a [GENERAL COMPOUNDED QUANTITY 5.5.5](#). Use the [CONTINUOUS YIELD FORMULA 5.4.17](#)—but with y standing for periods of weeks instead of years—to find a continuous weekly yield on fruit-flies.

Hint: Over a period of 1 week, we can take $B = 1$ and $S = 100$. Why?

ii) Now confirm the number of weeks you found in [PROBLEM 5.5.15 b](#)) by using the [TERM CONVERSION FORMULA 5.1.13](#) (with y standing for weeks again).

Hint: For S you can use the number of flies needed to cover the earth to a depth of 100 meters.

PROBLEM 5.5.17: Mexico City is the world's largest city. In 1980, the population of the metropolitan area was 8,000,000 and in the year 2000 it is 26,000,000.

i) Assuming that the population of Mexico City grows like a [GENERAL COMPOUNDED QUANTITY 5.5.5](#), estimate what the population was in 1990 and what it will be in 2010.

Hint: First use the [CONTINUOUS YIELD FORMULA 5.4.17](#) to deduce an annualized rate of growth from the two population figures given. Then, use this rate in the [GENERAL CONTINUOUS APPROXIMATION](#)



5.5 Applications of compounding

5.5.6 to produce the estimates: this is easiest if you measure time in year from 1980.

ii) What would the predicted population be in the year 2100? What conclusions do you draw from this answer?

Computation

Gordon Moore, one of the inventors of the integrated circuit and a founder of microprocessor giant Intel, is also famous for a rule which bears his name. He was trying to describe the rate at which the number of gates (basic logic elements) and transistors on the “latest” microprocessor grew as a function of time and he noticed that over the early period of their development these numbers had doubled roughly every 18 months. When he first enunciated this rule, few believed that this kind of *compounded* growth could continue. But it seems to have. In fact, many people regularly apply his rule to all kinds of computing equipment. Let’s state it in this general way.

MOORE’S LAW 5.5.18: *The complexity of the “latest” model of any type of computing equipment doubles every 18 months.*

Maybe it’s not obvious to you that this law says that the complexity of computing equipment is a **GENERAL COMPOUNDED QUANTITY 5.5.5**. There are no percent per year here at all. True, but remember that it is not the percent which make a **GENERAL COMPOUNDED QUANTITY 5.5.5**, it’s the existence of a “magic factor” which tells us how *any* initial quantity will change over some fixed period of time. In Moore’s Law, the period just happens to be 18 months and the magic factor is 2. Every 18 months the complexity of any piece of equipment doubles—that is, gets multiplied by 2—no matter how complex it was to begin with. In fact the mention of doubling should ring a bell. In **SECTION 5.4**, we saw that *continuously compounded* investment have a *doubling time* which depends only on the interest rate or *yield* of the investment, not on its size. The argument above

5.5 Applications of compounding

amounts to saying that converse is also true: if a quantity has a doubling time that does not depend on its initial value, it is a **GENERAL COMPOUNDED QUANTITY 5.5.5**. We can even say what p is: since the magic factor $(1 + p) = 2$ we see that $p = 1$. Put in another way, over a period equal to the doubling time, the “interest” which gets added to any starting amount of a compounded quantity equals the starting amount. If you look back to **PROBLEM 5.4.41**, you’ll see that it asks for the continuous yield which corresponds to a doubling time of 18 months. You should have found this to be 46.2% (because $r \simeq \frac{69.3}{y_2} = \frac{69.3}{1.5} = 46.2\%$).

PROBLEM 5.5.19: Let’s get an idea of what an enormously fast rate of growth this is. How much more complex is a 1999 processor than one which was start-of-the-art in

i) 1989?

Solution

All we need to do is use the **GENERAL CONTINUOUS APPROXIMATION 5.5.6** to continuously compound an initial value of 1 at a nominal rate of 46.2% for 10 years: we get $1 \cdot e^{(0.0146.2 \cdot 10)} = 101.4940321$. In other words, today’s processors are over 100 times as complex as those from just 10 years ago.

ii) 1979?

iii) 1969?

Of course, **MOORE’S LAW 5.5.18** does not apply to *software*. The best comment on this subject is an old hacker’s joke: “If the automobile industry could achieve the same kind of progress that the computer industry has, then cars would cost \$1,000, get a 1,000 miles to the gallon, run for 10 years without needing repair, and, once a month, go out of control and crash killing everyone on board.” If a computer ever ate your homework, you’ll agree.

PROBLEM 5.5.20: Here’s a short quote from a speech given by Gordon Moore about the history and future prospects for his law in 1997.



5.5 Applications of compounding

“This results in, you know, some rather dramatic changes in economics. The first planar transistors we sold in about 1959, the year the planar transistor was introduced, sold for several dollars. In fact, the first ones we shipped sold for \$1.50, I remember very clearly. But a good transistor in those days sold on the order for \$5 or \$6. Today, you can buy a 16 megabit DRAM for the same \$6. That’s something over 16 million transistors with all the interconnections and everything else for the price of a single transistor, something less than 30 years ago. This is really pretty dramatic.”

Let’s summarize this by saying that the price of a transistor has gone from \$1.50 apiece in 1959 to 16,000,000 for \$6 in 1997. Use this to calculate the “number of transistors per dollar” you could buy in 1959 and 1997. Let’s think of the value of this number in 1959 as a sort of buying price B and that in 1997 as a selling price S .

- i) What was the continuous “yield” on this number over the 28 years? Estimate the doubling time for this number and compare it with **MOORE’S LAW 5.5.18**.
- ii) How do your answers change if we use the other price Moore mentions of \$5 apiece in 1959?

PROBLEM 5.5.21: Monitors would seem to be an exception to **MOORE’S LAW 5.5.18**. In 1989, 12 inch monitors were common. In 1999, 21 inch monitors are the state of the art. That’s not even a single factor of 2. However, screen size is really a poor measure of complexity. We should really think about the amount of information the screen carries. This depends on two factors: the number of *pixels* or dots on the screen, and the number of *bits* of color information associated to each dot. The pixels tell us about the resolution of the screen. In 1989, resolutions of 640 by 480 were considered state-of-the-art for PC’s. A high-end monitor in 1999 had a resolution of



5.5 Applications of compounding

about 1,600 by 1,200. The number of bits b is related to the number of colors c which the monitor can show by the rule $c = 2^b$. For example, a black and white monitor carries just a single bit of information and shows $c = 2 = 2^1$ colors. A monitor which displays 16 colors—the latest in 1989, when these colors were actually shades of gray—carries 4 bits of information ($16 = 2^4$). A high-end monitor in 1999 used 32 bits of information per pixel. The “right” measure of the complexity of a monitor is the product of the number of pixels and the bits per pixel: that is, the total number of bits of information the monitor can display. In this problem, we’ll figure out how closely monitors have matched [MOORE’S LAW 5.5.18](#) based on these figures.

i) First, we want to calculate an annualized “yield” or rate of growth for monitor complexity. We can do this by viewing the complexity of a 1989 monitor as a “buying price” $B = 640 \cdot 480 \cdot 4$ and that of a 1998 monitor as a “selling price” $S = 1600 \cdot 1200 \cdot 32$ and then applying the [CONTINUOUS YIELD FORMULA 5.4.17](#). Show that the rate of growth is about 39% a year.

ii) Next, use the [RULE OF 69.3 5.4.30](#) and the answer to part a) to determine the actual *doubling time* for the complexity of monitors. How well does [MOORE’S LAW 5.5.18](#) seem to hold?

iii) How do the answers to i) and i) change if we assume that a high end monitor in 1989 could display 256 colors?

iv) How well did the rate of growth of i) hold up over the decade from 1999 to 2009?

PROBLEM 5.5.22: This problem looks at how [MOORE’S LAW 5.5.18](#) is reflected in the size and cost of computer memory or RAM.

i) In 1987, the author bought a state-of-the-art Macintosh II computer which came with 1-2Mb (megabytes) of RAM). In 1997, he bought a high end PowerMac which came with 128-256Mb of RAM. Using the larger of the two RAM figures for each year, determine the *doubling time* for quantity of RAM in Macintosh computers.

Hint: You will need to determine the *continuous yield* on RAM first,



5.5 Applications of compounding

then use this to get the doubling time: see [PROBLEM 5.4.39](#) for a similar example.

ii) Another way to try get a feel for [MOORE’S LAW 5.5.18](#) is to use it to track the decline in the cost of some piece of equipment of fixed complexity. In this exercise, we’ll carry this out for computer memory or RAM. Our fixed unit of RAM will be the megabyte. In 1984, this was a huge amount of memory: the first Macintosh appeared in that year and had 128K or only one-eighth this amount. The first year in which consumers could buy a 1Mb memory module was 1989. In 1996, you can no longer buy a new memory module this *small*. The table below shows representative costs per megabyte of general purpose RAM at 6 year intervals. Use this table to estimate the *continuous yield* on the price of a megabyte of RAM: since the prices have been falling, you’ll have a negative yield. Then use this yield to estimate the *halving time* for RAM prices.

Year	Size (mb)	Cost (\$)	Cost/mb
1985	0.5	399	798.00
1991	4	165	41.25
1997	32	104	3.25
2003	512	99	0.19
2009	2048	40	0.02

TABLE 5.5.23: SAMPLE HISTORICAL RAM SIZES AND PRICING

iii) The yield and halving time you get in b) depend on which two years from the table you use as the starting and ending years. How widely do these quantities vary? Try to suggest causes for this variation. Here’s one line of investigation. For mathematical reasons, memory modules come in sizes which are powers of 2 so we’d expect each generation to be twice as large as the previous one and to last for one of the doubling periods from part a). But, it’s a historical fact—I know of no really good reason why—that the size of common memory modules tends to increase by factors of 4 from one generation to the next. This is a typical example of an unexplained



5.5 Applications of compounding

divergence between a generally very good model and reality.

PROJECT 5.5.24:

i) Does **MOORE'S LAW 5.5.18** really describe how the power and scale of computer equipment evolves? This project asks you to make a practical test using some recent historical data from the personal computer industry. To begin with, go to your library and get out copies of a personal computing magazine like *Computer Shopper*, *PC World*, *MacWorld* etc. at six month intervals going back ten years. In each magazine, find a couple of ads for the “latest” model of PC and record basic data like: processor clock speed, amount of RAM, hard disk size in megabytes. Then ask yourself: do these quantities tend to double every 18 months as **MOORE'S LAW 5.5.18** claims. If not, do they grow faster or slower. How regular is the growth? Can you suggest any explanations for the irregularities? Predict what the hardware characteristics of the “latest” PC will be three years from now and six years from now.

ii) How does the *price* of the “latest” PC vary over this ten year period? **MOORE'S LAW 5.5.18** seems not to apply. Why?

iii) Suppose we decide that the correct way to measure the “cost” of a hard disk is by the price of each megabyte of storage (rather than by the total price of the drive). Use your magazines to determine a figure for this cost at yearly intervals over the past 10 years. You'll notice that it decreases quite rapidly like the cost of RAM discussed above. Use your figures to estimate the annualized rate of decrease of this cost and its *halving time*. How well do your figures match this compounded model. Is there a **MOORE'S LAW 5.5.18** for the price of a megabyte of hard disk storage?

iv) One quantity which definitely seems to double much less rapidly than every 18 months is processor “clock speed”. This exercise asks you to derive a possible explanation. Let's assume that the transistors on a processor are small squares which are laid out in a grid to form a larger square and processor clock speed is *inversely* proportional to the length of the side of this overall square. Although this



5.5 Applications of compounding

is a typical modeller's oversimplification which ignores many relevant factors, it contains a nugget of truth: The speed of a computer is roughly inversely proportional to how far the electrons which "do" the computation have to travel. Next suppose that every 18 months the length of the side of the small transistor square halves while the total number of transistors doubles. In other words both scale of a transistor and number of transistors obey Moore's law.

- Show that the factors by which the *area* of a single transistor, the area of a processor and the length of the side of a processor will change every 18 months are $\frac{1}{4}$, $\frac{1}{2}$ and $\frac{1}{\sqrt{2}}$ respectively.
- Given our assumptions about processor speed, the last answer to part iv)a says that processor clock speed increase by a factor of $\sqrt{2}$ every 18 months. Show—without using any formulas—that this implies that processor clock speed has a doubling time of 3 years.
- Confirm the answer to iv)b by using standard yield and doubling time formulae to compute first the "yield" on a quantity which increases by a factor of $\sqrt{2}$ every 18 months and then the doubling time which corresponds to this yield.
- How well do historical data on the clock speed of the "latest" Intel processor fit this predicted doubling time? Can you detect a point in time at which Intel began "over clocking" its processors? "Over clocking" is an innovation which allows chip makers to raise the published clock speeds of their processors without any essential change to the processors' design.

Radiation

With a few exceptions, just about every **GENERAL COMPOUNDED QUANTITY 5.5.5** we have looked at has gotten *larger* as time goes by. Put another way, the magic factor by which a starting amount gets multiplied over any period is *greater than 1* corresponding periodic rate has been *positive*. But, we can also consider cases where



5.5 Applications of compounding

amounts get *smaller* as time passes and where correspondingly, the periodic magic factors are *less than* 1 and the periodic rates are *negative*. In this subsection, we will look at one such example, radiation, or the decay of radioactive matter.

Here is the setup in a very simplified form. All matter is built up from a list of *atomic elements* which are characterized by the number of protons (positively charged particles) in the nuclei of atoms of that element. Hydrogen atoms contain one proton, helium atoms two, and so on up the *periodic table* of the elements. For example, carbon has 6 protons in its nucleus. However, the nucleus of an atom also contains neutrons (uncharged particles) and a single element can have different forms called *isotopes* whose nuclei contain different numbers of neutrons. Isotopes are generally named by giving the name of the corresponding element and the total number of protons and neutrons in the nucleus. For example, most carbon atoms have 6 neutrons; this common isotope is called carbon-12 (since 6 protons plus 6 neutrons gives 12). But a few carbon atoms have 7 or 8 neutrons; these rarer isotopes are called carbon-13 and carbon-14. (Why?)

Certain isotopes—usually those which contain a large number of neutrons—are *unstable*. Atoms of these isotopes tend to spontaneously undergo a process called *radioactive decay* or fission in which a particle in the nucleus splits into one or more smaller subatomic particles (releasing energy in the process). For example, carbon-14 undergoes β^- decay (essentially, a neutron is replaced by a proton and an electron) to nitrogen-14; potassium-40 decays into a mixture which consists of about 11% of argon-40 and 89% cadmium-40. The argon-40 isotope is stable but the cadmium-40 itself decays gradually.

Physicists and chemists who observed such decay found that although individual atoms appear to decay very sporadically, large groups of atoms decay in a very regular and predictable way. (We



5.5 Applications of compounding

can make an analogy to human populations: whether a single family (atom) will have a child is impossible to predict but the number of children born in an entire country can be accurately predicted by a birth rate.) The way physicists express this predictability is with a half-life. For each isotope, they have found that there is a period $y_{\frac{1}{2}}$ such that no matter what amount of the isotope you have to begin with, if you wait $y_{\frac{1}{2}}$ years you will end up with exactly half as much. For example, the half life of carbon-14 is $5.73 \times 10^3 \pm 40$ years. This means that if you have a gram today, then in 5,730 years (give or take 40 years) you will have a half a gram (the other half gram will have decayed into carbon-12 and carbon-13); if you have 10 grams today, you'll have 5 grams in 5,730 years. Likewise, potassium-40 has a half-life of 1.25×10^9 years so for 2 grams of potassium today to decay to 1 gram, you'll have to wait about one and a quarter *billion* years. Bring a lunch!

I hope you'll recognize the half-life of an isotope as the equivalent of a *doubling time* for a growing quantity. (We discussed this possibility earlier in the subsection on [THE RULES OF 69.3 AND 72](#) in [SECTION 5.4](#).) And just as having a doubling time—[MOORE'S LAW 5.5.18](#)—implies that the complexity of computer equipment is a [GENERAL COMPOUNDED QUANTITY 5.5.5](#), so having a half-life or halving time implies that the quantity of a radioactive element is a [GENERAL COMPOUNDED QUANTITY 5.5.5](#). The only difference, as suggested above, is that the *rate of decay* for an isotope will be negative. Moreover, since we are dealing with a physical law this is one situation in which we can hope to get accurate answers when we pass between half-lives and continuous rates so for the first time, we want to use the exact [DOUBLING TIME FORMULA 5.4.29](#) rather than one of the rules of thumb. Other than the need for a calculator, this is no harder than the rough conversions we have been doing, as the following example shows.

PROBLEM 5.5.25: If an isotope has a half-life of 29 years will it be



5.5 Applications of compounding

radioactive for exactly 58 years? Or was James right after all?

Bond His government's given you a bomb!

Goldfinger I prefer to call it an atomic device. It's small, but particularly dirty.

Bond Cobalt and iodine?

Goldfinger Precisely.

Bond Well, if you explode it in Fort Knox, the, uh, entire gold supply of the United States will be radioactive for... fifty-seven years!

Goldfinger Fifty-eight, to be exact.

EXAMPLE 5.5.26: Let's find the annualized rate of decay of carbon-14. We just rewrite the **DOUBLING TIME FORMULA 5.4.29** $y_2 = \frac{100}{r_c} \cdot \ln(2)$ as $r_c = \frac{100}{y_2} \cdot \ln(2)$ and plug in $y_2 = -5,730$ —we use the minus sign because we are dealing with a half-life and want the rate r_c to come out negative—to find: $r_c = -0.01209680943\%$ or about -0.0121% percent a year. Note that I rounded to three places because that is the accuracy to which we know the half-life of carbon-14. Also, I really mean about minus one-hundredth of a *percent* here, not the fraction minus one-hundredth.

We can go the other way even more easily. Suppose I know that an isotope has a decay rate of -2.831% a year. Then, I just plug in $r_c = -2.831\%$ in the **DOUBLING TIME FORMULA 5.4.29** and find that $y_2 = -24.48418158$: the negative value tells me I am dealing with a half-life of about 24.48 years. Question: Why did I choose the roundings above?

PROBLEM 5.5.27:

- Find the annualized rate of decay (expressed as a *negative* percent per year) of the isotopes potassium-40 with a half-life of 1.25×10^9 years—pay attention to that decimal point here—and tritium with a half life of 12.5 years.
- Estimate the half-lives of two isotopes which decay at rates of -0.00052% a year and -1.25% a year. Explain how you decided to round each half-life.



5.5 Applications of compounding

The fact that radioactive decay behaves in this predictable way has one surprising application: by measuring the quantities of certain isotopes which they contain, we can estimate the age of many astronomical, geological and archeological objects. Let me describe one application to archeology. The technique is based on the observation that the carbon in our atmosphere (mainly present in the form of carbon dioxide) contains a small and fixed percentage of the isotope carbon-14—historically, about 1 part in 1,013. Of course, the carbon-14 is always decaying but while it is doing so processes in the upper atmosphere which I won't enter into are producing new carbon-14 and the result is a steady level of carbon-14 in the air. The carbon in living organic material (our bodies, a tree, etc) contains the same small and fixed percentage of carbon-14 mainly because living material is constantly exchanging carbon with the environment and most of the inflow of carbon comes from the air. However, when something organic dies, this exchange ceases and nothing compensates for the decay of the carbon-14 present at death. The level of carbon-14 gradually decreases and by measuring this level we can tell how long something has been dead. For example, if a piece of wood contains half the level of carbon-14 present in living matter then the tree it came from must have died about one carbon-14 halflife ago. But, we can date any residual level of carbon-14 by simply applying the [TERM EQUATION 5.4.28](#) once we have the annualized rate of decay of -0.0121% found in [PROBLEM 5.5.27](#). Here are some examples.

PROBLEM 5.5.28:

- i) Estimate the age of a leather sandal which is found to contain .345 times the level of carbon-14 in present in living matter.

Solution

We want to use the [TERM EQUATION 5.4.28](#) $y = \frac{100}{r_c} \cdot \ln\left(\frac{S}{B}\right)$ to find the term y in years for which the carbon-14 in the sandal has been decaying. We know the rate of decay $r_c = -0.0121$ so the question we need to answer is: what values should we use for the buying price B and the selling price S ?

5.5 Applications of compounding

Since we are dealing with radioactivity rather than money, we want to use two carbon-14 levels. The level B is that at the start of the period we are trying to measure: this is just the baseline level of carbon-14 in living matter. The level S is that at the end of the period we are trying to measure: this is just .345 times the baseline level of carbon-14 in living matter.

The one missing ingredient seems to be the actual value of the baseline level of carbon-14. We do not know this level but a moment's thought shows that *we do not need to!* To use the formula, all we really need is the *ratio* $\frac{S}{B}$ and this we are *told* is 0.345. So we just plug in to find $y = \frac{100}{-0.0121} \cdot \ln(.0345) = 8,795.131091$. Note how the age came out positive as we'd like: the logarithm was negative and this cancelled the minus sign from r_c . How should I round this? I know the ingredients in my calculation to three places so I'd guess my answer is good to three places: the sandal is about 8.80×10^3 years old.

How can I check this answer? Well, if the sandal were one half-life old (that is 5,730 or so years old), I would expect to find half or .5 the baseline level of carbon-14 and if it were two half-lives old (about 11,460 years), I'd find about a quarter or .25 time this level. Since the calculated age comes out between these two ages, and the given carbon-14 level is between .25 and .5, my answer at least seems reasonable.

- ii) Estimate the age of a papyrus manuscript which is found to contain .822 times the level of carbon-14 in present in living matter.
- iii) Estimate the age of a bone which is found to contain .006 times the level of carbon-14 in present in living matter.

How much does our uncertainty about the exact half-life of carbon-14 affect these answers? In the next problem, we'll try to get a feel for this uncertainty.

PROBLEM 5.5.29: Recalculate the annualized rate of decay of carbon-14 assuming that



5.5 Applications of compounding

- i) the half life is 5,690 years;
- ii) the half-life is 5,770 years.
- iii) Recalculate the the age of the sandal in [PROBLEM 5.5.28 b\)](#) using each of the rates found in i) and ii). Based on these two ages, how accurate can we say the dating of the sandal is likely to be?

A second potential source of error is the accuracy with which we can measure carbon-14 levels. This limits how far back carbon-14 dating can take us: the next problem indicates why.

PROBLEM 5.5.30: Let's suppose that our equipment for measuring carbon-14 levels is accurate to ± 0.003 of the level of carbon-14 in present in living matter.

- i) The sandal in part i) of [PROBLEM 5.5.28](#) might actually contain between .342 and .348 times the level of carbon-14 in present in living matter. Show, by reestimating the age of the sandal based on these two levels, that the potential error in the measurement of the carbon-14 level in the sandal introduces an uncertainty of about ± 60 years in the dating of the sandal.
- ii) The bone in part iii) of [PROBLEM 5.5.28](#), which dated to about 42,300 years old accurate to three digits, might actually contain between .003 and .009 times the level of carbon-14 in present in living matter. Show, by reestimating the age of the bone based on these two levels, that the bone could be anywhere from 38,900 to 48,000 years old.

The moral here is that, while carbon dating is likely quite accurate for an object which is 10,000 years old, it becomes a very rough estimate for an object which is 40,000 years old. It's not hard to convince yourself that, beyond this point, carbon-14 pretty much breaks down. All we can say about objects over 50,000 years old is just that: they are over 50,000 years old. The carbon-14 levels of all such objects will measure as .000 to three places regardless of how old they actually are.

5.5 Applications of compounding

How then do geologists measure the age of rocks which can be *billions* of years old? The answer is to measure isotopes which have longer half-lives: potassium-40 is one favorite. The method is actually called argon dating after the stable end product. There's no standard baseline level of potassium-40, so instead, geologists measure the amounts of argon-40 *and* potassium-40 now and use a two-stage process. In the first stage, they then assume that all the argon is a product of potassium decay and use the ratio of the two amounts to deduce how much potassium-40 was in the rock when it was formed. In the second stage, they compare the formation and current potassium-40 levels to date the rock. The next exercise gives an idea of the way this second step works at geological time scales, leaving out the Why don't archeologists use argon dating too? See if you can figure out the answer before you do the exercise. I'll discuss it after.

PROBLEM 5.5.31: Potassium 40 has a half life of 1.25^9 years. In problem [PROBLEM 5.5.27](#), you should have found that its decay rate is $-5.55 \times 10^{-8}\%$ a year.

i) Find the fraction of the baseline level of potassium-40 which would remain in a rock after

a. 500,000,000 years.

Solution

We just use the [GENERAL CONTINUOUS APPROXIMATION 5.5.6](#) $S_y = Be^{(0.01r \cdot y)}$. Since we are interested in what fraction of the baseline level remains we can take $B = 1$. (Why?) Using the value of r above and $y = 500,000,000$, we find that $S_y = 0.758$ to three places. That is, a bit more than three-quarters of the potassium-40 remains.

b. 520,000,000 years.

ii) Find the fraction of the baseline level of potassium-40 which would remain in a bone after

a. 5,000 years.

b. 50,000 years.



5.5 Applications of compounding

Here what we see is that the residual potassium-40 levels of rocks 500,000,000 and 520,000,000 years old differ by many thousandths of the baseline level (.758 versus .749) while those of a bones 5,000 and 50,000 years old are the same to 4 places (both .9999). So if our ability to measure these levels is limited—say to 3 places—we can usefully easily distinguish the ages of the rocks but not those of the bones.

PROJECT 5.5.32: The explanation I have given above of how carbon-14 dating is performed bears no resemblance to what people who use this method actually do. One reason is that carbon-14 levels in the atmosphere actually rise and fall somewhat due to variations in solar activity (most of the carbon-14 in the atmosphere is produced by solar radiation). Not to mention that nuclear testing in the 1950's and 1960's almost *doubled* the levels of carbon-14 in the atmosphere. The variance is sufficiently great to introduce substantial errors in dates computed by our plug-and-chug method. But scientists interested in these dates don't just give up. They go to work. Find out how accurate dating using carbon-14 levels is still possible even though our basic assumption (that there is a constant baseline level in living matter) is wrong. The answer, based on sampling thousands of tree rings, is an excellent illustration of a statement from the opening discussion: “The hard part of mathematical modelling is not usually the math but the analysis of how far the model corresponds to reality and what its uses and limitations are and the work needed to overcome these limitations”. A good starting point is the site of the [Rafter Laboratory](#).

Recently, a brilliant set of experiments has turned the problem created by atmospheric nuclear testing around and used this very fact to overcome ethical restrictions on the study of human subjects and show how cells are renewed in the human heart. After the nuclear test ban treaty of 1963, the flow of “extra” carbon-14 into the atmosphere stopped and levels began to decay back towards the long

5.6 Amortization: saving

term baseline (due to solar irradiation). This means that there are slight, but detectable differences in the carbon-14 levels in cells in *living* bodies—like yours—depending on the year in which the body created the cell. Cells created in 1963 have the highest levels of carbon-14, and the level gradually decreases as the birth year of the cell gets closer and closer to the present. By studying carbon-14 levels in many cells in a single human heart, a team led by Dr Jonas Frisén of the Karolinska Institute in Stockholm has established that, contrary to general belief, the heart can and does regenerate cells, as had been conjectured by Dr Piero Anversa in the late 1980s. In a typical lifetime about half the cells in a human heart are renewed. You can read more about this work in [this excellent article](#).

5.6 Amortization: saving

It's now time to get my daughter's college fund together. I want to have \$120,000 available for her education when she turns 18. Our earlier discussions in [EXAMPLE 5.2.14](#) and [EXAMPLE 5.4.32](#) had two morals. The first is that the sooner I start planning for this the better. We'll come back to this point later. The second is that I'll never be able to reach this goal by investing a single lump sum: the amount I'd need is just out of my reach even if I make the investment when she's born. Almost everyone faces a number of similar problems in their life: you know you are going to need a sum of money many times your annual salary all at once. The most universal examples are the money you need to buy a house (or even a car) and the money you will live on when you retire.

How do people ever assemble these sums of money? You probably already see the answer: if the sum is too big to put together all at once the only solution is to put it together a bit at a time. People being creatures of habit, the only way most of us can be sure of

5.6 Amortization: saving

continuing to put aside the small sums needed is to regularly put aside a fixed sum. This process is called *amortization*. The goals of this section are to understand what the small sums will amount to over time and to draw some conclusions about planning for such major needs.

The future amortization formula

First let's define our terms.

AMORTIZATION 5.6.1: *A sum of money which is assembled by making a series of equal deposits at regular intervals into an account which earns a fixed interest rate, or a loan at fixed interest which is repaid by making a series of equal payments at regular intervals is said to be amortized.*

In this section, we'll consider amortizations which involve **saving**: we make regular deposits with the goal of having a lump sum of money available *after* the payments are made. **Loans**, where we receive the lump sum *before* making a series of payments will be discussed in [SECTION 5.8](#).

We'll continue to use many terms from the preceding section to describe such series of payments. For example, the *period* of the amortization will be the length of time between deposits or payments and we'll again use m to denote the number of periods in a year. The number of years in the *term* of the amortization will again be denoted y and the number of periods (or deposits or payments) will be T : these are still related by the [TERM CONVERSION FORMULA 5.1.13](#), $T = m \cdot y$. The fixed nominal interest rate will be r and the corresponding periodic rate will be $p = \frac{0.01 \cdot r}{m}$ by the [INTEREST RATE CONVERSION FORMULA 5.1.10](#).

SIMPLIFYING ASSUMPTION 5.6.2: In all the problems in this section, we will assume that the compounding frequency of the account

into which deposits are made or of the loan against which payments are made is equal to the frequency with which the deposits or payments are made and that deposits or payments are always made at the end of each compounding period.

The assumption that compounding frequency and deposit frequency are equal is definitely not true in many everyday amortizations and you wouldn't even want it to be. For example, you want to have the amount in your retirement account compounded daily (because you get more interest this way) even if you only make deposits into it once a month on payday. What you don't want is to have to learn the very complicated formulas which are needed to compute the amounts in such an account. (Of course, in some professions, these formulas are critical and there are entire courses devoted to them). Our assumption will allow us to work only with very simple formulas. What's more these simple formulas give answers close to, if not quite equal to, those from the more complicated ones so we can draw the important conclusions everyone should know about amortizations from them. In a similar vein, there are lots of amortizations in which the deposits are made at the start rather than the end of each period. The corresponding formulas are only a touch more complicated than the ones we'll use but they are significantly harder to remember and it's easy to confuse the two. We'll leave them to the professionals too.

We will use two new terms for the regular amounts involved:

DEPOSIT 5.6.3: *The common amount of each deposit will be denoted D although we'll use the term **payment** when a loan is being amortized.*

SUM AND BALANCE 5.6.4: *The lump sum of money being assembled in a **savings** account will be denoted S and be called the **sum** of the account or loan. The intermediate sum after the end of the i^{th} period—that is, the amount which has accumulated in the account in the first i deposits will be denoted by S_i . In particular, when a sum is being*

5.6 Amortization: saving

assembled we will have $S = S_T$, the final intermediate sum after all T periods.

*The lump sum of money being repaid in a loan will be denoted B and be called the **balance** of the loan. The intermediate balance after the end of the i^{th} period—that is, the balance outstanding on a loan after the first i payments will be denoted by B_i . In particular, when a loan is being repaid, we will have $B = B_0$ the initial intermediate balance before the start of the first period.*

As I hope you'll have guessed from the uses of these letters in earlier sections, we use S for the intermediate sums in a savings account because these are all basically future values and use B for the intermediate balances on a loan because these are all basically present values.

For the rest of this section, we discuss only savings. To start, let's get back to my daughter's college fund. What I'd like to do is open an account which pays a fixed rate of interest—let's say 3.9%—and then make a small deposit into the account at the end of every month from my daughter's birth until she turns 18 at which time I'd like to have \$120,000.00 in the account. Since I am making monthly payments, $m = 12$, and since I make them for 18 years, $y = 18$. Thus the term $T = m \cdot y = 12 \cdot 18 = 216$. Further since $r = 3.9\%$, the periodic rate $p = \frac{0.01 \cdot r}{m} = \frac{0.01 \cdot 3.9}{12}$. The final sum $S = S_{216}$ I want to reach is \$120,000.00 The only thing I don't know is how big a deposit D to make every month.

Let's just leave this as an unknown for now and try to understand how the money in the account builds up. To do so, I'll use S_i to stand for the sum at the end of i months. Let's compute the first few intermediate sums S_i (remember this is the amount in the account at the end of i months) and see whether we can spot a pattern.

After one month, all that's in the account the first deposit D which I have just made so

$$S_1 = D.$$



5.6 Amortization: saving

That was easy. The second month isn't much harder. The [ONE PERIOD PRINCIPLE 5.2.3](#) says that adding the interest earned in the second month to the sum S_1 just multiplies it by the magic factor $(1+p)$ to $S_1(1+p)$ to which we add the second deposit D to get the second month's sum

$$S_2 = D + (1+p)S_1 = D + (1+p)D = D(1 + (1+p)) .$$

(Wondering why I factored out the D and then did *not* rewrite the $(1 + (1+p))$ as $2+p$? Hindsight! This helps makes the final pattern easier to spot as you'll see in a moment.)

The third month isn't much harder either. Again, the [ONE PERIOD PRINCIPLE 5.2.3](#) says that adding the interest earned in the third month to the sum S_2 just multiplies it too by the magic factor $(1+p)$ to $S_2(1+p)$ to which we add the third deposit D to get the third month's sum

$$\begin{aligned} S_3 &= D + (1+p)S_2 \\ &= D + (1+p)D(1 + (1+p)) \\ &= D(1 + (1+p) + (1+p)^2) . \end{aligned}$$

It's not too hard to see the pattern which is emerging here. Passing from one month's sum to the next is always the same. Suppose that our sum after v months is S_v . (Why I have called the variable v will be clear in a moment.) Then the [ONE PERIOD PRINCIPLE 5.2.3](#) says that adding the interest earned in the next— $(v+1)^{\text{st}}$ —month to the sum S_v just multiplies it by the magic factor $(1+p)$ giving $S_v(1+p)$ to which we add the next or $(v+1)^{\text{st}}$ deposit D to get the sum S_{v+1} at the end of $v+1$ months:

$$S_{v+1} = D + S_v(1+p) .$$

Moreover, it's easy to see that the final expanded formula for S_v to which this leads is

$$S_v = D(1 + (1+p) + (1+p)^2 + \cdots + (1+p)^{v-1}) .$$

5.6 Amortization: saving

All the formulas begin with a common factor of D which multiplies a sum in which all but the first two terms are powers of the magic factor $(1 + p)$ and the highest exponent which appears is 1 less than the number i of months we are working with. We can actually think of *all* the terms as powers of $(1 + p)$ by remembering that for any positive base x , we have $x^0 = 1$ and $x^1 = x$. Taking $x = (1 + p)$, this lets us write the last formula as:

FIRST SUM FORMULA 5.6.5:

$$S_v = D \left((1 + p)^0 + (1 + p)^1 + (1 + p)^2 + \cdots + (1 + p)^{v-1} \right).$$

We can check our guesses for both patterns by plugging the [FIRST SUM FORMULA 5.6.5](#) into the equation relating S_v and S_{v+1} above:

$$\begin{aligned} S_{v+1} &= D + (1 + p)S_v \\ &= D \cdot 1 + (1 + p)D \left((1 + p)^0 + (1 + p)^1 + \cdots + (1 + p)^{v-1} \right) \\ &= D(1 + p)^0 + D \left((1 + p)^1 + (1 + p)^2 + \cdots + (1 + p)^v \right) \\ &= D \left((1 + p)^0 + (1 + p)^1 + (1 + p)^2 + \cdots + (1 + p)^v \right) \end{aligned}$$

Here the highest power of $(1 + p)$ is v which is again one less than the number of months in S_{v+1} so we have verified the prediction of the [FIRST SUM FORMULA 5.6.5](#).

SECOND APPROACH 5.6.6: There's another way to think about the [FIRST SUM FORMULA 5.6.5](#) which sometimes comes in handy. Having a second way to think about any problem is never a bad thing because often something which looks difficult or messy from one point of view becomes very simple when we use the other.

If we track that very first deposit through the calculation it shows up in the sum for S_v as the term $D(1 + p)^{v-1}$. This term is what the [COMPOUND INTEREST FORMULA 5.2.4](#) would give for the future value after $v - 1$ periods of an amount $A_0 = D$ at a periodic interest rate of p . A moment's thought reveals that this is not an accident: that first payment was made at the end of the first month so at the end

5.6 Amortization: saving

of the v^{th} month it's been earning interest for $v - 1$ months and has accrued to $D(1 + p)^{v-1}$.

The same is true for all the other terms in the sum: the term $D(1 + p)^{v-2}$ corresponds to the future value of the second deposit which has earned $v - 2$ months of interest, and so on down to the terms $D(1 + p)^1$ and $D(1 + p)^0$ which correspond to the last two deposits which have earned 1 month and 0 months of interest respectively. In other words, we can reach the formula *either* by computing successive sums as we did above *or* by tracking deposits individually to the end of some period and then adding.

Great! I have two ways to reach a formula for the final sum S in my daughter's college fund because $S = S_{216}$, the intermediate sum after 216 months. There's just one problem with this formula as it stands. The sum of powers of $(1 + p)$ in the formula for S_v has v terms. This means that to use it to figure out what will my sum will be after 216 months, I'd need to compute and sum 216 powers of $(1 + p)$. No thanks. I might not finish before my daughter's 18th birthday.

What we have here is a classic, messy summation like those discussed in [Section 1.3](#). What we need, therefore, to eliminate the need to compute and total all 216 powers is a **closed form** formula for this kind of sum. Although, at first glance, it doesn't look like it, we already have such a formula in hand. The [FIRST SUM FORMULA 5.6.5](#) is a **geometric summation**—each term is the constant D times a power of the magic factor $(1 + p)$ —so we can apply the [GEOMETRIC SUMMATION FORMULA 1.3.3](#).

All we need to do is match up the elements in the [FIRST SUM FORMULA 5.6.5](#) to those in the [GEOMETRIC SUMMATION FORMULA 1.3.3](#). if we substitute $r = (1 + p)$, $u = v - 1$. (See why I used the v ? I knew that the months variable in our formula was not quite going to match up with the [GEOMETRIC SUMMATION FORMULA 1.3.3](#) and wanted to avoid any confusion.) Therefore, we can replace the sum in the [FIRST SUM FORMULA 5.6.5](#) by what we get on the right hand

5.6 Amortization: saving

side of the [GEOMETRIC SUMMATION FORMULA 1.3.3](#) when we make the same substitutions. Of course, if $u = v - 1$ then $u + 1 = v$. With $r = (1 + p)$, we find that $(1 - r^{u+1}) = (1 - (1 + p)^v)$ and $(1 - r) = (1 - (1 + p)) = -p$. Combining all this, we get

SECOND SUM FORMULA 5.6.7: $S_v = D \left(\frac{1 - (1+p)^v}{-p} \right) = D \left(\frac{(1+p)^v - 1}{p} \right).$

By setting $S = S_{216}$, this gives me a formula relating the final sum S in my daughter's college fund to the deposit D that I make and the periodic interest rate p . But there's really nothing special about this example except that we fixed the term T to be 216 periods. So, let's write down the general formula before seeing what it implies for my daughter's college fund.

FUTURE AMORTIZATION FORMULA 5.6.8: *If a deposit is made at the end of each of T periods into an account which earns compound interest at a periodic rate p , then the final sum S in the account at the end of the T^{th} period and the amount D of each deposit are related by*

$$S = D \left(\frac{(1+p)^T - 1}{p} \right) \quad \text{and} \quad D = S \left(\frac{p}{(1+p)^T - 1} \right).$$

Working with the future amortization formula

The [FUTURE AMORTIZATION FORMULA 5.6.8](#) is amazingly simple once we've put the geometric series formula to work. This means that, although we had to huff and puff a fair bit to get to them, applying it is a piece of cake.

EXAMPLE 5.6.9: We did all the work above: if my account earns 3.9% for 18 years compounded monthly, then $p = \frac{0.01 \cdot r}{m} = \frac{0.01 \cdot 3.9}{12}$ and $T = 216$. If I want to have a final sum S of \$120,000.00 then the equation for D tells me I need to deposit $D = S \left(\frac{p}{(1+p)^T - 1} \right)$. Plugging in we find

$$D = \$120,000.00 \left(\frac{\left(\frac{0.01 \cdot 3.9}{12} \right)}{\left(1 + \left(\frac{0.01 \cdot 3.9}{12} \right) \right)^{216} - 1} \right) = \$384.05$$

5.6 Amortization: saving

every month. That's a big expense but it is still a sum I can think about including in my monthly budget (especially as such deposits generally are not taxed).

We'll come back to consider other ways to use the [FUTURE AMORTIZATION FORMULA 5.6.8](#) in a moment but first let's formalize what we did above with a method. As always, the first two steps are the same (find m and use it to get p and T) and the third just involves plugging values into the formula.

We'll come back to consider other ways to use the [FUTURE AMORTIZATION FORMULA 5.6.8](#) in a moment but first let's formalize what we did above with a method. As always, the first two steps are the same (find m and use it to get p and T) and the third just involves plugging values into the formula.

METHOD FOR FUTURE/SAVINGS AMORTIZATIONS 5.6.10:

Step 1: Determine the periods in the problem (that is, the units in which the term is measured) and the value of m , the number of periods per year.

Step 2: Use the [INTEREST RATE CONVERSION FORMULA 5.1.10](#) to find the periodic interest rate p from the nominal interest rate r and the [TERM CONVERSION FORMULA 5.1.13](#) to find the term T in periods from the term in years y .

Step 3: Apply the appropriate [FUTURE AMORTIZATION FORMULA 5.6.8](#) to find whichever of the the deposit D and the sum S is to be determined.

Here are a few exercises for you to try which involve what are generally called *sinking funds*. These are accounts into which businesses make regular deposits which accumulate towards the purchase of some high-ticket item. The purpose, as with my daughter's college fund, is to spread out the corresponding expense over the term of the fund and avoid having a large charge on the books in any accounting period. I have worked a few of the examples.



5.6 Amortization: saving

PROBLEM 5.6.11: An insurance company wants to make monthly deposits into an account which interest compounded monthly to fund the purchase of a computer which will cost \$135,000.00 How much should each deposit be if:

- i) the account has a nominal rate of 5% and the payments are made over 5 years?

Solution

Step 1: The periods are months so $m = 12$.

Step 2: $p = \frac{0.01 \cdot r}{m} = \frac{0.01 \cdot 5}{12}$ and $T = my = 12 \cdot 5 = 60$.

Step 3: We know the sum $S = \$135,000.00$ so using the formula

$D = S \left(\frac{p}{(1+p)^T - 1} \right)$ we find the deposit is

$$D = \$135,000.00 \left(\frac{\left(\frac{0.01 \cdot 5}{12} \right)}{\left(1 + \left(\frac{0.01 \cdot 5}{12} \right) \right)^{60} - 1} \right) = \$1,985.12.$$

- ii) the account has a nominal rate of 8% and the payments are made over 3 years?

- iii) the account has a nominal rate of 3% and the payments are made over 5 years?

PROBLEM 5.6.12: A University is planning to make annual deposits of \$10,000.00 into a sinking fund on which interest is compounded annually to fund the purchase of a statue to commemorate a distinguished faculty member. How much can they afford to pay for the statue if:

- i) the account has a nominal rate of 4.8% and the payments are made over 4 years?

Solution

Step 1: The periods are years so $m = 1$.

Step 2: $p = \frac{0.01 \cdot r}{m} = \frac{0.01 \cdot 4.8}{1}$ and $T = my = 1 \cdot 4 = 4$.

Step 3: Here we know the deposit amount $D = \$10,000.00$ so we plug into the formula $S = D \left(\frac{(1+p)^T - 1}{p} \right)$ to find that the final

sum is

$$S = 10,000.00 \left(\frac{\left(1 + \left(\frac{0.01 \cdot 4.8}{1}\right)\right)^4 - 1}{\left(\frac{0.01 \cdot 4.8}{1}\right)} \right) = \$42,973.27$$

- ii) the account has a nominal rate of 3.2% and the payments are made over 3 years?
- iii) the account has a nominal rate of 6.6% and the payments are made over 7 years?

How can we check such amortization calculations? Basically, both the [SIMPLE INTEREST APPROXIMATION 5.3.3](#) and the [CONTINUOUS APPROXIMATION 5.3.11](#) can be souped up for use in checking amortizations. The latter is actually more straightforward. We simply approximate the exponential $(1 + p)^T$ in the [FUTURE AMORTIZATION FORMULA 5.6.8](#) by the slightly larger exponential $e^{(p \cdot T)} = e^{(0.01r \cdot y)}$ from [CONTINUOUS APPROXIMATION 5.3.11](#). In the formula for D where this appears in the denominator and we are now dividing by a larger quantity, we get an approximation slightly smaller than the exact value.

FUTURE AMORTIZATION-CONTINUOUS APPROXIMATION 5.6.13:

The final savings S is a bit less than $D \left(\frac{e^{(0.01r \cdot y)} - 1}{p} \right)$.

This check lacks one feature of the [CONTINUOUS APPROXIMATION 5.3.11](#): there's still a periodic rate p in each formula. If you forgot to convert the nominal rate when using the [FUTURE AMORTIZATION FORMULA 5.6.8](#), you'll probably use r for p here too. Fortunately, the different numerator will lead to a different answer and let you catch your mistake. You may also notice that I haven't colored it: this is one formula which you don't really need to learn. You can make the necessary approximations if you just remember to use the [CONTINUOUS APPROXIMATION 5.3.11](#) to replace the $(1 + p)^T$.

EXAMPLE 5.6.14: Let's check the calculation in [EXAMPLE 5.6.9](#). Here we had $r = 3.9\%$ and $y = 18$ years and since we were compounding

5.6 Amortization: saving

monthly $m = 12$ and $p = \frac{0.01 \cdot 3.9}{12}$. Our D was \$384.05 so the sum S of \$120,000.00 should be a bit smaller than

$$D \left(\frac{e^{(0.01 \cdot 3.9)} - 1}{p} \right) = 384.05 \left(\frac{e^{(0.013 \cdot 3.9 \cdot 18)} - 1}{\left(\frac{0.01 \cdot 3.9}{12} \right)} \right) = \$120,270.78.$$

and it is. You can check a sum calculation the same way.

PROBLEM 5.6.15: Use the continuous approximation to check your answers to problems [PROBLEM 5.6.11](#) and [PROBLEM 5.6.12](#).

As with [SIMPLE INTEREST APPROXIMATION 5.3.3](#), the simple interest approximation to an amortization calculation is both better and worse. It's better because the calculation is easier—you can often do it in your head—and experience shows that we're much likelier to make an easy check than a hard one—but worse because it's less accurate. (It's like the difference between a cheap point-and-shoot camera you can put in your pocket and reflex camera with lots of lenses. The reflex camera takes much better pictures but that's not an advantage if you leave it in the hotel room because the case is so heavy.)

The idea is very simple. First add up all the deposits getting an amount $A = T \cdot D$, then add some simple interest to this. The only question is how much simple interest to add. All the deposits earn simple interest for differing numbers of periods so we really ought to add a different amount of interest to each. The problem with doing so is that you get a formula more complicated than the one you are trying to check. The solution is to group pairs of deposits moving inwards from the start and end of the term. The first deposit earns $T - 1$ periods interest and the last 0 periods. Together the two earn a total of $T - 1$ periods interest and an average of $\frac{T-1}{2}$ periods. The second deposit earns $T - 2$ periods interest and the second last 1 period of interest. Again, the total is $T - 1$ periods and the average is $\frac{T-1}{2}$ periods. The third deposit earns $T - 3$ periods interest and the third last 2 periods of interest. Again, the total is $T - 1$ periods and the average is $\frac{T-1}{2}$ periods.

5.6 Amortization: saving

So the pattern is that an average deposit earns $\frac{T-1}{2}$ periods of interest. To keep things simply, we replace this with $\frac{T}{2}$ periods which is an average term of $\frac{y}{2}$ years. Now using years as periods, the periodic rate is $0.01r$ so the [SIMPLE INTEREST FORMULA 5.1.6](#) says the interest earned should be roughly $0.01r \cdot A \cdot \frac{y}{2}$. If we then add the amount A , we should get a total $A + A0.01r\frac{y}{2} = A(1 + 0.01r\frac{y}{2})$ somewhat smaller than the actual final sum S —smaller because we are ignoring the effect of compounding. Note that I didn't collect the two fractions because I think it's easier to remember the formula when we think of them separately. In fact, this is another formula where it's better to remember the idea—an average deposit earns interest for half the term—than the formula. The example following the formula illustrates this.

As with the [SIMPLE INTEREST APPROXIMATION 5.3.3](#), this approximation only gives reasonable accuracy for short terms. When the term is longer, use the [FUTURE AMORTIZATION-CONTINUOUS APPROXIMATION 5.6.13](#) instead.

FUTURE AMORTIZATION-SIMPLE INTEREST APPROXIMATION 5.6.16:
 S is larger—possibly quite a bit larger—than $T \cdot D \left(1 + 0.01r\frac{y}{2}\right)$.

EXAMPLE 5.6.17: Let's recheck the deposit $D = \$384.05$ I calculated in [EXAMPLE 5.6.9](#). Here we had $T = 216$, $r = 3.9\%$ and $y = 18$ years so so $0.01r = 0.039$ and $\frac{y}{2} = 9$. Our final sum S should be somewhat greater than

$$T \cdot D \left(1 + 0.01r\frac{y}{2}\right) = 216 \cdot 384.05(1 + 0.039 \cdot 9) = \$112,071.93$$

and it is. As we learned to expect in [SECTION 5.3](#), this approximation is cruder than the previous one. However, by making it a bit cruder still, we could use it without getting our calculator out. In a class, I'd say something like this. "Well, 216 times 384.05 is about 200 times 400 or 80,000. And, 3.9% times half of 18 is about 4 times 9 or about 36%—lets say 40%. So add 40% of \$80,000 which is \$32,000 to get \$112,000. That's a bit less than the sum of \$120,000.00 so

5.6 Amortization: saving

I'm happy." The point to notice is that I was happy to replace 216 by 200 or 3.9 by 4 or 36 by 40 to simplify the arithmetic. I know this approximation isn't going to give me a lot of decimals of accuracy anyway. I just want to make sure that I didn't do something stupid—which I do all the time, just a bit less often than most of you, I hope. If I did, my approximation and my answer would probably be off by a factor of 2 or 200, I'd be worried and I'd go look for my mistake; here they're within 10% of each other which is reasonable agreement.

PROBLEM 5.6.18: Use the simple interest approximation to check your answers to [PROBLEM 5.6.11](#) and [PROBLEM 5.6.12](#).

PROBLEM 5.6.19: You make monthly deposits of \$200.00 into a retirement account.

- i) How much will you have in the account at the end of 3 years if the account earns interest of
 - a. 2%?
 - b. 6%?
 - c. 10%?
- ii) How much will you have in the account at the end of 30 years if the account earns interest of
 - a. 2%?
 - b. 6%?
 - c. 10%?
- iii) Use the [FUTURE AMORTIZATION-SIMPLE INTEREST APPROXIMATION 5.6.16](#) and [FUTURE AMORTIZATION-CONTINUOUS APPROXIMATION 5.6.13](#) to check your answers to i)b and ii)b and compare the accuracy they give. Why is the agreement so much better in i)b than in ii)b?

Any smokers out there? Like to quit? Here's some motivation.

PROBLEM 5.6.20: Let's suppose that cigarettes cost \$6.00 a pack and that you smoke a pack a day. We'll call this a monthly expense of \$180.00 on cigarettes. Suppose you are 21 now and that you quit

5.7 Planning for future needs

smoking and put that \$180.00 a month into a retirement account which invests in stocks and which yields 8% every month. If you continue investing your cigarette money until you are 65, show that you will have over \$870,000.00 in the account! Exactly how much will you have?

5.7 Planning for future needs

Let's look at other options for my daughter's college fund. Since I have an 18 year period over which to build up the account, I can afford to consider investments which have higher risks of short term losses but offer higher average returns over the long term. Suppose I want to test the possibility of investing in bonds and I think these will give me an annualized yield of 6%. Since my formula only applies when my deposit period and compounding period match, I'll consider making 18 annual deposits. (I could consider monthly deposits too but why do the extra arithmetic when I'm only trying to test possibilities.) In this case, $m = 1$ so $p = \frac{0.01 \cdot r}{m} = \frac{0.01 \cdot 6}{1}$ and $T = m \cdot y = 1 \cdot 18 = 18$ and I find that

$$D = S \left(\frac{p}{(1+p)^T - 1} \right) = \$120,000.00 \left(\frac{\left(\frac{0.01 \cdot 6}{1} \right)}{\left(1 + \left(\frac{0.01 \cdot 6}{1} \right) \right)^{18} - 1} \right) = \$3,882.78$$

or about \$325 a month. If I thought that by investing in stocks, I could earn an 8% yield, I could calculate that I'd need to contribute

$$D = S \left(\frac{p}{(1+p)^T - 1} \right) = \$120,000.00 \left(\frac{\left(\frac{0.01 \cdot 6}{1} \right)}{\left(1 + \left(\frac{0.01 \cdot 6}{1} \right) \right)^{18} - 1} \right) = \$3,204.25$$

or about \$265 a month.

There are several things "wrong" with the last two calculations but the conclusions they lead to itself are still basically correct. Let's

5.7 Planning for future needs

discuss the bond calculation. First, even if the bonds I bought were compounded annually, I'd probably want to make monthly deposits rather than trying to come up with a sum close to \$4,000 once a year. This takes us outside our assumption of equal deposit and compounding periods. On the other hand, the difference is small and if I did make monthly deposits in a bank account to make up the annual sum, they'd earn a bit of interest during the year so I'd be that small amount further ahead.

PROBLEM 5.7.1:

- i) Suppose you wish to save \$120,000.00 by making *monthly* deposits into an account which earns 6% compounded monthly for 18 years. How much must you deposit each month? Show that your total deposits over a one year period will be slightly less than the \$3,882.78 computed above.
- ii) Suppose you wish to save \$120,000.00 by making *monthly* deposits into an account which earns 8% compounded monthly for 18 years. How much must you deposit each month? Show that your total deposits over a one year period will be slightly less than the \$3,204.25 computed above.
- iii) What approximation would give the best check of your two preceding answers? Perform this check.

Second, and more importantly, I would not be able to purchase bonds in such uneven amounts and I would probably not want to even if I could. Instead, I'd probably want to make regular deposits to a bond mutual fund: this would not only let me deposit whatever sum I chose each month, but would allow me to own shares of a range of bonds. Having small amounts of many assets is called *diversification* and makes your investments somewhat less prone to disaster: if I own one kind of bond and the issuer goes bankrupt, I suddenly earn no interest and possibly lose all my principal. If I hold many different bonds, as I do through shares in a bond mutual fund,

5.7 Planning for future needs

then the impact of any given issuer going bankrupt would be limited to a fraction of my daughter's college fund.

On the other hand, bond mutual funds do not promise fixed yields. They are continually buying and selling bonds with different yields for one thing. For another, the prices of bonds themselves are not fixed but go up and down in response to various demand factors. So a yield figure like 6% is at best an estimate based on historical yields of such investments. The bottom line is that the assumption which runs through this entire chapter—that interest rates or yields are constant—does not apply. Once again, however, as long as I feel that the actual rates I am likely to earn are *close* to the 6% figure, then my calculation gives me a good *estimate* of what I need to be putting into the college fund each month.

Both the same objections apply even more forcefully to the stock purchase scenario. Stocks do not really compound at all. Their prices fluctuate in response to market forces in ways which are very violent in the short term—often the value changes by 1 or 2 percent in a single day and the fall of 2008 swings of 6 – –10% in a day were not uncommon. If that does not seem like much to you, the problem below will give you some idea how wild such swings are.

PROBLEM 5.7.2: Show that the annualized yield of an investment that increases in value by 1% a day is about 3,678%. Show that a \$1,000.00 investment which decreases in value by 1% a day will be worth \$25.52 at the end of 1 year.

However, over longer terms the ups and downs of stock prices largely cancel. It's rare that prices rise or fall by more than 25% in a year. In most years, average prices rise but there are many years when they fall. When we speak of expecting an 8% return on stocks, we need to have in mind a very long term—at least a decade and better two. My daughter's college fund with its 18 year term is a reasonable example. Even so the assumptions that underlie the [FUTURE AMORTIZATION FORMULA 5.6.8](#) only apply very roughly to investing

5.7 Planning for future needs

in stocks. The calculation provides a useful estimate of what I am *likely* to have in the fund when my daughter is 18 but it's far from the the accurate estimate I get for a bank account.

However, even with a bank the same basic objection applies. I can lock in a fixed rate over a future term with my bank for a CD that I purchase today. However, next month when I go to purchase another I will almost certainly be offered a different rate—slightly higher or lower. Over longer periods, the changes in these rates can be large: in 1982, savings bank CD rates peaked at close to 14% while in 2008 they were down below 2%. That's why the S.E.C. makes them say "Past history is no guarantee of future performance" in the small print at the bottom of the ads.

MORAL 5.7.3: *When discussing the future performance of investments, you can only estimate and even the wisest estimates can only reveal what is likely to happen not what is going to happen.*

PROJECT 5.7.4: Look up the year-end levels of the Dow Jones Stock Index over the past century and calculate the annualized yields of this index over various the 18 periods. What 18 year period had the lowest yield? What 18 year period had the highest yield? What has the annualized yield been over the past century? What conclusions do you draw from your calculations?

Is there any point in trying to make mathematical calculations in the face of this kind of uncertainty? Absolutely. The estimates we get from a formula like the [FUTURE AMORTIZATION FORMULA 5.6.8](#) with all its limitations are much better than no estimates at all. After all, I do need to decide somehow what to do about my daughter's college fund, how to plan for my retirement, whether I can afford that new home Our formulas give us a good handle on these problems: they tell us what's likely to happen. It's just a handle we should use with caution aware that it might come unstuck. One prudent strategy is to ask "What's the worst that can happen?", and to try to devise a fall-back strategy to handle such worst cases.

5.7 Planning for future needs

Suppose I decide that I am willing to put my daughter's college fund in the stock market and to assume an 8% average yield which will let me get away with putting away about \$265.00 a month or \$3,180.00 a year. Suppose the worst average yield I can imagine over the 18 years is only 2%. What will the final sum be then?

$$S = D \left(\frac{(1+p)^T - 1}{p} \right) = \$3,180.00 \left(\frac{\left(1 + \left(\frac{0.01 \cdot 2}{1}\right)\right)^{18} - 1}{\left(\frac{0.01 \cdot 2}{1}\right)} \right) = \$68,091.15.$$

I'll have just over half what I'll need. If that happens, she'll have to contribute too. Maybe she'll get a scholarship. If not she'll have to work part-time. If worst comes to worst, I might have to take out a second mortgage on my house to help with her junior and senior years.

Suppose I can imagine a negative average yield of -2% ?

$$S = D \left(\frac{(1+p)^T - 1}{p} \right) = \$3,180.00 \left(\frac{\left(1 + \left(\frac{0.01 \cdot -2}{1}\right)\right)^{18} - 1}{\left(\frac{0.01 \cdot -2}{1}\right)} \right) = \$48,473.48.$$

Now I am \$70,000.00 short: if this happens, she'll have to live at home and go to a state school. If that's a thought I just can't live with then I better make larger monthly deposits. Either the worst case will come to pass and then I'll be very unhappy that my investment has turned out so poorly but very glad because my daughter can still go to the school of her choice, or, I'll do better than I feared and then I'll have a lot more than \$120,000.00 in my daughter's college fund and I can use the extra to buy that boat (or fulfill whatever other dream I might have).

Here are some problems dealing with retirement funds that will let you play around with using the [FUTURE AMORTIZATION FORMULA 5.6.8](#) in this way and with the [FUTURE AMORTIZATION-SIMPLE INTEREST APPROXIMATION 5.6.16](#) and [FUTURE AMORTIZATION-CONTINUOUS APPROXIMATION 5.6.13](#).

5.7 Planning for future needs

PROBLEM 5.7.5: Suppose that you are self-employed and wish to have a fund of \$500,000.00 available when you retire at age 65. To reach this goal, you plan to make a series of equal monthly deposits into a Self-Employed Pension or SEP.

i) First, let's suppose that you are now 25 years old. Determine how much you need to deposit each month for the next 40 years to reach your goal if the SEP account earns interest at

- a. 3%.
- b. 6%.
- c. 9%.

ii) Next, let's suppose that you are now 45 years old. Determine how much you need to deposit each month for the next 20 years to reach your goal if the SEP account earns interest at

- a. 3%.
- b. 6%.
- c. 9%.

iii) Check your answers to part i) using the [FUTURE AMORTIZATION-CONTINUOUS APPROXIMATION 5.6.13](#).

iv) Check your answers to part ii) using the [FUTURE AMORTIZATION-SIMPLE INTEREST APPROXIMATION 5.6.16](#).

The answers to this problem deserve some examination as they contain an important message: *The first commandment of saving for retirement is: "Start young!"*. Suppose that, instead of asking you to work this problem, I had just asked you, "How much more do you have to deposit each month if you save for retirement for 20 years instead of 40?". Your first guess would probably be, "Twice as much since you are saving for only half as long." The problem shows that this is wrong: the real figures are from 3 to 7 times as much! The discrepancy is due to the effect of compounding and the answers illustrate two basic facts we have already seen. The longer you save for retirement the more compounding helps your deposits mount up: that's why no matter what the interest rate the deposits you

5.7 Planning for future needs

make starting at age 45 need to be substantially *more* than twice those you make starting at 25. And, the higher the interest rate, the more pronounced the effects of compounding become. At 3%, you need to put in \$1,520.46 a month starting at 45 but only a third as much—\$538.77—starting at 25. At 9%, both deposits are smaller. But the deposit starting at age 45 has only decreased by about half to \$742.63 while the deposit starting at age 25 have gone down by a factor of *five* to \$105.34 (less than a seventh of \$742.63).

I wish someone had made me work this problem when I was your age. The same principle applies to any kind of goal oriented saving. For example, I can kick in much less each month for my daughter's college fund if I start when she is born than if I wait until she is 9 or 10.

FIRST RULE OF SAVING 5.7.6: *Start now!*

Yes, the rate of interest you earn also effects how your savings grow but the effect is less dramatic. Note that a 45 year old whose deposits earn 9% interest still has to deposit one and a half times as much as a 25 year old whose deposits earn only 3%. What's more the 25 year old is actually a better candidate for investments combining high yields with high risk. A few down years have much less of an effect on the final sum when the saving occurs over a long period. In the next problems, you can check this for yourself (with a little help from me if you want).

Suppose that instead the SEP account into which the retirement funds are going has a fluctuating yield. We'll imagine that there are good decades in which the annualized yield is 12% and bad decades in which the yield is 0%. If there are three good decades for every bad one then this is roughly like getting a 9% yield.

PROBLEM 5.7.7: Explain why the assumption above is only *very* roughly like getting a 9% yield. Illustrate your answer with a couple of example scenarios.

5.7 Planning for future needs

PROBLEM 5.7.8: A 25 year old puts \$105.34 a month each month for 40 years into a SEP account which he *hopes* will have a 9% yield while a 45 year old puts \$742.63 a month each month for 20 years into a SEP account which she *hopes* will have a 9% yield. If they are right, each will have a final sum of \$500,000.00 in the account. Instead the account follows the pattern above with good and bad decades.

i) Suppose that for each account the first decade is a bad one and the rest are good. How much will each final sum be?

Solution

We do not have a formula to fit this scenario but we can track it by combining two formulas. I'll work things out for the 25 year old and leave the 45 year old to you. The first step is to ask what the sum is at the end of the 10 years of zero yield. That's easy: there have been 120 payments of \$105.34 which has earned no interest and so amount to \$12,640.80.

The second and key step is to separate this sum from the rest of the payments. It will now earn 12% interest compounded monthly (so $p = 0.01$) for 30 years (so $T = 360$) amounting to $\$12,640.80 \cdot (1 + 0.01)^{360} = \$454,432.23$. The rest of the payments form, in essence, a retirement fund with monthly deposits of \$105.34 made for 30 years and earning a yield of 12%. These have a final sum B of

$$S = D \left(\frac{(1+p)^T - 1}{p} \right) = \$105.34 \left(\frac{\left(1 + \left(\frac{0.01 \cdot 1}{1}\right)\right)^{360} - 1}{\left(\frac{0.01 \cdot 1}{1}\right)} \right) = \$368,159.52.$$

The two give a combined final sum of \$822,591.75 not very close to \$500,000.00: this certainly confirms the previous problem.

It also confirms the **FIRST RULE OF SAVING 5.7.6**. After the bad decade the 25 year old had barely 2% of his goal saved, but because he had those 3 good decades left before retiring that sum had time to grow to over 90% of his final target—and to more than his savings over the all three of the good decades.

5.7 Planning for future needs

ii) Suppose that for each account all the decades are good except the *last*. How much will each final sum be?

Hint: You will need a different (and slightly simpler) two part approach to handle this part.

If you followed the steps in the solution correctly, you should get a combined final sum of \$464,949.59 for the 45 year old's sum after a decade of 0% yield and then one of 12% yield. In other words, the swings in yields worked out well for the 25 year old—who had more time to recover from the bad decade and benefit from the good ones—while leaving the 45 year old 10% short of the goal. When the bad decade comes at the end of the term, the effects are more pronounced. Now there is no time for either person to recover and the calculation shows both wind up short of their goals. The 25 year old has \$380,800.32 and the 45 year old has \$259,949.23.

This confirms the principle that you can tolerate risk better when you have a longer time horizon than when you have a shorter one in two ways. The 45 year old falls twice as far short of the goal as the 25 year old: the down decade had more of an effect because it represented a greater fraction of the term. But after the three good decades, with only 10 years to go in his pension contributions, the 25 year old really only had a short 10 year horizon. The fact that the final bad decade leaves him almost 25% short of his goal shows this. He would have been wise to have begun moving his money to some less risky investment at about this time since he goal was too close to allow him to recover from a bad period.

RULE OF THUMB FOR RISK VERSUS TERM 5.7.9: *Risky, up-and-down investments with higher average yields are better choices for savers or investors with long term goals. As our goals approach, vehicles with lower but more certain yields become better choices. When you are retired or close to it, your savings become blood money.*

PROBLEM 5.7.10: Perhaps, you are wondering what happens if the 45 year old lucks out and hits two good decades with a 12% yield. If



5.7 Planning for future needs

he deposits \$742.63 a month each month for 20 years, what will his final sum be in this case?

A few final comments about this example. Note the spread between the worst case here (12% yield in the first 10 years and 0% yield in the last 10 years and a final sum of \$259,949.23) and the best (12% yield through the entire 20 years and a final sum of \$734,650.71): almost a factor of three, from half of the goal of \$500,000.00 to one and half times the goal. When we speak of *risk* we basically mean the possibility of that our actual final sum will differ greatly from our expectation or goal. Since the possibility of surpassing your goal is a pleasant one and that of failing to reach it unpleasant, you might think that this simply adds a mild element of spice to saving. It seldom works out that way.

The reason is that what we can buy with money beyond what we need for our goals is generally much less valuable to us. The 45 year old who lucks out may be able to take a cruise every year with the extra money in his retirement fund. Or, if my daughter's college fund hits \$150,000 instead of just \$120,000.00, then I can buy her a car for a graduation present. But the 45 year old who only has half of what he needs in his retirement fund may wind up having to eat cat food when he is 75, and if my daughter's college fund falls to far short, she'll have to settle for an inferior college with possible lifelong consequences for her career. That's why when assessing risk, the question we ask is not "What's the *best* that can happen?", but "What's the *worst* that can happen?".

Have you noticed a huge gap in these discussions of long term saving? We haven't said a word about *inflation*. The \$120,000.00 sum we have been using is 4 times the *current* \$30,000.00 cost per year of sending a child to a good private university. There's a major problem with that computation. My daughter isn't going to college today, she'll be going in 18 years. So the amount I really need to provide for is the 4 times the annual cost of sending a child to a good private

5.7 Planning for future needs

university *in 18 years*. The project for this section tries to fill in this gap a bit.

PROJECT 5.7.11: The question for this project is: How might inflation and risk combine to affect a saving's goal like my daughter's college fund? You'll get to practice pretty much everything we have done in this section. For simplicity, let's assume annual compounding throughout the problem. Assume that a college education today costs \$120,000.00

- i) What will the cost of a college education be in 18 years if, during the interim, inflation averages
- a. 3%.
 - b. 6%.
 - c. 9%.

Check your answers using the [CONTINUOUS APPROXIMATION 5.3.11](#). The first rate, 3%, is close to what has been the general rate of inflation for the past few years. However, the cost of a college education has been rising faster than the general inflation rate for over a decade. The rate is closer to 6%. Right now the third rate, 9% is well beyond anything we see. But, when we ask “What’s the worst that *can* happen?” we should always allow for possibilities worse than the worst possibilities that *are* happening.

- ii) Next make a table showing how much I will have to contribute every year to my daughter's college fund (as usual, from birth to age 18) to reach *each* of these sums if the college fund has a yield of
- a. 4%.
 - b. 8%.
 - c. 12%.

This table will have 9 entries: check each using either the [FUTURE AMORTIZATION-SIMPLE INTEREST APPROXIMATION 5.6.16](#) or [FUTURE AMORTIZATION-CONTINUOUS APPROXIMATION 5.6.13](#).

- iii) Generally speaking, higher yields are offered to compensate for higher risk. If I could get 8% guaranteed every year, I wouldn't be interested in an investment that offered me 4% some years and 12%

5.7 Planning for future needs

others. Although in the long run they might appear roughly equal, over shorter periods, the second has a much worse “worst case” than the first. Let’s suppose that the fund which yields 8% has periods of 6 years in which the annualized yield is 5% alternating with periods of 6 years in which the yield is 11%. What are the best and worst outcomes for the college fund? To keep the arithmetic from getting out of hand, let’s assume the 6% inflation rate in college costs.

Partial Solution

I am going to work out one scenario to guide you. Once again we can’t just apply a formula but so what? The formulas we have let us work out what happens in each 6 year period so all that is needed is a bit of calculator grease. The only difference from [PROBLEM 5.7.8](#) is that now there are three periods instead of two. Here there are two possibilities: 5% yield for the first 6 years, then 11% for the second 6 and 5% for the last 6, or 11%, then 5%, then 11%. I’ll work things out for the first and you can imitate what I do for the second. You should have found that at 6% inflation I’ll need \$342,520.70 in 18 years to pay for my daughter’s college and that the annual payment will need to be \$8,507.99 (a bit more than \$700 a month) if the fund earns an 8% yield. (Of course, since we’re estimating the future, the to-the-penny precision of these amounts is misleading. Nonetheless, I’ll keep using it just so you can tell whether you have duplicated my calculations correctly when you try them yourself.)

The first six year period is just like a six year amortization with annual payments of $D = \$8,507.99$ earning 5% a year (so $p = \frac{0.01 \cdot 5}{1}$ and $T = 6$). At the end of 6 years, the sum is

$$S = D \left(\frac{(1+p)^T - 1}{p} \right) = \$8,507.99 \left(\frac{\left(1 + \frac{0.01 \cdot 5}{1}\right)^6 - 1}{\left(\frac{0.01 \cdot 5}{1}\right)} \right) = \$57,870.62.$$

Again, the key idea for handling the second 6 years is to keep this lump sum separate from the annual payments. It just earns compound interest of 11% (compounded annually, of course) for

5.7 Planning for future needs

6 years amount to: \$108,242.05. The second six years of payments are just like the first six except that the rate of 5% is now 11%. Making this change in the calculation of S above we see that these payments will amount to \$67,322.55. So after 12 years the account holds a total of \$175,564.60.

Over the last 6 years, this lump sum compounds at 5% to become \$235,273.36. Meanwhile the last 6 annual payments earn 5% too so the formula above applies exactly and these amount to \$57,870.62. The final sum is the sum of these last two amounts, or \$293,143.98. The fund ends up about \$50,000 short—remember at 6% inflation we decided we needed to save \$342,520.70—but since this is only about 15% of the goal we'd probably feel we could cross this bridge if we came to it.

iv) Finally suppose that the fund which yields 12% has 6 year periods with annualized yields of 0%, 12% and 24% and that the 18 year term is made up of one of each. What are the best and worst cases? Again, let's assume that inflation is 6%.

Hint: There are six possible orders for the 0%, 12% and 24% periods.

v) What would you do if you wanted to save for your daughter's college fund? There is, of course, no "right" answer here. Different people have different resources for such savings, different tolerances for risk and soon. Explain why the plan you chose appeals to you, what alternative strategies you rejected and why. What is the worst you outcome you can imagine for your plan?

By the way, parts [iii\)](#) and [iv\)](#) of this project start to give you a feel for how quickly things get complicated when you do not assume that yields are constant.

Let's sum up. I've put a lot of thought into coming up with the financial planning equivalent of Michael Pollan's brilliant 7 word summary of sound dietary principles: "Eat food, mostly plants, not too much". Here it is:

5.8 Amortization: borrowing

DR I'S ADVICE TO SAVERS 5.7.12: *Start now, diversify, allocate by horizon, invest don't speculate.*

We've already discussed most of these points. Time is your greatest ally as a saver and you'll never get back any you lose by not starting now. Diversification is the best way to reduce risk with a single asset class. Allocating your savings across multiple asset classes is the best way to tune your risk. In general, it's smart to concentrate on riskier assets with higher yields when your investment goal or horizon is far off and you can booth profit most from those higher yields and recover best from downward swings. You should move towards safer assets when your horizon nears and you have less to gain from higher yields and more to lose from higher risks.

Only the last point is new. It's also the least universally accepted. The main idea is that smart savers buy assets and hold them for long periods, selling only to make the sort of adjustments to their allocations needed to match their approach of their horizon. This keeps fees and taxes low. The alternative is to try "Buy low and sell high." This sounds much better but a lot of studies have shown that such market timers do not achieve better yields in practice: after factoring in the higher costs and taxes they incur, they fare worse. The classic account is Burtin Malkiel's [A Random Walk down Wall Street](#). But the injunction against speculation also contains a warning against greed. If it sounds too good to be true, it probably is. Ask any of [Bernie Madoff](#)'s clients.

5.8 Amortization: borrowing

In this section, we want to look at the second kind of [AMORTIZATION 5.6.1](#), [loans](#), where we receive the lump sum *before* making a series of equal payments at regular intervals.



5.8 Amortization: borrowing

Here our model example will be a home mortgage. Suppose that, in addition to my down payment, I need \$100,000.00 to purchase a house. The bank will lend me this money as a lump sum mortgage—say at a nominal rate of 7.5% a year—and in return I agree to make payments at the end of every month for a term of 30 years. The basic question here is, “What should my monthly payment be?”. Closely related is the question I might ask myself before starting to look for a home: “If I can afford to budget about \$800.00 a month for my mortgage, how big a mortgage will the bank give me?”.

Why did I separate the **savings** and **loan** amortizations? The answer is not that they are mathematically very different. In fact, we’ll be able to apply the formula derived for savings accounts with very little change. What’s different about loans is that, in most cases, they do not involve the uncertainty of savings accounts. When we take out a loan, we general know going in what the interest rate will be and hence exactly each payment must be to amortize the loan. This kind of predictability is rare for savings and investments, as we have seen. The result is that we can answer precisely questions about loans in which for an investment, we’d have to use a guess or estimate.

The Present amortization formula

To begin with, we are going to figure out a formula for relating the balance B and payment D of a loan *with no calculation!* Before reading on, go back and look at the derivation of the **FUTURE AMORTIZATION FORMULA 5.6.8**. We had to work hard to get that formula even if it turned out to be very simple at the end. So you should be impressed that we can figure out the analogous formula for loans with no work. I’m lying, of course: what we’ll really do is shake the **FUTURE AMORTIZATION FORMULA 5.6.8** until the **PRESENT AMORTIZATION FORMULA 5.8.1** drops out, making the work we have already done pay double dividends. But to a mathematician, that’s always

5.8 Amortization: borrowing

like getting something for nothing, the charm of the subject. This is also where introducing the sum and balance notation (S and B) will pay off.

So, how are the payment and the balance of a loan related? The easiest way to see the answer is to ask: “What’s the difference between an investment and a loan if the interest rate, period, term and deposit of each are the same?”. In other words, suppose we make equal payments or deposits of $\$D$ at equal periods—say m times a year—for a term of T periods into an account that earns a periodic interest rate p . So far, we could either be describing what we do to accumulating money in a savings or investment account or what we do to repay a loan.

What’s the difference? Not a whole lot. In both cases, what we want to calculate is a single lump sum of money which represents the combined value of all the deposits. The only difference is *the point in time* at which that lump sum exists. In a savings account, the lump sum is the *final sum* $S = S_T$ in the account at the *end* of the term of the amortization. In a loan repayment, the lump sum is the *initial balance* $B = B_0$ of the loan at the *start* of the term of the amortization. To repeat, both these lump sums represent the combined value of the series of T deposits of $\$D$. Lets call this A to be neutral. The only difference is the moment in time at which this combined value is computed. The initial balance B of a loan is the value of the amount A at the *start* of the term of the amortization, after 0 periods have passed: in other words, $B = A_0$. The final sum S in a savings account is the value of the amount A at the *end* of the term of the amortization, after all T periods have passed: in other words, $S = A_T$.

But the [COMPOUND INTEREST FORMULA 5.2.4](#) tells us how A_0 and A_T are related: $A_T = A_0 \cdot (1+p)^T$. Thus, we conclude that $S = B \cdot (1+p)^T$, or multiplying both sides by $(1+p)^{-T}$, that $B = S \cdot (1+p)^{-T}$. Bingo! Now all we have to do is plug in the value of S given by the [FUTURE](#)

5.8 Amortization: borrowing

AMORTIZATION FORMULA 5.6.8, munge a few exponents and we find that

$$B = S \cdot (1+p)^{-T} = D \left(\frac{(1+p)^T - 1}{p} \right) \cdot (1+p)^{-T} = D \left(\frac{1 - (1+p)^{-T}}{p} \right)$$

PRESENT AMORTIZATION FORMULA 5.8.1: *If a loan which earns compound interest at a periodic rate p is repaid by a series of payments made at the end of each of T periods, then the initial balance B of the loan and the amount D of each payment are related by*

$$B = D \left(\frac{1 - (1+p)^{-T}}{p} \right) \quad \text{and} \quad D = B \left(\frac{p}{1 - (1+p)^{-T}} \right).$$

Before we use this, a warning is in order. The key observation that let us derive this formula with no calculation was that $B_0 = A_0$ and that $S_T = A_T$. What about all the intermediate balances B_i and sums S_i after i deposits or payments are made? Can we relate these to the value A_i of the amount A after i periods? No! The point is that A and hence any A_i involves *all* the payments and this is only true of the *initial* balance B_0 and *final* sum S_T . But not to worry. In a moment, we'll see that the **PRESENT AMORTIZATION FORMULA 5.8.1** actually can tell us about all the intermediate balances too.

Working with the present amortization formula

Once again, the formula is beautifully simple to use. Here are the answers to the mortgage questions at the start of this section.

EXAMPLE 5.8.2: I am going to borrow \$100,000.00 from the bank at 7.5% interest and make monthly payments for thirty years. Since the payments are monthly, $m = 12$ and $p = \frac{0.01 \cdot 7.5}{12}$. Since the term is $y = 30$ years, we have $T = m \cdot y = 12 \cdot 30 = 360$. Thus, plugging into $D = B \left(\frac{p}{1 - (1+p)^{-T}} \right)$, we find that

$$D = 100,000.00 \left(\frac{\left(\frac{0.01 \cdot 7.5}{12} \right)}{1 - \left(1 + \left(\frac{0.01 \cdot 7.5}{12} \right) \right)^{-360}} \right) = \$699.2145093$$



5.8 Amortization: borrowing

and my monthly payment will be \$699.21.

If I think I can budget \$800.00 a month for my mortgage payment, then using $B = D \left(\frac{1-(1+p)^{-T}}{p} \right)$, we find that we can repay a mortgage with a balance of

$$B = \$800.00 \left(\frac{1 - \left(1 + \left(\frac{0.01 \cdot 7.5}{12} \right) \right)^{-360}}{\left(\frac{0.01 \cdot 7.5}{12} \right)} \right) = \$114,414.1017$$

or about \$114,000 dollars.

We'll come back to consider other ways to use the **PRESENT AMORTIZATION FORMULA 5.8.1** in a moment but first let's formalize what we did above with a method.

As always, the first two steps are the same (find m and use it to get p and T) and the third just involves plugging values into the formula.

METHOD FOR SOLVING PRESENT/LOAN AMORTIZATIONS 5.8.3:

- Step 1: Determine the periods in the problem (that is, the units in which the term is measured) and the value of m , the number of periods per year.
- Step 2: Use the **INTEREST RATE CONVERSION FORMULA 5.1.10** to find the periodic interest rate p from the nominal interest rate r and the **TERM CONVERSION FORMULA 5.1.13** to find the term T in periods from the term in years y .
- Step 3: Apply the appropriate **PRESENT AMORTIZATION FORMULA 5.8.1** to find whichever of the the deposit D and the balance B is to be determined.

Here are a few exercises for you to try which involve mortgages and other consumer loans. Since such loans are almost invariably paid on a monthly basis, I have *not* explicitly stated the payment frequency unless it is *other* than monthly. As usual, we want to work to the nearest cent. In loans, the rounding error of a fraction of a cent

5.8 Amortization: borrowing

per payment gets multiplied by a number of payments which can be large—in the hundreds—and this can cause an discrepancy of a few dollars between the exact value of all the rounded payments and the balance of the loan. We'll just ignore this but—surprise, surprise—the bank doesn't. The final payment is usually adjusted to make things balance exactly. I have worked a few of the exercises as further examples.

PROBLEM 5.8.4: What is the monthly payment on a mortgage with a balance of \$100,000.00 if

- i) the nominal interest rate is 7.5% and the term of the mortgage is a. 25 years?

Solution

Step 1: The periods are months so $m = 12$.

Step 2: $p = \frac{0.01 \cdot r}{m} = \frac{0.01 \cdot 7.5}{12}$ and

$$T = my = 12 \cdot 25 = 300.$$

Step 3: Here we know the loan balance $B = \$100,000.00$ so we plug into find the payment $D = B \left(\frac{p}{1 - (1+p)^{-T}} \right)$ getting

$$D = 100,000.00 \left(\frac{\left(\frac{0.01 \cdot 7.5}{12} \right)}{1 - \left(1 + \left(\frac{0.01 \cdot 7.5}{12} \right) \right)^{-300}} \right) = \$738.99.$$

b. 20 years?

c. 15 years?

- ii) the nominal interest rate is 9% and the term of the mortgage is

a. 30 years?

b. 20 years?

c. 15 years?

Solution

Step 1: The periods are months so $m = 12$.

Step 2: $p = \frac{0.01 \cdot r}{m} = \frac{0.01 \cdot 9}{12}$ and $T = my = 12 \cdot 15 = 180$.

Step 3: Here we know the loan balance $B = \$100,000.00$ so

5.8 Amortization: borrowing

we plug into $D = B \left(\frac{p}{1 - (1+p)^{-T}} \right)$ to find the payment

$$D = 100,000.00 \left(\frac{\left(\frac{0.01 \cdot 9}{12} \right)}{1 - \left(1 + \left(\frac{0.01 \cdot 9}{12} \right) \right)^{-180}} \right) = \$1,014.27.$$

PROBLEM 5.8.5: A car dealer offers to sell you a new car, “No money down and easy monthly payments”. You calculate that you can afford to make car payments of \$250.00 a month. How much can you afford to pay for the car if

- i) the nominal interest rate on the loan is 4.5% and the term is
- 2 years?

Solution

Step 1: The periods are months so $m = 12$.

Step 2: $p = \frac{0.01 \cdot r}{m} = \frac{0.01 \cdot 4.5}{12}$ and $T = my = 12 \cdot 2 = 24$.

Step 3: Here we know the payment $D = \$250.00$ so we plug into $B = D \left(\frac{1 - (1+p)^{-T}}{p} \right)$ to find the loan balance B

$$B = \$250.00 \left(\frac{1 - \left(1 + \left(\frac{0.01 \cdot 4.5}{12} \right) \right)^{-24}}{\left(\frac{0.01 \cdot 4.5}{12} \right)} \right) = \$5,727.66.$$

b. 3 years?

c. 4 years?

- ii) the nominal interest rate is 12% and the term is

a. 2 years?

b. 3 years?

c. 4 years?

Solution

Step 1: The periods are months so $m = 12$.

Step 2: $p = \frac{0.01 \cdot r}{m} = \frac{0.01 \cdot 12}{12}$ and $T = my = 12 \cdot 4 = 48$.

Step 3: Here we know the payment $D = \$250.00$ so we plug into $B = D \left(\frac{1 - (1+p)^{-T}}{p} \right)$ to find the loan balance B

$$B = \$250.00 \left(\frac{1 - \left(1 + \left(\frac{0.01 \cdot 12}{12} \right) \right)^{-48}}{\left(\frac{0.01 \cdot 12}{12} \right)} \right) = \$9,493.49.$$

5.8 Amortization: borrowing

I hope you can guess what's coming next. Yes, I'm going to ask how we can check such calculations. This time there will be a version of the [CONTINUOUS APPROXIMATION 5.3.11](#) and of the [SIMPLE INTEREST APPROXIMATION 5.3.3](#) and something new.

To use the [CONTINUOUS APPROXIMATION 5.3.11](#), we simply approximate the exponential $(1 + p)^{-T}$ in the [PRESENT AMORTIZATION FORMULA 5.8.1](#) by the slightly smaller exponential $e^{-(p \cdot T)} = e^{(0.01r \cdot (-y))}$ getting a slightly larger numerator. In the formula for D where this appears in the denominator and we are now *dividing* by a larger quantity, we get an approximation slightly smaller than the exact value.

PRESENT AMORTIZATION-CONTINUOUS APPROXIMATION 5.8.6: B is a bit less than $D \left(\frac{1 - e^{(0.01r \cdot (-y))}}{p} \right)$.

There's still a periodic rate p in the formula as in the [FUTURE AMORTIZATION-CONTINUOUS APPROXIMATION 5.6.13](#). If you forgot to convert the nominal rate when using the [PRESENT AMORTIZATION FORMULA 5.8.1](#), you'll probably use r for p here too. Fortunately, the different numerator will lead to a different answer and let you catch your mistake. This is another formula which you don't really need to learn. You can make the necessary approximations if you just remember to use the [CONTINUOUS APPROXIMATION 5.3.11](#) to replace the $(1 + p)^{-T}$.

EXAMPLE 5.8.7: Let's check the calculation in [EXAMPLE 5.8.2](#). Here we had $r = 7.5\%$ and $y = 30$ years and since we were compounding monthly $m = 12$ and $p = \frac{0.01 \cdot 7.5}{12}$. Our D was \$699.21 so the balance B of \$100,000.00 should be a bit less than

$$D \left(\frac{1 - e^{(0.01r \cdot (-y))}}{p} \right) = 699.21 \left(\frac{1 - e^{(0.0175 \cdot (-30))}}{\left(\frac{0.01 \cdot 7.5}{12} \right)} \right) = \$100,082.2093$$

and it is. You can check a loan balance calculation the same way.

5.8 Amortization: borrowing

PROBLEM 5.8.8: Use the **PRESENT AMORTIZATION-CONTINUOUS APPROXIMATION 5.8.6** to check your answers to part b) of each of **PROBLEM 5.8.4** and **PROBLEM 5.8.5**.

To apply the **SIMPLE INTEREST APPROXIMATION 5.3.3** to loans, we use the same basic idea as we did for future amortizations. However, we apply it to the balance rather than the deposits because the arithmetic is then easier. We ask, “What’s the average outstanding balance on the loan over its term?” and we answer “About half the initial balance”. So the simple interest on the balance should be roughly $0.01r \cdot y \cdot \frac{B}{2}$. (In the next section, we’ll see that this underestimates the average balance, but not by an amount we need to worry about here).

The total of all the T payments of $\$D$ each should match the original balance plus this interest. In fact, they should amount to somewhat more as the actual compound interest on the outstanding balance will be greater than the simple interest approximation we have used.

PRESENT AMORTIZATION-SIMPLE INTEREST APPROXIMATION 5.8.9: $T \cdot D$ is larger—possibly quite a bit larger—than $B \left(1 + 0.01r \cdot \frac{y}{2}\right)$.

Once again, the best feature of this check is that you can do it roughly in your head. Here are a couple of examples.

EXAMPLE 5.8.10: Let’s check the calculation in **EXAMPLE 5.8.2**. Here we had $r = 7.5\%$ and $y = 30$ years and since we were compounding monthly $m = 12$ and $T = 360$. Our D was $\$699.21$ so $T \cdot D = \$251,715.60$. This should be somewhat larger than

$$B \left(1 + 0.01r \cdot \frac{y}{2}\right) = \$100,000.00 \left(1 + 0.01 \cdot 7.5 \cdot \frac{30}{2}\right) = \$212,500.00$$

and it is. As usual, the agreement is not very good here because the term was fairly long.

We’ll get better agreement if we check a problem with a shorter term like i) of **PROBLEM 5.8.5**. Here we had $r = 4.5\%$, $y = 2$, $m = 12$, $T = 24$, $D = \$250$ and $B = \$5,727.66$. So we expect $T \cdot D = \$6,000.00$ to be a bit larger than



5.8 Amortization: borrowing

$$B \left(1 + 0.01r \cdot \frac{Y}{2} \right) = \$5,727.66 \left(1 + 0.01 \cdot 4.5 \cdot \frac{2}{2} \right) = \$5,985.40.$$

Since the term was short we get excellent agreement.

PROBLEM 5.8.11: Use the **PRESENT AMORTIZATION-SIMPLE INTEREST APPROXIMATION 5.8.9** to check your answers to part b) of each of **PROBLEM 5.8.4** and **PROBLEM 5.8.5**.

One nice feature of loan amortizations is that when the term is long—exactly when **PRESENT AMORTIZATION-SIMPLE INTEREST APPROXIMATION 5.8.9** is way off—there is another and even easier way to check your answer. The idea is that when T is big the negative exponential $(1 + p)^{-T}$ will be small. For example in **EXAMPLE 5.8.2**, we have $T = 360$ and $p = \frac{0.01 \cdot 7.5}{12}$ so $(1 + p)^{-T} = 0.1061398302$. So we don't lose too much by ignoring this.

We call this the **PRESENT AMORTIZATION-INTEREST APPROXIMATION 5.8.12** because the estimate it give for the payment D simply equals one periods *interest* on the initial balance of the loan. If you made this payment every month, your outstanding balance would never change: each month you'd pay off the preceding month's interest but you'd never reduce the outstanding balance on the loan. So you'd be paying the loan forever. One consequence is that this approximation would be exactly correct if the term were infinite and is accurate only when the term is fairly long. We get an estimate for B which is bit large and one for D which is a bit small and both are stunningly simple.

PRESENT AMORTIZATION-INTEREST APPROXIMATION 5.8.12: *If an amortized loan has a long term, the the balance B is slightly smaller than $\frac{D}{p}$ and the payment D is slightly larger than $B \cdot p$.*

EXAMPLE 5.8.13: Let's check the calculation in **EXAMPLE 5.8.2**. Here we had $p = \frac{0.01 \cdot 7.5}{12}$ and $B = 100,000.00$ so we'd expect D which turned out to be \$699.21 to be a bit larger than $100,000.00 \cdot \left(\frac{0.01 \cdot 7.5}{12} \right) = \625.00 . Note how simple the check is.



5.9 Things everyone should know

The downside here is that for loans with short terms the approximation is very poor. For part i) of [PROBLEM 5.8.5](#), where $p = \frac{0.01 \cdot 4.5}{12}$ and $B = \$5,727.66$ we get the estimate $D \simeq \$5,727.66 \cdot \left(\frac{0.01 \cdot 4.5}{12} \right) = \21.48 which is *much* smaller than the exact payment of \$250.00.

PROBLEM 5.8.14: Compute the monthly payment on a \$80,000.00 mortgage at a rate of 8.4% and compare it with approximation given by the [PRESENT AMORTIZATION-INTEREST APPROXIMATION 5.8.12](#) for terms of

- i) 20 years.
- ii) 30 years.
- iii) 40 years.

5.9 Things everyone should know

This section is a short introduction to issues that you'll almost certainly confront in managing your personal finances. You'll learn some handy tricks—like how to compute the balance on a mortgage or any other amortizing loan—but, more important, I hope you'll learn how to maximize your spending power and how to avoid common traps in managing your personal credit.

Equity

“Why rent when you can own? Paying rent is just throwing money away. If you buy a house, you'll be building equity.” This is a classic real estate agent's argument to convince renters to become homeowners. Does it make sense? For the real estate agents, definitely. When a property changes hands, the agents for both the buyer and the seller collect several percent of the the selling price. That's their business so they are interested in generating the greatest number of



5.9 Things everyone should know

sales possible both by convincing new buyers to enter the market and by convincing existing owners to “trade up”. The first question I want to look at in this section is what your position as a potential buyer should be.

Let’s begin by trying to understand what “building equity” means. The idea is simple. In discussing the [PRESENT AMORTIZATION-INTEREST APPROXIMATION 5.8.12](#), we saw that the actual periodic payment D you make on a loan is somewhat greater than the periodic interest $B_0 \cdot p$ on the initial loan balance B_0 . (By the [SIMPLE INTEREST FORMULA 5.1.6](#), the interest in one period is $I = p \cdot A \cdot T = p \cdot B_0 \cdot 1 = pB_0$.) For example on a loan of $B_0 = \$100,000.00$ at 7.5% interest which is paid monthly (so $p = \frac{0.01 \cdot 7.5}{12}$), we’d owe interest of $B_0 \cdot p = \$625.00$ after one month. The difference $D - B_0 \cdot p$ is the part of the *first* payment which is used to reduce the outstanding balance on the loan.

Note that while the interest owed is the same whatever the *term* of the loan—in our example, \$625.00—this difference *will* depend on the term. We saw in [PROBLEM 5.8.4](#) that a loan with a 30 year term the payment $D = \$699.21$ and hence the difference $D - B_0 \cdot p = \$74.21$ while for a loan with a 15 year term we have a payment of $D = \$927.01$ and a difference $D - B_0 \cdot p = \$302.01$.

In any case, at the end of the first period, your outstanding balance B_1 has decreased a bit. Let’s track this in the case of the 30 year mortgage. The reduction in balance is $(D - B_0 \cdot p) = \$74.21$ so $B_1 = B_0 - (D - B_0 \cdot p) = \$99,925.79$. In the second month you incur periodic interest of $B_1 \cdot p = \$99,925.79 \frac{0.01 \cdot 7.5}{12} = \624.54 . Since B_1 is a bit smaller than B_0 , this interest is a bit smaller (46¢ smaller to be precise) than the periodic interest due in the first month, so the difference $D - B_1 \cdot p$ between what you pay and the interest you owe is a bit bigger (46¢ bigger). This difference is also the reduction in your outstanding balance in the second month so $B_1 - B_2$ is a bit bigger than $B_0 - B_1$. I’ll leave you to check that B_2 comes out to

5.9 Things everyone should know

be \$99,851.14. This process continues. We won't go any further with our example because it's already pretty clear both that we *could* compute more balances if we wanted to, and that we don't much want to unless we can find a better approach. Imagine finding the balance after 5 years—or 60 months. We have to make 60 calculations like those above. No thanks!

Let's just summarize what we've seen from the first two calculations. Each payment you make reduces your outstanding balance and each reduction is a bit *bigger* than the preceding one. That is, the amount of the reduction gets bigger and bigger every payment. Equity is the name for these reductions.

EQUITY 5.9.1: *The difference $E_i = B_0 - B_i$ between the initial balance B_0 of a loan and the intermediate balance B_i after i payments is called your equity after i months. The equity increases with each payment made and the size of the increase also grows from payment to payment.*

We tend to think of equity as a form of savings or investment particularly with respect to mortgage loans against real estate. Let's assume that *the value of the property on which the mortgage is held is equal to the initial balance of the mortgage at all times*: this is wildly unrealistic but we'll ignore this for now. Assuming this, then the equity E_i is the amount of extra cash you could realize if you sold the property and paid out the outstanding balance after i periods. You'd get back the initial balance B_0 as the sale price, pay the bank the outstanding balance B_i and be left with the equity $E_i = B_0 - B_i$ in your hot little hand. In particular, at the end of the term of the mortgage, when the outstanding balance $B_T = 0$, your equity equals the value B_0 of the property.

Mortgage based equity was traditionally an important component of overall savings in the United States, particularly for households approaching retirement age. You'd buy a house in your 20's or early 30's, pay off the 30 year mortgage at roughly age 60 and the equity

5.9 Things everyone should know

in the house would be your “nest egg” for retirement. Two things have altered this fact. The first is the increased role of targeted retirement funds and pensions. The second is the increased mobility of American families. It is this second effect which I’d like you to understand.

To do so, we’ll need a formula for the equity E_i after i periods in terms of the basic quantities B and D . We can unwind the calculations we did above in the case of the 30 year mortgage to get such formulae. Indeed, we already found a formula for E_1 above: $E_1 = B_0 - B_1 = (D - B_0 \cdot p)$. We could also have found a formula for E_2 : the reduction in balance in the second month was $D - B_1 \cdot p$ so $E_2 = E_1 + (D - B_1 \cdot p)$ and since $B_1 = B_0 - (D - B_0 \cdot p)$ and $E_1 = (D - B_0 \cdot p)$,

$$E_2 = (D - B_0 \cdot p) + (D - (B_0 - (D - B_0 \cdot p)) \cdot p).$$

Did that go by a bit quick? Not to worry because it is pretty clear that even if we understand E_2 in this way, things are going to be too complicated to get a general formula. We can see that each payment reduces the balance a bit more but the exact amount of the reduction in the i^{th} payment depends on the values of the reductions in *all* the preceding payments. It looks hopeless.

What we need is another way to think about E_i . The basic formula $E_i = B_0 - B_i$ provides just this. We know B_0 just equals B . So, if we can find a simple formula for the outstanding balance B_i we can use it to get a simple formula for E_i . We have just such a formula: the [PRESENT AMORTIZATION FORMULA 5.8.1](#)! The trick making this do the work is to ask, “What would the bank be willing to accept as a substitute for the lump sum balance B_i *at the end of the i^{th} period?*”

We know one answer: the remaining $(T - i)$ periodic payments of $\$D$ which we would have to make if we did *not* pay of the balance after i periods. In other words, the balance B_i is the answer to the question, “How big a loan can I take out at a periodic interest p if I am willing

5.9 Things everyone should know

to make $T - i$ periodic payments of $\$D$?” Thus, we can think of B_i as the *initial* balance of a loan with payment $\$D$, periodic rate p and *shortened term* of $T - i$ periods. The [PRESENT AMORTIZATION FORMULA 5.8.1](#) tells that this initial balance is

$$B_i = D \frac{(1 - (1 + p)^{-(T-i)})}{p}$$

and hence that

$$E_i = B_0 - B_i = B - D \frac{(1 - (1 + p)^{-(T-i)})}{p}.$$

This is another of those cases where it is much easier to remember the idea than to learn a new formula. After all, we really just used the [PRESENT AMORTIZATION FORMULA 5.8.1](#).

BALANCE AND EQUITY PRINCIPLE 5.9.2: *The intermediate balance B_i outstanding after exactly i payments have been made on a loan with initial balance B , payment D , periodic rate p and term T periods equals the initial balance of a loan with payment D , periodic rate p and shortened term $T - i$ periods. The equity E_i after i payments is just the difference between the initial balance B and the i^{th} intermediate balance B_i .*

We can put this more informally: to get the balance after i payments, just shorten the term by i periods. Because we’re using a formula we already understand, we can jump right in and work equity and balance problems. Here’s the method.

METHOD FOR FINDING BALANCE AND EQUITY 5.9.3:

Step 1: We start by using the [METHOD FOR SOLVING PRESENT/LOAN AMORTIZATIONS 5.8.3](#) to find the payment amount D on the loan as usual—often, we will already know this and so can skip this step.

Step 2: Use the [PRESENT AMORTIZATION FORMULA 5.8.1](#) with this payment and the *same* periodic rate to compute the intermediate balance B_i remembering to replace the original term T with the number of payments $T - i$ *not yet made* on the

5.9 Things everyone should know

loan. In other words, we just *subtract* the number i of payments *made* from T : remember that i is in units of *payments* or periods not units of *years*. If we're asked for a balance after z years we'll have to convert the term of z years to an equivalent number i of periods or payments.

Step 3: Compute the intermediate equity E_i by subtracting the intermediate balance B_i from the initial balance B .

EXAMPLE 5.9.4: Here's a typical example of how we use this. Let's find the outstanding balance and equity on a \$100,000.00 mortgage at interest of 7.5% with a term of 30 years at the end of 5 years.

Solution

Step 1: Finding the payment for this mortgage was part **i)** of **PROBLEM 5.8.4** so we can just borrow the values from solution given there. As usual $m = 12$ and we saw that $p = \frac{0.01 \cdot 7.5}{12}$, $T = 360$ and $D = \$699.21$.

Step 2: Over the first 5 years of the mortgage we will make $i = 5 \cdot m = 60$ payments, leaving $T - i = 360 - 60 = 300$ payments to go, so using the formula $B_i = D \left(\frac{1 - (1+p)^{-(T-i)}}{p} \right)$ we find that

$$B_{60} = \$699.21 \left(\frac{1 - \left(1 + \left(\frac{0.01 \cdot 7.5}{12} \right) \right)^{-300}}{\left(\frac{0.01 \cdot 7.5}{12} \right)} \right) = \$94,616.83.$$

Step 3: The desired equity E_{60} is given by $E_{60} = B - B_{60} = \$100,000.00 - \$94,616.83 = \$5,383.17$.

Here are some exercises for you to try.

PROBLEM 5.9.5: Find the outstanding balance and equity on a mortgage with an amount of \$100,000.00 at interest of 7.5% with a term of 15 years at the end of:

- i) 5 years.
- ii) 10 years.

Solution



5.9 Things everyone should know

Step 1: Here we need to figure out the payment first. We again have $B = \$100,000.00$, $m = 12$ and $p = \frac{0.01 \cdot 7.5}{12}$ but now $T = 15 \cdot 12 = 180$ so $D = B \left(\frac{p}{1 - (1+p)^{-T}} \right)$ gives

$$D = 100,000.00 \left(\frac{\left(\frac{0.01 \cdot 7.5}{12} \right)}{1 - \left(1 + \left(\frac{0.01 \cdot 7.5}{12} \right) \right)^{-180}} \right) = \$927.01.$$

Step 2: Over the first 10 years of the mortgage we will make $i = 10 \cdot 12 = 120$ payments, leaving $T - i = 180 - 120 = 60$ payments to go, so the intermediate balance we want is

$$B_{120} = \$927.01 \left(\frac{1 - \left(1 + \left(\frac{0.01 \cdot 7.5}{12} \right) \right)^{-60}}{\left(\frac{0.01 \cdot 7.5}{12} \right)} \right) = \$46,262.72.$$

Step 3: The desired equity is $E_{120} = B - B_{120} = \$100,000.00 - \$46,262.72 = \$53,737.28$.

iii) 15 years. (No calculator allowed!)

PROBLEM 5.9.6: Find the outstanding balance and equity on a mortgage with an amount of \$100,000.00 at interest of 7.5% with a term of 30 years at the end of

- i) 10 years.
- ii) 15 years.
- iii) 20 years.
- iv) 25 years.

Stop for a moment and stare at your answers to these problems before going on. They are striking in several ways.

The first worked example describes a scenario which the author has seen played out many times by friends. A young family buys a house with a 30 year mortgage (usually they cannot afford the higher payment which would go with a shorter term) and then moves after 5 years (because of a job change or a divorce, to buy a bigger house etc.). Note that the equity such a family accumulates over the first 5 years they pay the mortgage before reselling the house is only about

5.9 Things everyone should know

\$5,000.00. Closing costs on the purchase of a house of this value—closing costs are the commissions and fees you pay to the real estate agent, banks, title search agency, and lawyers, are likely to be at least this large. In other words, such a family will not build up enough equity to recover its initial closing costs in buying the house—not to speak of what the yield on this outlay might have amounted to over the 5 years period nor of the closing costs involved in selling the house at the end of the term.

The second worked example might describe a family who can afford to take out a mortgage with a term of 15 years and then stay in the house for 10 years before moving. This is different in two ways. The first difference is that the term is only half as long (15 years instead of 30). This means a higher payment but the extra amount is surprisingly small. If we ignored interest, we'd need a payment twice as big to pay off a loan in half the time. The effect of compounding is to compress this gap strikingly: the actual difference of about \$228 a month is less than 25% of the 15 year payment. The second difference in this example is that the family stays twice as long. We'd expect both factors to lead to a larger equity. One guess might be to double the equity from about \$5,000 to about \$10,000 to account for doubling the payments made and then add 25% to account for the higher payment: this would give equity of \$12,500 or so. The actual equity is more than *four times* this guess and close to *ten times* what the first family would accumulate!

What are we missing? The answer is that we are making a simple mistake in two ways. Instead of comparing the payments for the two mortgages we should be comparing the amount of each payment which *remains* after deducting interest owed. These are the amounts that yield equity. We saw above that for the first payment the remainder was \$74.21 for the 30 year mortgage and \$302.01 for the 15 year mortgage. The remainder after we remove the interest in the 15 year mortgage is *four times* the remainder in the 30 year mortgage. So we

5.9 Things everyone should know

can improve our guess and say that the second family should have 8 times the equity of the first—a factor of two from the number of payments and a factor of 4 for the size of the remainder.

So, why is the actual multiple closer to 10 than 8? For the same reason in another guise. Another way to say that the equity on the 15 year mortgage is growing faster is to say that the outstanding balance is shrinking faster. Therefore, the interest due each month on the outstanding balance is also dropping faster making the difference between the two remainders at the ends of the respective periods is even bigger. After 5 years the balance outstanding on the 30 year mortgage will be \$94,616.83 so the interest owed in the 61st month will be \$591.36 and the remainder in this payment will be $\$699.21 - \$591.36 = \$107.85$. The balance after 120 months on the 15 year mortgage is \$46,262.72 so the interest owed in the 121st month will be \$289.14 and hence the remainders in this month will be $\$927.01 - \$289.14 = \$637.87$. The ratio of the two remainders thus ends up at 6 not 4. On the average this ratio should be about 5 which with the factor of 2 for the term gives us the factor of 10 we see in the equities.

To really get a feel for how equity build up, a picture is worth a thousand words. Each of the graphs in [FIGURE 5.9.7](#) plots the number of payments made on the x-axis against the equity built up after that many payments on the y-axis. In both graphs, the interest rate is 7.5% but on the left I have used a term of 30 years while on the right the term is 15 years. Actually, I have plotted the equity as a *fraction* of a initial loan balance B which is why the ticks on the y-axis go from 0 (no equity) to 1 (no outstanding balance). For example, after 60 payments on the 30 year mortgage the graph is at height about 0.05: on a balance of \$100,000.00 I would have equity of about $0.05 \cdot \$100,000.00 = \$5,000.00$. Likewise, after 120 payments on the 15 year mortgage the graph is at height about 0.54: on a balance of \$100,000.00 I would have equity of about

5.9 Things everyone should know

$0.54 \cdot \$100,000.00 = \$54,000.00$. Remember that you can use the magnifier tool in **Acrobat** to zoom in on the graph.





FIGURE 5.9.7: Relative equity by number of payments made

PROBLEM 5.9.8: Check your answers to [PROBLEM 5.9.6](#) and [PROBLEM 5.9.5](#) against the graphs.

CHALLENGE 5.9.9: Is it really OK to ignore the loan balance? More precisely, the way I drew the graph above implicitly claims that the equity accrued as a fraction of initial balance only depends on the rate r and term T of the loan and the number i of payments made, not on the initial balance B .

- i) First, check this by computing a few such fractions for a loan of \$200,000.00 at 7.5% interest and comparing with known answers for a loan of \$100,000.00 (You can save half the work if you use as examples a loan with a 30 year term after $i = 60$ payments are made and one with a 15 year term after $i = 120$ payments are made.)
- ii) Next, use [PRESENT AMORTIZATION FORMULA 5.8.1](#) to show that the fractions are the same for any two balances B and B' if r , T and i are the same.
- iii) Finally, explain why the principle of [EQUALITY OF DOLLARS 5.1.4](#) implies that the fractions in the graph are independent of the initial loan balance.

5.9 Things everyone should know

Let me make a point about these checks. The check in [i\)](#) looks easiest—it's just a standard concrete balance problem—but is the most work. The check in [ii\)](#) is a bit scary—you've got to write down the quotient of two medium sized formulas—but when you do the B 's just cancel. But the conceptual argument in [iii\)](#) is the easiest of all. Moral: an ounce of inspiration is worth a pound of perspiration.

What conclusions can we draw from these graphs? First, staying 10 years won't help the family with the 30 year mortgage much. Their equity will be just over a tenth of the value of their mortgage after 10 years—actually about 0.13—so on a balance of \$100,000.00 they will only have accumulated about \$13,000 in equity. They will need to keep their house for over 20 years before they build up equity of more than \$50,000.

PROBLEM 5.9.10: Use the [METHOD FOR FINDING BALANCE AND EQUITY 5.9.3](#) to check that after 22 years this family still owes a touch more than \$50,000 on their mortgage.

Not many people stay in the same house this long today. So, we conclude that if you need to take out a 30 year mortgage to finance a home purchase, you should probably *not* expect to build up much equity before moving. Of course, you *might* stay put for 30 years and build up \$100,000.00 in equity but while this was common thirty or forty years ago, you'd be a rare exception if you did this today.

A second conclusion we can draw is that even if you can afford to take out a 15 year mortgage, you need to stay for most of the term to build substantial equity. After 5 years, your equity fraction will be about 0.22 and after 10 (as we've seen) it's about 0.53. In other words, you create roughly 4 times as much equity in the last 10 years as the first 5 and almost as much in the last 5 years as you do in the first 10.

Don't get me wrong. I have nothing against buying a home per se. It is sometimes cheaper than buy than to rent even when you include

5.9 Things everyone should know

closing costs, insurance, maintenance, taxes and so on and even if you are not planning to stay in the place you are buying “forever”. One important factor that I haven’t mentioned is that mortgage interest is often deductible from your income for federal tax purposes. But, if you are concerned about money, this is the calculation you should be making: which is cheaper *all costs included*.

Buying does not turn out to be cheaper all that often. It’s true that there are intangible values to owning your home—no landlord, you can fix it up any way you want, ... —the whole “It’s our very own place” factor; or, the kind of home you want may be impossible to rent. But these work both ways: if the roof starts leaking, *you* have to fix it; if you want to move, you have to find a buyer. You have to decide how much these intangibles matter to you. My advice is simply to have a clear idea of relative costs to balance these other factors against. And, unless you can afford a mortgage with a short term and expect stay put for most of that term, don’t make “building equity” one of those intangibles.

PROJECT 5.9.11: This project involves a lot of calculation so you might want to work with some friends and divide up the grunt work.

i) Another common term for mortgages is 20 years. Compute the equity built up in a \$100,000.00 mortgage at 7.5% at 2 year intervals. Plot these on a graph like those above. How do the three graphs compare?

ii) Interest rates also affect how equity grows. Make plots like those above assuming terms of 15 and 30 years but changing the interest rate to 6% and to 9%. Discuss how higher or lower rates seem to affect the growth of equity.

MAJOR PROJECT 5.9.12: Which is cheaper in your area, renting or owning? Imagine that you have an income of \$45,000.00 a year and that mortgage interest rates are 7.5%.

i) Most banks have a rule that your mortgage payment can not be greater than 28% of your income. How much can you afford to spend

5.9 Things everyone should know

on a house if you do not have any money for a down payment? (In real life, you are usually required to include costs like taxes in the 28% figure. Moreover, if you do not have a down payment of 20% of the purchase price, this percentage would be reduced somewhat and the interest rate raised. We'll ignore such facts.)

ii) Write down specifications for a home you'd "like" to buy (number of beds and baths, amenities etc.). Then find current ads for several homes of this type in your area which are for rent and for sale. Compute the *total* cost renting or of owning each property for 3 years, for 5 years and for 10 years. Be sure to include costs like closing costs, maintenance, insurance, utilities and taxes in your calculations in addition to rent or mortgage costs. In budgeting like this, the federal mortgage interest deduction is significant so be sure to include it in your calculation. Remember you must stay within the mortgage budget from i). You may need to research some of these costs with a realtor and or a bank in your area.

iii) Wait, I used to hear many students saying a few years ago (circa 2004-2006): "Real estate is booming". For some reason they stopped in 2007 and 2008. But it's true. I have so far completely ignored the possibility that the value of your home will *increase* while you own it. When you sell, this increased value is like "extra" equity you have accumulated. But prices may also drop and then you may wind up with less or even "negative" equity—upside down as it's known in the trade.

Repeat the calculation of part i) (including the same costs) but incorporating changing rents and real estate values in your area. To do this you will need to find data for homes in your area which are for sale today and were for rent or sale 3, 5 and 10 years ago so you can calculate these changes. To make it easier to find examples, we'll remove the ceiling on your mortgage budget. To be fully fair to renting, you should also include a down payment. For simplicity, suppose that the buyer was required to make a 20% down-payment on each house when you purchased it while the renter had this money

5.9 Things everyone should know

to invest in the stock market over the same period. Use the S&P 500 index to measure the renter's yield. Who comes out ahead now over each of the three terms?

A FINAL TIP 5.9.13: We've seen how much more effective a way to save paying a 15 year mortgage is than paying a 30 year one. What if your bank will only approve you for a 30 mortgage (with its lower payment of \$699.21), but you think you can come up with the 15 year payment of \$927.01? Just make the \$927.01 payment every month and you'll clear the mortgage in 15 years! That's because there are no penalties for prepaying mortgages (unlike many other loans). More generally, if you want to use a mortgage to save, then you're smart to increase your payments by as much as you can afford to shorten the term. The only downside to be aware of is that mortgage interest—but not reduction of principal—can be deducted from your Federal income tax. Since the interest is already covered by the 30 payment amount, any extra you pay off will not get any tax break.

“Money talks, nobody walks”

The title of this subsection is an old used car dealer's slogan. “I'll do anything to sell this car”. If you think the salesman wants the money in your pocket, you're wrong. You don't have to have *any* money in your pocket now because “I'll finance your car with no money down”. The money the dealer wants is the money that is *going* to be in your pocket in the next few years. If he can get you to take out a loan to pay for the car, he'll make a sale and that is what he really wants.

Of course, there is a substantial risk. You have his car and he has only your word that you'll make all the payments. What if you default? That's easy: buyers are charged such a high interest rate that the even after figuring in the losses involved in repossessing from those who do not pay up the loans offer an excellent yield. In fact, in

5.9 Things everyone should know

most cases, the car dealer will immediately sell your loan to a collection agency. He receives less than the face value of the loan but has that loquacious money in his hand. In other words, he has achieved his goal and sold a car—that's his job. The collection agency has the difference between the face value and the sale price of the loan as an extra margin to compensate it for delinquent borrowers as it collects the payments—that's its job. You have the car—that was your goal. It seems like everyone comes out ahead. What's the catch?

The catch comes over the next few years when you make also those easy monthly payments. (Ever noticed how *all* monthly payments are easy?) In this section, I'd like to compare what the costs of such a loan are to those of saving to buy the car. The principles we'll discover apply to many other kinds of consumer credit loans: for furniture, electronics, appliances and so on.

EXAMPLE 5.9.14: Let's say you want a car that sells for \$4,800.00. The salesman is likely to ask whether you can afford to pay \$150.00 a month. "I think so", you say. "Fine", he answers, "We can arrange for you to make 48 easy monthly payments of just \$149.93 a month". What interest rate are you being charged? If you are like most people, you have no idea. There are two ways to find out. The easy one is to look at the finance contract. Down in the small print somewhere you'll find the interest rate because it's a legal requirement to put it there: it's 21.5%. In the next section, [SECTION 5.10](#), we'll learn how to find the interest rate given the balance, payment and term.

Now let's ask what you'd need to do to save for the car and pay cash. I claim that if you put away \$90.00 a month for 48 months you'll have the same \$4,800.00. Before we check this, let's note the main point. If we use the loan payment as a reference, then this means that by saving the money to borrow the car instead of borrowing it, I can reduce my monthly payment by about 40%. If I use the savings deposit as a reference the difference is even more striking: I have to pay fully *two-thirds* again as much to borrow for the car as to save.

5.9 Things everyone should know

Another way to get a feel for the difference between saving and borrowing is to ask how long you'd have to make monthly deposits of \$150.00 a month into a savings account to accrue \$4,800.00 I claim that it would take about 30 months. In other words, you have to make the same payments for an extra *year and a half* if you borrow the money. For short term investing of this type, safety is more important than yield. The last problem shows that even a very high yield shortens the term by only a couple of months. But to get a high yield you'll have to risk substantial losses and if these occur during such a short term they will delay reaching the goal considerably.

Before we go on, let's check the numbers in [EXAMPLE 5.9.14](#).

PROBLEM 5.9.15: Show that the monthly payment on a loan with a 48 month term at interest of 21.5% is \$149.93.

PROBLEM 5.9.16: Check the savings deposit amount given above. Of course, one ingredient is missing: the interest rate.

- i) Show that at a rate of 5.3% the monthly deposit needed to accumulate \$4,800.00 in 48 months is \$89.99.
- ii) Paradoxically, here the rate is not critical because the term of the amortization is so short. Show that at rates of 2% or 10% the monthly deposit you will need to make is within 10% of \$90.00
- iii) Compare the potential plusses and minuses of various kinds of investments for this kind of a short-term savings goal. What would you choose?

PROBLEM 5.9.17:

- i) Show that at a rate of 5.3% a monthly deposit of \$150.00 for 30 months accumulates to \$4,800.00 to the nearest dollar.
- ii) Here the rate you get on your savings matters even less because the term of the amortization is even short. Show that at rates of 1% to 12%, the number of months you need to deposit \$150.00 to accumulate \$4,800.00 changes by no more than 2.

5.9 Things everyone should know

Did I hear an objection? I hope so. I have slurred over two issues in [EXAMPLE 5.9.14](#). Perhaps, *I* should have been a used car salesman. Let's discuss them in turn. The basic moral of the example is that you pay a very heavy penalty for borrowing to make a consumer goods purchase like that new car—one amortized over a term of a few years at a high interest rate—instead of saving for the purchase. The main source of this penalty is that high interest rate.

On the other hand, in [PROBLEM 5.9.16](#) and [PROBLEM 5.9.17](#) we checked that the amount and/or term for which you'd have to save are relatively little affected by the rate your savings earn. Why does the rate matter for the loan but not for the savings? It's *not* that there is some basic difference between savings and loans as the following problem lets you check.

PROBLEM 5.9.18:

- i) Check that the monthly payment needed to amortize a loan of \$4,800.00 over 48 months at an interest rate of 6% is \$112.73.
- ii) Show that at rates of 2% or 10% the monthly deposit you will need to make is within 10% of this amount.

What these problems illustrate is that for both savings and loan amortizations, when the *term is short* and the *interest rate is fairly low* (say less than 10%), then the deposit D is relatively insensitive to the the interest rate. It's basically the final sum (if you're saving) or the initial balance (if you're borrowing) divided by the number T of payments.

This is exactly correct if the interest rate is 0%. If you make T payments of $\$D$ and no interest is involved they amount to a sum of $\$S = T \cdot \D or pay off a loan of $\$B = T \cdot \D . In [EXAMPLE 5.9.14](#), this zero interest deposit D is just $\$100.00 = \frac{4800.00}{48}$. For each extra percent of interest paid, $\$D$ changes by a few dollars. i) and ii) suggest that D changes about \$2 for each percent change in the nominal rate. When we save D goes down because we're being paid the interest: at 5.3%, D went down about \$10 to about \$90 Likewise, D goes



5.9 Things everyone should know

up when we borrow and pay the interest: it went up about \$12 to about \$112 when we were charged 6% interest. But in both cases we stay close to the zero-interest deposit of \$100.00.




FIGURE 5.9.19: Payment D versus rate r in [EXAMPLE 5.9.14](#)

Pictorially what I am claiming is that if we plot D against the nominal rate r we'll see something that looks a lot like a straight line with slope 2. [FIGURE 5.9.19](#) is the proof. The graph curves up a bit—the slope is closer to 3 than to 2 at the right side—but at a glance looks straight. Saying the rate is low means we're on the left side of the graph where the deposit is close to \$100.00.

But this is *not* true if *either* of the italicized assumptions fails to hold. We have already seen lots of examples of how sensitive to interest rates the final sums of long term amortizations can be: the

5.9 Things everyone should know

answers to part ii) of [PROBLEM 5.6.19](#) are excellent ones. (By the way, the answers to part i) of this problem give another example of how over a short term this need not be so). In the example of the used car loan, the term is short but the interest rate on the loan is no longer “fairly low”. It’s big enough to make a substantial difference even over a short term. We have moved out to the right side of [FIGURE 5.9.19](#) where the deposit is no longer close to \$100.00.

Consumer loan rates—rates on credit card balances, loans for consumer goods like furniture, home electronics etc, and in general loans made by people in order to sell you something generally have rates in the range 18%–24%. Such rates are high enough to substantially affect the deposit D need to amortize a given balance or sum even over short terms. This applies whether saving or borrowing. The following problem shows, for example, that if you could earn 20% on your savings you could reduce the \$90.00 deposit in [PROBLEM 5.9.16](#) by more than 25% to about \$66 a month.

PROBLEM 5.9.20: Suppose you could find a savings account that earned 20% interest. Check that the monthly payment needed to amortize a loan of \$4,800.00 over 48 months would be \$66.07 and that depositing \$149.93 a month into this account you would accrue slightly over \$4,800.00 in 26 months.

So the first clarification that needs to be made about [EXAMPLE 5.9.14](#) is that it’s borrowing at high consumer loan interest rates which is pernicious not the act of borrowing per se. A good thing too, since we wouldn’t want to have to save to save for 25 years—meanwhile paying rent—to buy a house. (Note, however, that there *are* countries in which this is more or less what you have to do to buy a house). The difference in a mortgage is the substantially lower rate as you can confirm in the following problem which asks what would happen if mortgages had rates like consumer loans.

PROBLEM 5.9.21: What is the monthly payment on a mortgage with a balance of \$100,000.00 and an interest rate of 20% if the term is



5.9 Things everyone should know

- i) 30 years?
- ii) 15 years?

What's the second clarification? Suppose we ignore the high interest rate charged on consumer loans and imagine they are available at, say, 6%. Look at the answers to [PROBLEM 5.9.16](#) and [PROBLEM 5.9.15](#): for the same amount and term and for our hypothetical interest rate of 6%, the savings payment is about 20% less than the loan payment. Is this a second factor—besides the high rate on consumer loans—which argues for saving to buy rather than borrowing to buy? No! The reason is a key fact which I ignored all through [EXAMPLE 5.9.14](#). You've probably spotted it long ago. The borrower has that car *now* and through the four years of the loan. The saver has to wait until the end of the 4 year period to buy the car. The use of the car during those intervening 4 years, like the use of money over a period of time has a value and the higher loan payment reflects the fact that the borrower acquires this use while the saver does not.

As a consequence, all my statements as to how much more it was costing the borrower than the saver are exaggerated: they all failed to take account of the fact that what the borrower is buying—a car today—is worth somewhat more than what the saver is buying—a car four years from now—and he should expect to pay somewhat more. My basic point, however, stands. The lion's share of the premium that the borrower pays is attributable to the high interest rate charged and not to the greater value of what is purchased.

Here's final striking way to view this difference. Suppose that both the borrower and the saver buy a \$4,800.00 car every 4 years. The borrower gets his first car today, pays off the loan at \$150.00 a month over 4 years, and repeats this cycle. The saver takes the bus for 4 years and puts away \$90.00 a month for a car and \$60.00 a month for his retirement (both at 6% interest). He then pays cash for his car but continues to deposit \$90.00 a month so he can pay cash for his next car in 4 years and to save \$60.00 a month for his

5.9 Things everyone should know

retirement. What are the plusses and minuses? Both are out \$150.00 a month. For the first 4 years, the borrower has a car but the saver does not. Thereafter, they both drive the same car. Eight cars later the saver has a retirement account with over \$90,000.00 in it. The borrower has nothing

PROBLEM 5.9.22: Show that after 36 years of making monthly payments of \$60 a month into a retirement account which earns 6% interest, you'll have a sum of \$91,495.13.

Here are a few more problems to illustrate the principles in this discussion.

“No interest and no payments until 2010!” This, or something like it, is a come-on often used in selling electronics, furniture etc. In the next problem, we analyze how such offers work and what their effect is. Payments on these loans generally come at the start of each period but we'll pretend they come at the end to avoid having to introduce more formulas.

PROBLEM 5.9.23: Let's suppose that what you want is to buy that big-screen TV for the 2009 Superbowl. Scanning the newspaper on January 15, 2009, you see an ad for that big screen TV you been hankering for, “This weekend only: \$2,499.99 or just 36 easy payments of \$99.99 a month. Plus no interest and no payments until 2010!”. In [PROBLEM 5.10.22](#), we'll see that the hidden interest rate in this offer is 25.45%.

i) Check that this is the correct rate. Then check that at 5.3% interest, you'd have to put away about \$64.22 a month for the same 36 month period.

ii) What does “No interest and no payments until 2010!” mean. If you think it means that you can make the first of your 36 \$99.99 payments on January 1, 2010 go to the back of the class. If you look at the small print in the loan contract, you'll see something like the following language, taken verbatim from the fine print of a promotion of this type.



5.9 Things everyone should know

“FINANCE CHARGES accrue on a promotional purchase from the date of purchase and all accrued FINANCE CHARGES for the entire promotional period will be added to your account if the purchase is not paid in full by the end of the promotional period or if you default under your card agreement.”

What that means, in plain English, is that you are not obliged to pay anything until January 1, 2010. Second, you may, if you can and wish, pay off the *entire* \$2,499.00 before January 1, 2010 and you will not owe the store anything. But, on January 1, 2010, you're not *obliged* to do anything but make that first \$99.99 payment. So far so good. Third—and here's the kicker—on January 1, 2010, the store will add 11.5 months interest (from January 15, 2009 to January 1, 2010) on whatever part of the balance you have not paid off to that balance. Suppose that you're like most people and pay off nothing before January 1, 2010. Show that your outstanding balance on January 1, 2010 will be \$3,183.88.

iii) Show that it will take you not 36 but 54 easy monthly payments of \$99.99 to pay off this debt at 25.45%. You'll still be paying for that TV during the NBA Finals of 2014!

iv) How much would you have to save each month in an account which earned 5.3% interest to accumulate \$2,499.99 in 54 months?

In fact, if you just put \$99.99 under your mattress for 54 months, you'd have enough cash to buy that TV and *another* just like it for the bedroom, with enough left over to host a major Superbowl party!

PROBLEM 5.9.24: You lose your job and to cover your expenses while looking for a new one use a credit card which charges a nominal rate of 23.65% on outstanding balances to pay most of your bills. By the time you find a new job get your financial life sufficiently organized to confront your outstanding balance, it has run up to \$17,500.00 You cut up the card and go to work to pay this off.

i) Show that when you start the additional interest you owe each month on the debt is \$344.89.

5.10 The hunt-and-peck method

- ii) Show that if you pay off this debt at \$500.00 a month it will take you about 60 months to amortize it, if you can only afford \$400.00 a month it will take 102 months (that's almost 9 *years*).
- iii) The minimum payment on a credit card is usually 2% of the outstanding balance. That's \$350.00 a month for your balance of \$17,500.00. Making the minimum payment, show that it will take you more than 18 years (!) to pay off your card.
- iv) If you can't afford to pay more than \$350.00 a month, how much more debt could you assume before you'd never be able to repay?

5.10 The hunt-and-peck method

A lengthy example

If you are trying to evaluate a consumer loan—say in a newspaper ad, you will probably not have the contract and the rate in front of you. Is there some way to estimate the interest rate? At first, it seems not. If we write down what we know about the car loan in [EXAMPLE 5.9.14](#) the [PRESENT AMORTIZATION FORMULA 5.8.1](#) $B = D \left(\frac{1 - (1+p)^{-T}}{p} \right)$ becomes

$$(5.10.1) \quad \$4,800.00 = \$149.93 \left(\frac{1 - (1+p)^{-48}}{p} \right).$$

There is a problem here: the periodic rate occurs in two places in the equation and there is no obvious way to isolate it and solve. Fine, as mathematicians we should look for some non-obvious way to get a formula for p and hence r in terms of B , D and T .

The bad news is that no one has found one. This is actually pretty typical. Most equations are impossible to solve *exactly*, that is, by giving a formula for the desired quantity in terms of the others in the

5.10 The hunt-and-peck method

equation. This may come as news to many of you. Your math texts over the years have probably given you a lot of equations to solve exactly. What the authors of your texts probably did not emphasize was that these equations were carefully selected from the very few which *can* be solved exactly.

The good news is that, in a problem like finding the interest rate where we do not really want a *formula* or exact solution for p but an *approximate value* or approximate solution, we can often substitute a bit of perspiration for a lack of inspiration to solve problems of this type. The basic idea is pretty simple. Make a reasonable guess for p , plug it into the [PRESENT AMORTIZATION FORMULA 5.8.1](#) and by comparing the two sides decide whether the guess is too big or too small. This lets us make a *better* guess for p and we then repeat the process until we know p as accurately as we wish.

I call this the [HUNT AND PECK METHOD 5.10.18](#). The fancy word for what we are doing is *iterating* (which just means repeating). Of course, all this guessing, plugging in and comparing is a fair bit of work. But by thinking a bit about what we see, we can greatly reduce the number of guesses we need. In effect, we substitute thought for calculation and that means we are still being mathematicians. Since most equations can *only* be solved by hunting and pecking, mathematicians have thought about better ways to hunt and peck quite a bit. The idea behind the particular method we'll use goes back to the great Isaac Newton.

Before we formalize the [HUNT AND PECK METHOD 5.10.18](#), let's look at how we can use it to solve equation [\(5.10.1\)](#) for p . This discussion is going to be a rather long one because we'll eventually come up with 3 separate methods, each with its own plusses and minusses. So get comfortable before we start, and just try to follow the flow of the main ideas. When we've completely discussed, the examples I'll lay out a detailed but concise method for attacking this and other similar equations in [HUNT AND PECK METHOD 5.10.18](#).



5.10 The hunt-and-peck method

Let's say our goal is to know the nominal rate r on the loan to the nearest tenth of a percent. This is more accuracy than you'd need in real life to evaluate the loan but it will illustrate the principle that with enough work we can get pretty much whatever accuracy we need. Since $p = \frac{0.01 \cdot r}{12} = \frac{0.01r}{12}$, to know r to within a tenth, it seems like we need to know p to about 4 places. This is sometimes a lie as the next problem below shows: so for safety I will guess p to 5 places and then compute r as $1,200 \cdot p$ rounded to the nearest tenth.

PROBLEM 5.10.2: If $r = 12.3\%$, then $p = \frac{r}{1200} = 0.01025$. Show that *no* rounding of p to 4 places gives the correct value of r . How likely is this problem to occur?

OK, let's review the plan. Make a first guess. Then, keep making better guesses until one of them is so good that it's the answer we're after. The details will take some time because repeated guessing means repeated calculation. But the ideas are very simple. The only trick is in making sure we improve our current best guess. As long we do this, we'll get a good answer eventually.

How do we come up with a first guess? Often information we already have about the problem tells us a good first guess. If not, we have two options. The first is to use an approximation, for example, the [PRESENT AMORTIZATION-SIMPLE INTEREST APPROXIMATION 5.8.9](#). Remember that this is not so much a formula as an idea: the average balance of an amortized loan over its term is about half the initial balance (we saw in the preceding subsection how far for accurate this idea is but it still gives a useful first approximation). Multiplying this average balance by the term of the loan and the periodic interest rate— $\frac{B}{2} \cdot T \cdot p = \frac{B}{2} \cdot \gamma \cdot \frac{r}{100}$ —should give us a simple interest estimate of all the interest in the loan. Another way to estimate this total interest is to subtract the starting loan balance from the total value of all payments getting $T \cdot D - B$ (this time ignoring the fact that these amounts live at different times and so cannot properly be merged like this). That is, we should have an approximate equality

5.10 The hunt-and-peck method

$$\frac{1}{2}B \cdot T \cdot p \simeq T \cdot D - B.$$

PROBLEM 5.10.3: The actual [PRESENT AMORTIZATION-SIMPLE INTEREST APPROXIMATION 5.8.9](#) was $T \cdot D \simeq B \left(1 + 0.01r \cdot \frac{y}{2}\right)$. Check that substituting $r = 100 \cdot m \cdot p$ and $y = \frac{T}{m}$ in this approximation gives the approximation $\frac{1}{2}B \cdot T \cdot p \simeq T \cdot D - B$ above.

This may be pretty crude but it is better than nothing and gives us an equation we *can* solve for p to get started. First, we can isolate p by dividing by $B \cdot T$ to find

$$p \simeq 2 \frac{T \cdot D - B}{B \cdot T} = 2 \left(\frac{D}{B} - \frac{1}{T} \right)$$

Then, we just plus in to obtain the starting guess

$$p \simeq \left(\frac{149.93}{4,800} - \frac{1}{48} \right) = 0.02080416667 = 0.02080.$$

This is one situation in which we need to work with a periodic rate p that is *not* given by the [INTEREST RATE CONVERSION FORMULA 5.1.10](#) $P = \frac{0.01 \cdot r}{m}$. In addition, we're trying to *guess* p approximately. This means that, for once, there's no harm in rounding p and above I have rounded it to 5 places since that's the accuracy I am going to work with.

Let's note that $p = 0.02080$ corresponds to $r = 25.0\%$ (I have rounded r to the nearest tenth of a percent) so my first guess is not too bad since in [EXAMPLE 5.9.14](#) we used a rate of 21.5%. I am off by 3.5%. Of course, in real life, if we knew r we wouldn't be wasting time trying to find p . My point is that the simple interest approximation gives us a decent first guess.

The bad news is that I had to come with a new version of the [PRESENT AMORTIZATION-SIMPLE INTEREST APPROXIMATION 5.8.9](#) to get this guess. The good news is that you won't need to learn this variant. A good initial guess is a nice thing to have—the better our first guess for p is, the less work we'll have to do to arrive at an accurate final value of p —but we'll soon see that with a bit more elbow grease we can make do with just about any first guess.

5.10 The hunt-and-peck method

This brings us to option two for making a first guess for p : *Wing it!* That is, use whatever you intuition or knowledge you have about the problem and make the best first guess you can based on it. In many of the problems, I'll actually supply you with a first guess, not because you couldn't find p without it, but just to save you some work and keep the focus on the main idea of the [HUNT AND PECK METHOD 5.10.18](#) which is to keep *improving* your latest guess. So let's just take the starting point of $r = 25.0\%$ or $p = 0.02080$ as given.

Now the real work begins. I plug this guess into the [PRESENT AMORTIZATION FORMULA 5.8.1](#) and compare the two sides. We find

$$\$4,800.00 \simeq 149.93 \left(\frac{1 - (1 + 0.02080)^{-48}}{0.02080} \right) = \$4,524.85.$$

Let me emphasize again that, because we are trying to guess p , we have no alternative to plugging in the approximation 0.02080 above—and [PERIODIC RATE RULE 5.1.11](#) does not apply. I rounded the trial balance to the penny to avoid a sea of decimals and I'll keep doing so in this section to keep it readable.

We can note with satisfaction that we are not too far off: the two sides differ by a few hundred dollars which confirms that our guess for p is in the ballpark. Now comes the first key question? Is $p = 0.02080$ too big or too small? Looking at the quantity $149.93 \left(\frac{1 - (1+p)^{-48}}{p} \right)$, it is not obvious whether it gets bigger or smaller as p gets bigger.

PROBLEM 5.10.4:

- Show, by plugging in the values $p = 0.1$, $p = 0.2$ and $p = 0.3$, that both the numerator and denominator get bigger as p gets bigger.
- Consider the fractions $\frac{1+p}{1+p^4}$ and $\frac{1+p^3}{1+p^2}$. Show, by plugging in the same values that all the numerators and denominators individually grow when p grows but that first *ratio* decreases while the second increases. In other words, part i) tells us nothing about how the ratio $\frac{1 - (1+p)^{-48}}{p}$ changes as p increases



5.10 The hunt-and-peck method

Fortunately, if we think for a moment about the *loans* which the **PRESENT AMORTIZATION FORMULA 5.8.1** describes, rather than about the formula itself, the answer becomes clear. *For a fixed loan balance B and term T , the higher the interest rate the larger the payment D . Turning this around, for a fixed loan payment D and term T , the higher the interest rate the smaller the balance B which can be amortized.* This is the key idea which lets us keep making better guesses.

PROBLEM 5.10.5:

- i) Confirm the first statement above by computing the payment on a \$100,000.00 mortgage with a 20 year term at interest rates of
 - a. 6%.
 - b. 7.5%.
 - c. 9%.
- ii) Confirm the second statement above by computing the balance on a mortgage with a payment $D = \$600.00$ and a 15 year term at interest rates of
 - a. 6%.
 - b. 7.5%.
 - c. 9%.

OK! We are looking for the periodic interest rate which lets a payment of \$149.93 amortize a loan of \$4,800.00 over 48 months. Since with $p = 0.02080$, we can only amortize \$4,524.85 we need a *smaller* rate p . How much smaller? It looks a like a few percent should do it but exactly how many percent is not clear. If it was, we wouldn't be hunting-and-pecking. But the beauty of the method, is that at this stage any reasonable second guess will do. Let's try $r = 20\%$ or $p = 0.01\dot{6} \simeq 0.01667$. Then we recompute

$$149.93 \left(\frac{1 - (1 + 0.01667)^{-48}}{0.01667} \right) = \$4,926.64.$$

Now we are almost home. First, the new balance of \$4,926.64 is high which means our second guess for p is low. The first key point to



5.10 The hunt-and-peck method

note is that we now know that p is definitely bigger than 0.01667 and smaller than 0.02080. The fancy way to say this is that our two guesses **bracket** the true value of p . I'll call the smallest range or interval in which we know the rate we're looking for lies the **interval of uncertainty**: here it's the interval $[.01667, 0.02080]$ in terms of p , or the interval $[20\%, 25\%]$ in terms of r .

I should warn you that, in working later problems, your second guess may not bracket the solution as mine did. This is not a serious problem. I'll explain how to handle it at the end of this example.

The second key point is that to 2 decimal places both guesses give $p = 0.02$. If we can just keep improving our guesses—that is, if we can shrink the interval of uncertainty one more place—we'll eventually find two guesses for p which both correspond, after rounding, to the desired value $r = 21.5\%$ and we'll be done.

I'm going to provide three methods which achieve our goal of shrinking the interval of uncertainty, or **buffing** a solution. They provide a trade-off between the *number* of extra guesses we'll need to make to get the accuracy we're after (wear and tear on your calculator) and the *thought* required to come up with each guess (wear and tear on your brain).

Let's take the simplest method first. The *average* of two guesses which bracket the value we're provides a new guess which cuts the interval of uncertainty in half. Here the average of 0.01667 and 0.02080 is 0.018735. To keep our guesses for p to 5 places, I'll have to round to 0.01874. Plugging in this guess for p , we recompute

$$149.93 \left(\frac{1 - (1 + 0.01874)^{-48}}{0.01874} \right) = \$4,718.99.$$

Since this amount is *less* than \$4,800.00 we know that the guess $p = 0.01874$ is *more* than the true value so our new interval of uncertainty for p is now $[.01667, .01874]$ which is *half* as big as our old one and gives us an interval of uncertainty for r of $[20.0\%, 22.5\%]$ (to the nearest tenth of a percent).



5.10 The hunt-and-peck method

A moment's thought shows that this will always be the case. Our two best guesses are the endpoints of the interval of uncertainty so their average is the midpoint or bisector of this interval. Our guess has bisected the interval of uncertainty—cut into two equal parts each half as big as the original—and the value we are after must lie either in the lower or in the upper of the two halves. For this reason, the method of guessing the average of our two best guesses so far is usually called the **bisection method**.

The bisection method is the tortoise of our trio of methods: slow but steady. Each extra guess shrinks the interval of uncertainty by a factor 2, 2 guesses shrink it by a factor $4 = 2 \cdot 2$, 3 guesses shrink it by a factor $8 = 2 \cdot 2 \cdot 2$, 4 guesses shrink it by a factor $16 = 2 \cdot 2 \cdot 2 \cdot 2$ and so on. To get one extra decimal place of accuracy, we need to shrink the interval of uncertainty by a factor 10 which will take either 3 and 4 guesses since 10 is between 8 and 16.

To get 2 more places we need to shrink the interval of uncertainty by a factor of 100 which will take either 6 or 7 guesses since 100 is between $2^6 = 64$ and $2^7 = 128$. I could continue but, fortunately, we do not need to know exactly how many guesses will be needed. We just keep guessing until our interval of uncertainty for r is small enough that we know r to the nearest 0.1%. Then we're done. [TABLE 5.10.6](#) summarizes what happens when carry this out.

guess	low	high	high	low	average	new	new intervals of uncertainty for	
	p	B	p	B	p	B	p	r
1	.01667	4,926.64	.02080	4,524.85	.01874	4,718.99	[.01667, .01874]	[20.0%, 22.5%]
2	.01667	4,926.64	.01874	4,718.99	.01771	4,820.68	[.01771, .01874]	[21.3%, 22.5%]
3	.01771	4,820.68	.01874	4,718.99	.01823	4,768.95	[.01771, .01823]	[21.3%, 21.9%]
4	.01771	4,820.68	.01823	4,768.95	.01797	4,794.71	[.01771, .01797]	[21.3%, 21.6%]
5	.01784	4,820.68	.01797	4,794.71	.01791	4,800.69	[.01791, .01797]	[21.5%, 21.6%]
6	.01791	4,800.69	.01797	4,794.71	.01794	4,797.70	[.01791, .01794]	[21.5%, 21.5%]

TABLE 5.10.6: FINDING p BY THE BISECTION METHOD

5.10 The hunt-and-peck method

A few remarks are in order. We stop after the sixth bisection because, after it, we know the value of r to the desired accuracy of 0.1%: we have “discovered” the rate 21.5% which we knew all along. Note, also, that before we started we knew p to 2 places and [PROBLEM 5.10.2](#) says that we’ll probably need to know p to 4 places before we quit. Thus, the 6 bisections it took is just what I predicted before the table would be required. Let me repeat, for emphasis, however, that knowing this was an unneeded luxury. The reason we can stop after 6 bisections is that by then we know r to the nearest 0.1%.

Finally, we could have eliminated much of the table and still have carried out the calculations. There was no real need to repeatedly list the “old” p ’s and B ’s. I did this just to make it clearer how the method was working. The only new values in each line of the table are the new p which is the average of our two best guesses so far, and the corresponding r and B given by plugging this into the [TERM CONVERSION FORMULA 5.1.13](#) and the [PRESENT AMORTIZATION-SIMPLE INTEREST APPROXIMATION 5.8.9](#).

Now let’s ask a mathematician’s question. Is there any way we to get our answer with *fewer* guesses (saving work) by making *better* guesses (being cleverer)? I hope you’ll realize by now that if I ask this, the answer must be yes. Moreover, only a modest amount of cleverness is called for. To see what the extra idea is, let’s go back and look a bit more carefully at what we knew before we started bisection. We knew that the first guess $p = 0.02080$ produces a trial balance $B = \$4,524.85$, the second guess $p = 0.01667$ produces a trial balance $B = \$4,926.64$ and hence that the value of p we’re after—which would produce a balance $B = \$4,800.00$ —is *between* these two guesses.

The key observation to make is that the second balance \$4,926.64 is a lot *closer* to \$4,800.00 than the first balance \$4,524.85. They are off by \$126.64 and \$275.15 respectively. So we should expect that the second guess 0.01667 is a lot closer to the actual p than the



5.10 The hunt-and-peck method

first guess 0.02080. Instead of blindly taking as our third guess the average of the first two (i.e., bisecting), we can try to *eyeball* a better guess. I call this the **eyeball** method of improving our guesses.

How do we eyeball? Well, the first balance is off by a bit more than twice as much as the second. I'll try to make a third guess for p guess that's about twice as close to the second guess as to the first. A nice round guess that seems to do this is $p = 0.01800$. Why did I come up with that number? By eyeball! I guessed! By this I mean that I have no formula in my mind—I'm not doing anything definite like taking an average as I do when bisecting. I just picked a simple decimal that *looks* like it's about right.

I almost certainly won't nail down the right p with this guess but I don't really need to. All that really matters is that my guess generally turns out *better than guessing the average*. If it does, I'll find p with fewer guesses, hence correspondingly less work. [TABLE 5.10.7](#) below shows what happens if I carry this out.

guess	low p	high B	high p	low B	eyeball p	new B	new intervals of uncertainty for	
							p	r
1	.01667	4,926.64	.02080	4,524.85	.01800	4,791.73	[.01667, .01800]	[20.0%, 21.6%]
2	.01667	4,926.64	.01800	4,791.73	.01790	4,801.69	[.01790, .01800]	[21.5%, 21.6%]
3	.01790	4,801.69	.01800	4,791.73	.01792	4,799.69	[.01790, .01792]	[21.5%, 21.5%]

TABLE 5.10.7: FINDING p BY THE EYEBALL METHOD

First point: I saved half the work as I'd hoped. Second point. There's nothing up my sleeve. I used no magic recipe to make those guesses of $p = 0.01790$ and $p = 0.01792$. I eyeballed a best guess but in both cases I really did just *guess*. For the first, I just noticed that the trial balance for my eyeball guess of $p = 0.01800$ was only off by about \$8 so I figured p had to be a *lot* closer to 0.01800 than to 0.01667 and I picked 0.01790 as a “round” guess that did this. This p gives a balance in error by only about \$2 and so I guessed $p = 0.01792$

5.10 The hunt-and-peck method

because it differs from 0.01790 and 0.01800 by 0.00002 and 0.00008 respectively, which is roughly proportional to the errors of \$2 and \$8 in the balances.

But, if I'd been a bit less clever, I'd still have saved steps. Once again, all I have to do to show a profit is be able to make a guess that's better than the average of my 2 best guesses so far. Indeed, in some ways these guesses are easier to make because I do not need to find an average in the calculator, I just look at the errors in the balances and try to spilt the differences in the p 's correspondingly.

However, in teaching this method over the years, I have learned that the idea of *just guessing* causes many students to freeze up. They'd prefer to use a method like the **bisection method** in which you have a plug-and-chug formula—the average—for each guess even if it means doing a *lot* more work.

Linear Interpolation

This subsection is optional. It describes a third method of making improved guesses that lets us have our cake and eat it too—it combines the benefits of both our existing ones. But you can skip to the next section and just ignore references to the **LINEAR INTERPOLATION FORMULA 5.10.9** if you just want to get started working problems. This turns out to be a *formula* for eyeballing the next p ! Choosing successive p 's by using the interpolation formula lets us enjoy the speed of the eyeball method—in fact, it's even faster—with no guesswork.

The idea is very simple. What was I trying to do when I eyeballed? My goal was to have the differences between by new guess for p and my prior low and high guesses be roughly proportional to the differences between the correct balance B and the prior high and low balances. (Why did I reverse high and low here?) To hell with roughly, then. The **LINEAR INTERPOLATION FORMULA 5.10.9** is just a



5.10 The hunt-and-peck method

formula which makes the differences between new p and my prior low and high guesses *exactly* proportional to the differences between the correct balance B and the prior high and low balances.

The formula really has nothing to do with rates and balances. Since we'll want to use interpolation to solve other kinds of equations later on, I'm going to digress for a moment to explain the simple but general idea before showing how to use the formula to solve our example yet again. The word linear in the name is the key and the picture below shows why.




FIGURE 5.10.8: Linear interpolation

In general we'll be trying to find the unknown value of some quantity like p which solves an equation, that is, produces a known value of some other quantity B . What we'll know are two prior guesses for p (which I've called p' and p'' in the picture) and the corresponding trial values of B (which I've called B' and B''). I've plotted these as two points (p', B') and (p'', B'') on the picture. I've deliberately drawn a picture where the right B is much closer to B'' than to B' and where B decreases as p increases but the calculation which follows works in all cases.

Another way to express the goal of finding the guess for p which makes the differences between my two prior guesses *exactly proportional* to the differences between the known balance B and the prior balances is to say that the point (p, B) lies on the line joining the points (p', B') and (p'', B'') . Yet another way to state this is:

5.10 The hunt-and-peck method

the slopes of the line segment joining (p, B) and (p', B') equals the slope of the line segment joining (p, B) and (p'', B'') . Using the fact that slope equals rise over run, this gives the slope equation

$$\frac{p - p'}{B - B'} = \frac{p - p''}{B - B''}.$$

This is an equation we *can* solve with a bit of algebra. Clearing denominators gives

$$(p - p')(B - B'') = (p - p'')(B - B')$$

and then expanding gives

$$pB - pB'' - p'B + p'B'' = pB - pB' - p''B + p''B'.$$

Isolating all the p terms to the left side gives

$$p(B'' - B') = p'(B - B'') - p''(B - B')$$

and dividing by $(B'' - B')$ finally yields the first equality below:

LINEAR INTERPOLATION FORMULA 5.10.9: *If we are trying to find a rate p which corresponds to a balance B and we have two guesses p' and p'' which produce respective trial balances B' and B'' , the guess p which makes the differences from $p - p'$ and $p - p''$ proportional to the differences $B - B'$ and $B - B''$ is given by the linear interpolation formula*

$$p = \frac{p'(B - B'') - p''(B - B')}{B'' - B'} \quad \text{or} \quad p = \frac{p''(B - B') - p'(B - B'')}{B' - B''}.$$

I've given two versions of the formula but they are equivalent as the problem below shows. The upshot is that we don't need to worry which of our two previous guesses to take for p' and which for p'' ; we get the same new guess either way. (We do, however, have to match each balance with the corresponding periodic rate.)

PROBLEM 5.10.10: Obtain the second formula above by isolating all the p terms to the *right* side after expanding the slope equation.

5.10 The hunt-and-peck method

Now let's see how to use the formula in practice. In our example, $(p', B') = (0.02080, \$4,524.85)$ and $(p'', B'') = (0.01667, \$4,926.64)$. So the [LINEAR INTERPOLATION FORMULA 5.10.9](#) says we should try

$$\begin{aligned} p &= \frac{p'(B - B'') - p''(B - B')}{B'' - B'} \\ &= \frac{0.02080(4,800.00 - 4,926.64) - 0.01667(4,800.00 - 4,524.85)}{4,926.64 - 4,524.85} \\ &= 0.01797173275. \end{aligned}$$

PROBLEM 5.10.11: Confirm that plugging into the second version of the [LINEAR INTERPOLATION FORMULA 5.10.9](#) also yields the value $p = 0.01797173275$.

As usual, we'll round this guess to $p = 0.01797$. From here on, everything proceeds just as with the other two methods. We calculate the trial balance B corresponding to our new guess for p and repeat. I've called this method the **linear interpolation** method after the formula it uses. [TABLE 5.10.12](#) summarizes the next two repetitions.

guess	low p	high B	high p	low B	new p	new B	new intervals of uncertainty for	
							p	r
1	.01667	4,926.64	.02080	4,524.85	.01797	4,794.71	[.01667, .01797]	[20.0%, 21.6%]
2	.01667	4,926.64	.01797	4,794.54	.01792	4,799.69	[.01667, .01792]	[20.0%, 21.6%]
3	.01667	4,926.64	.01792	4,799.69	.01792	4,799.69	[.01667, .01792]	[20.0%, 21.6%]

TABLE 5.10.12: FINDING p BY THE LINEAR INTERPOLATION METHOD

What happened in that third line? We apparently made no progress so there's not much point in repeating any more. The problem comes when we round p to 5 places in the third line. An *unrounded value* for the p given by the [LINEAR INTERPOLATION FORMULA 5.10.9](#) in the third line is $p = 0.01791694762$ and this value produces a trial balance $B = 4,799.996521$ which is \$4,800.00 to the nearest penny allowing us to conclude that the corresponding r which is

5.10 The hunt-and-peck method

21.50033714% or 21.5% to the nearest 0.1% is what we're after. Once again, the moral is **never round anything except a final answer**.

In fact, if we look a bit more closely at the second line we can see that blindly following our procedure was not very smart. The two guesses 0.01797 and 0.01792 give balances off by less than \$6 and less than \$1 respectively.

Why not use these values even though they *both* give balances on the *low* side of \$4,800.00. This was something we couldn't permit when using the **bisection** method because bisecting an interval that does *not* contain the answer will only give us a smaller interval that also does not contain the answer. But this objection does not apply when using linear interpolation to make guesses. Of course, the picture in this case is a bit different from that shown in [FIGURE 5.10.8](#). Here the point we're looking for no longer lies between the two points we have already found.

But, so what? We can still equate slopes and once we do so we'll still get the [LINEAR INTERPOLATION FORMULA 5.10.9](#) (because the algebra we did never used the fact that the "answer" was between the guesses). So rejecting these two p 's because both balances are a bit low and instead using the rate 0.01667 which gives a balance which is high but off by a whole \$126 wasn't very clever. Here the moral is that *it never hurts and often helps to think a bit when making a calculation*. If we're willing to do so, we can give up the security of having the solution bracketed and, as a bonus, get an accurate answer much more quickly.

Just to confirm these ideas let's redo the calculation in the third line *without rounding* and *using the guesses 0.01797 and 0.01792*.

$$\begin{aligned} p &= \frac{p'(B - B'') - p''(B - B')}{B'' - B'} \\ &= \frac{0.01797(4,800.00 - 4,794.71) - 0.01792(4,800.00 - 4,799.69)}{4,794.54 - 4,799.69} \\ &= 0.01791688755. \end{aligned}$$

5.10 The hunt-and-peck method

and the corresponding trial balance is \$4,800.002520, correct to the penny.

PROBLEM 5.10.13: OK, but we *were* off by a quarter penny. Why? *Again*, mainly due to rounding error, this time the error caused by rounding our trial balances to the nearest penny! This problem asks you to confirm this.

- i) Show that the unrounded balance which can be amortized at a periodic rate of $p = 0.01797$ is \$4,794.714028 and at a rate of $p = 0.01792$ is \$4,799.692301.
- ii) Show that the interpolation formula applies to these values predicts a value of $p = 0.01791690958$ and that the trial balance corresponding to this p is $B = 4,800.000314$, which is accurate to three-hundredths of a penny!

When good guess go bad

Hang on, we're almost done. That's pretty much everything we'll need to know about making improved guesses. But before, I write down a detailed procedure for solving equations by hunting and pecking, I need to go back and address two problems which might arise when making your initial guesses. I ignored these at the time because I wanted to show you that we really could come up with accurate values of p by guessing. I hope you're convinced. However, all three of our methods for improving guesses require us to come up with the first two guesses and I need to say a few words about problems that might arise with these.

Maybe you think, I was lucky in making that second guess $r = 20\%$ or $p = 0.01667$ which I did pretty much blind. But when I said any reasonable second guess will do, I meant it. Suppose I had decided to drop down only from $r = 25\%$ to $r = 24\%$ or $p = 0.0200$. Then, I'd find a trial balance

$$B = 149.93 \left(\frac{1 - (1 + 0.0200)^{-48}}{0.0200} \right) = \$4,598.82.$$



5.10 The hunt-and-peck method

Now I have two guesses for p which are *both* too high, i.e., give trial balances which are too *low*. One thing I definitely can *not* do with these guesses is start bisection. The average of two guesses which are both high (or both low) will be worse than the best current guess. We don't need to carry out a detailed calculation to see this. The average of 24% and 25% is 24.5% which is farther from the right r of 21.5% than 24%.

There are two ways to handle this problem. The simplest approach is just to making blind guesses until we find one on the “other” (here the low) side of the solution—which we'll recognize by the fact that it gives a trial balance on the “other” side of the known balance (here a balance *greater* than \$4,800).

As usual, we can save work by thinking a bit. Instead of guessing, next an r of 23%, then 22% and finally 21% before reaching the “other” side, we could notice that when we dropped our guess 1%, we raised the balance by about \$75. So we'd expect that to raise the balance another \$200 or so, we'd need to drop r about 3% or so. That is, we should jump to 21% right away. While not strictly necessary, this would save 2 of the 3 extra guesses—and 2 of the 3 extra trips through the [PRESENT AMORTIZATION-SIMPLE INTEREST APPROXIMATION 5.8.9](#).

But doesn't this speedup seem a bit familiar? What I really did to decide I should drop my guess for r by 3% was make an *eyeball* best guess. In applying the eyeball method, we only used it to make improved guesses between guesses which *already* bracketed the correct value but it can work equally well to improve any two guesses.

We already used linear interpolation approach to do the same thing above. There, we were improving two very good guesses (0.01797 and 0.01792) but the method *generally* works just as well on not-so-good guesses. Let's try it. Our two guesses for r —25% and 24%—correspond to p 's and B 's of 0.02080 and \$4,524.85 and 0.02000 and \$4,598.82 respectively. Plugging the p 's into the [LINEAR INTER-](#)



5.10 The hunt-and-peck method

INTERPOLATION FORMULA 5.10.9, gives an new guess of $p = 0.017824$ and a new trial balance of \$4,809.25. Ka-ching!

PROBLEM 5.10.14: Confirm these figures above by carrying out the interpolation and amortization calculations involved yourself.

Why is that ominous word *generally* in bold type above? Well, the bad news about the **LINEAR INTERPOLATION FORMULA 5.10.9** is that there are cases where it gives you an guess which is *worse* than the two you started with. Moreover, there's no easy way to predict in advance whether this will happen. The good news is that such problems are very rare and I have chosen the examples in this course to avoid such difficulties. Still, if you ever decide to try this formula on a real-life problem—and, as we'll soon see, it is very widely applicable—you should just keep the possibility that things will go wrong in the back of your mind. In the method, outlined below I ask you, as a safety measure, to always *bracket* the solution with guesses before trying to make really accurate guesses. If interpolation ever takes you “outside the brackets”, something has gone wrong. If so, just fall back on a less efficient but more sure method like bisection.

One last point. You're taking the final exam in the course and you've been asked to solve for an interest rate by the **HUNT AND PECK METHOD 5.10.18**. You make a first guess for r and p . So far so good. But when you go to make the second guess you find you've forgotten whether raising the rate lowers the balance or vice-versa. (This is not uncommon, in my experience, because later in this section we'll work with savings amortizations where the rule is that rates and amounts go up together and it's easy to confuse the two cases). Should you panic? Never! Knowing in advance whether you need to raise or lower a guess is another timesaving plus in using the **HUNT AND PECK METHOD 5.10.18** but it's basically a luxury you can do without in a pinch.

Suppose that after guessing $r = 25\%$ and getting a trial balance of \$4,524.85, I went the *wrong* way and next guessed $r = 27\%$. Then,

5.10 The hunt-and-peck method

I'd compute a trial balance of \$4,373.39. "Whoops!", you say, "I'm further from the desired balance of \$4,800.00 than I was the first time." No problem. If your balance went the "wrong way", that just means your guess did too. You should have guessed a rate less than 25% not greater. Just do so and you're back on track.

The method

Finally, we are ready to summarize what we have learned from our example. First a few pieces of terminology.

BRACKETING A SOLUTION 5.10.15: *We say that two guesses **bracket** the solution to an equation if plugging one guess into the equation gives a left hand side greater than the right hand side and plugging the other in gives a left hand side less than the right hand side.*

BETTER AND WORSE GUESSES 5.10.16: *If the two sides of an equation are closer to being equal when the one guess is plugged in than they are when a second guess is plugged in, we say that the first guess **better guess** than a second and the second is a **worse guess** than the first.*

BUFFING A SOLUTION 5.10.17: *Given two guesses which bracket a solution, any method for repeatedly making improved guesses until a solution is found to the desired accuracy is called a method for **buffing** the solution.*

HUNT AND PECK METHOD 5.10.18:

Step 1: **Make a first guess.** Try to use any information you are given, or any ideas suggested by the problem to make the best guess you can. Plug in your guess and compute both sides of the equation you are trying to solve.

Step 2: **Make a second guess.** If you can use your knowledge of the problem to predict whether your first guess was high or



low, use this information to head in the “right” direction—towards the solution. If not, do not worry. Plug in your guess and compute both sides of the equation you are trying to solve. If your second guess went the “wrong” way (you’re further from agreement than you were the first time), repeat this step making a new guess in the “right” direction.

Step 3: Bracket the solution. If you’re lucky, your first two guesses will already bracket the solution. If not, you can look for bracketing guesses in three ways:

i) Brute force: make any new guess which lies on *opposite* side of your better guess to your worse guess. Repeat, if necessary, until your new guess and your better guess *do* bracket the solution.

ii) Eyeball: compare *how much closer* to equality the two sides of the equation are when you plug in the better guess than the worse guess. Use this to predict “by eye” *how far* from the better guess (on the side away from the worse guess) your new guess needs to be to bracket the solution. Repeat, if necessary, until your new guess and your better guess *do* bracket the solution.

iii) Interpolate: plug your two guesses into the [LINEAR INTERPOLATION FORMULA 5.10.9](#) to get a new guess. If you do not find a bracketing guess the first time you try this, fall back on one of the two other methods.

Usually, you’ll be able to bracket the solution by third guess at worst. If you do need to make more than two guesses, you can discard all but the best guess on each side of the solution before proceeding to step 4.

Step 4: Buff the solution Use one of the three methods below as often as is necessary to produce a new and improved guess for the solution. After each new guess, discard the old guess which lies on the *same* side of the solution as the new guess.

5.10 The hunt-and-peck method

Repeat until you have bracketed the solution to within the accuracy you are seeking.

i) **Bisection**: Use the average of the two solutions as the new guess. *opposite* to your worse guess. Repeat, if necessary, until your new guess and your better guess *do* bracket the solution.

ii) **Eyeball**: compare *how close* to equality the two sides of the equation are when you plug in the two bracketing guesses. Use this to predict “by eye” a new guess close to the solution and between the brackets.

iii) **Interpolating**: plug your two guesses into the [LINEAR INTERPOLATION FORMULA 5.10.9](#) to get a new guess.

I expect that many of you are a bit uncomfortable with the idea of a method which has so many variants. Please try to view this as a good rather than a bad thing: you can pick the method that’s easiest for you. In the next few problem, I *will* ask you to solve for few interest rates using each of the three basic approaches: **bisection**, **eyeball** and **interpolation**.

My goal in doing so is mainly to give you a feel for how each of these methods goes in practice. You’ll probably develop a preference for one of the three approaches in doing these three problems. Good! After you get through them, you can do the other problems which ask you to solve equations by hunting and pecking whichever way you like best.

Here are a few comments you might find helpful to keep in mind. The bisection method is the easiest—once you have brackets you just keep averaging your two best guesses—but it takes a lot of repetitions.

The interpolation method is very fast (the solutions to [PROBLEM 5.10.19](#) are good examples) but it has one big drawback: if you can’t remember the relatively messy [LINEAR INTERPOLATION FORMULA 5.10.9](#), you’re toast.

5.10 The hunt-and-peck method

Eyeballing is in the middle. It can be almost as fast as interpolating if you have a good eye and there's no need to remember any formula at all. It's my preference but I should say that if you don't find it easy in the examples which follow you should probably choose one of the other two.

Let's get to it! I have done the first problem below (without rounding) to give you a second model for each method. To save a few trees (or LEDs for those of you reading this online), I have not repeated any of the calculations using the [PRESENT AMORTIZATION FORMULA 5.8.1](#). Once I have a guess for p , I simply state what value of B comes out by plugging into the formula. Also, for each method, I have made different initial guesses to try to convince you that these don't really matter that much.

You'll probably have guessed that doing these problems is a fair bit of work. Don't worry, I'm not going to ask you to do very many. But I would like you to try out all three methods with care a few times.

PROBLEM 5.10.19: Use each of the three basic methods (bisection, eyeballing and interpolation) to find the nominal interest rate, to the nearest tenth of a per-cent, which is being charged on a loan with an initial balance of \$3,000.00

- i) if the term is 2 years and
 - a. the monthly payment is \$147.18.

Bisection Solution

Step 1: This time I'll just dive in and guess that $r = 18\%$ or $p = 0.015$ and I find $B = \$2,948.075058$.

Step 2: Our first guess is *high* since the balance is low. Lets try $r = 12\%$ or $p = 0.01$. This gives $B = \$3,126.601737$.

Step 3: We already have the solution bracketed since so there's nothing to do here.

Step 4: Now I'll start bisection. I'll just make a table of my new guesses for p and the corresponding r and B then make a few comments below it.

5.10 The hunt-and-peck method

p	r	B
0.00125000000	15.00000000%	\$3,035.474833
0.01375000000	16.50000000%	\$2,991.320447
0.01312500000	15.75000000%	\$3,013.282627
0.01343750000	16.12500000%	\$3,002.272962
0.01359375000	16.31250000%	\$2,996.789592
0.01351562500	16.21875000%	\$2,999.529491
0.01347656250	16.17187500%	\$3,000.900874

TABLE 5.10.20: PROBLEM 5.10.19 BY BISECTION

I started with $p = 0.125$ the average of my guess 0.01 and 0.015. Since this gave a B of 3,035.474833 which was too *high*, I discarded the *low* guess and averaged 0.125 and 0.15. And so on.

I stopped when my last two values if r both rounded to 16.2% since I then know that this must be the rate in the loan to the nearest 0.1%. This time I needed 7 repetitions, again fairly typical.

Eyeball Solution

Step 1: This time I'll guess that $r = 20\%$ or $p = 0.016666667$ and I find $B = \$2,891.790568$

Step 2: Our first guess is *high* since the balance is low. Lets try $r = 10\%$ or $p = 0.08333333333$. This gives $B = \$3,189.516306$.

Step 3: We already have the solution bracketed since so there's nothing to do here.

Step 4: Now I'll start eyeballing. My first guess here is quite a bit better: the first B is off by about \$108, the second by about \$190 almost twice as much. Let's try guessing twice as close to 20% as to 10%, i.e., $16\frac{2}{3}\% = 16.66666667$. In doing this problem, I kept guessing a new r and de-

5.10 The hunt-and-peck method

rived the corresponding p . From here on, I'll just give a table of my new guesses for r and the corresponding p and B and again make a few comments below it.

r	p	B
16.66666667%	0.01388888889	\$2,986.470943
16.0%	0.01333333333	\$3,005.936413
16.2%	0.01350000000	\$3,000.077896
16.24%	.0135333,333	\$2,998.908061

TABLE 5.10.21: PROBLEM 5.10.19 BY EYEBALLING

I was pretty sure that my third guess would be on to the nearest tenth from eyeballing the first two rows in the table. Since the trial balance was only off by 7 cents, I could reasonably have stopped there but just to be sure, I checked 16.24% gave a low balance (so bracketing the exact value with two r 's both rounding to 16.2%). Note that even though my first two guesses were much worse I did far fewer repetitions with this method and I avoided taking all those averages.

I'll interpolate *twice* since this method has the messiest formula in it.

First Interpolation Solution

- Step 1: This time I'll guess that $r = 15\%$ or $p = 0.0125$ and I find $B = \$3,035.474833$
- Step 2: Our first guess is *low* since the balance is high but not too low. Lets try $r = 17\%$ or $p = 0.01416666667$ This gives $B = \$2,976.805484$, low but also quite close.
- Step 3: Once again, my first two guesses bracket the solution.
- Step 4: Now, let's use the [LINEAR INTERPOLATION FORMULA 5.10.9](#) to buff these guesses. Plugging my first two

5.10 The hunt-and-peck method

guesses in, I get

$$\frac{0.0125(3,000.00 - 2,976.805484) - 0.01416666667(3,000.00 - 3,035.474833)}{2,976.805484 - 3,035.474833}$$

which gives $p = 0.01350776166$, then $r = 16.20931399\%$ and finally $B = \$2,999.805516$. One more round should do it. I discard my second guess to keep the solution bracketed and using my first and third I interpolate to find $p = 0.01350226692$, $r = 16.20272030\%$ and $B = \$2,999.998337$. The trial balance is off by a fraction of a penny and common sense says I've found r to the nearest 0.1% as 16.2%.

Second Interpolation Solution

Step 1: This time I'll guess that $r = 15\%$ or $p = 0.0125$ and I find $B = \$3,035.474833$

Step 2: Our first guess is *low* since the balance is high but not too low. Lets try $r = 16\%$ or $p = 0.01333333333$ This gives $B = \$3,005.936413$: we're almost there.

Step 3: However, we do *not* have the solution bracketed yet. Since I am going to interpolate anyway, let's use the [LINEAR INTERPOLATION FORMULA 5.10.9](#) to try to get a bracketing guess. Plugging my first two guesses this time, I get

$$\frac{0.0125(3,000.00 - 3,005.936413) - 0.01333333333(3,000.00 - 3,035.474833)}{3,005.936413 - 3,035.474833}$$

This time $p = 0.01350081049$ so $r = 16.20097259\%$ and I recompute $B = \$3,000.049361$.

I *still* don't have the solution bracketed but since I'm off by less than a nickel in the trial balance, common sense tells me I've nailed down r to within 0.1%. However, just to see how powerful this method really is you should interpolate one more time. When you do, you should obtain $p = 0.01350221473$, $r = 16.20265768\%$ and $B = \$3,000.000209$, off by two-hundredths of a cent.

b. the monthly payment is \$154.16. (A good starting guess here would be 16%. Why?)



5.10 The hunt-and-peck method

- ii) if the term is 4 years and
 - a. the monthly payment is \$89.07.
 - b. the monthly payment is \$83.49.

Try to come up with your own starting guesses for these parts.

Here are a few more practice examples. We will also apply the method further in the following subsections.

PROBLEM 5.10.22: In [PROBLEM 5.9.23](#) we discussed an ad for a big screen TV offering it “This weekend only: \$2,499.99 or just 36 easy payments of \$99.99 a month”. What is the hidden interest rate on the loan option? Here you might start with a nominal rate of 24%.

PROBLEM 5.10.23: A home equity loan is essentially a second mortgage on your home. If you default the home equity lender will get paid *after* the holder of your primary mortgage but, if you have sufficient equity in your home, the lender may feel confident of being able to recover his money eventually. For this reason such loans are easier to get than general consumer loans and usually have a somewhat lower rate. You get a letter in the mail from a home equity lender telling you that “You’re pre-qualified for a home equity loan \$10,000.00 with easy monthly payments of just \$140.00 a month for 10 years”. What interest rate will you be paying if you accept this offer? These loans tend to have rates a bit higher than mortgages so try starting at 10% here.

More About Linear Interpolation

Time for a brief break for some optional enrichment material. The rest of you can skip to the next subsection.

I hope though that you have been asking, “Where does the [LINEAR INTERPOLATION FORMULA 5.10.9](#) come from?”. The idea, remember, which can get lost in the complexity of the formula is simple. Given two guesses and their trial balances, we get a better guess by joining

5.10 The hunt-and-peck method

the corresponding points by a line segment and solving for the p value *on the line* whose B equals the desired balance. You might want to go back and stare at [FIGURE 5.10.8](#) again. Essentially, the formula amounts to assuming that: *if we graph rate versus balance what we'll see is a straight line.*

On the other hand, you can also go back and look at an actual graph of rate versus balance in [FIGURE 5.9.19](#). What you see is *somewhat* straight—that was my point in showing it there—but still visibly curved. Another example is shown in [FIGURE 5.10.24](#) below.




FIGURE 5.10.24: Graph of Balance versus Rate

Once again: *The graph is **not** straight!*. How can the [LINEAR INTERPOLATION FORMULA 5.10.9](#) give us such amazingly accurate answers if it is based on an assumption that is only very roughly true? The answer is one of the few really key ideas in calculus.

If we zoom in on a portion of the graph, what we see becomes straighter. This point is illustrated by the circle on the right showing a blownup image of a small part of the graph. Moreover, the more we zoom in the straighter what we see becomes. You can test this out by using the **Acrobat** magnifier tool to zoom in on any part of the graph. In effect, once we are interpolating two guesses which both fairly close to the solution the part of the graph we are “looking at” is very small and hence “very straight”.

5.10 The hunt-and-peck method

Understanding what “approximately” or “very” straight means and how to “zoom in” with a calculator instead of a picture is the subject of most of the first semester of calculus and I won’t try to say any more here. I just wanted to sneak another snippet of calculus by you so that you could see that it’s really not so hard and you can do some surprising things with it.

Finding yields using the hunt-and-peck method

So far we have focussed on finding unknown interest rates on amortized loans. The value of being able to do so might seem questionable given that when we take out a loan we almost always know the rate—the lender is obliged to spell it out in the contract. However, in the next subsection [POINTS](#), we’ll see how being able to solve for unknown interest rates can provide helpful information even when we “know” the nominal rate of a loan.

This said, the hunt-and-peck method shows its value more clearly when when we apply it to amortized savings accounts, things like college funds, retirement accounts and so on. Two factors contribute to this. First, most real-life savings accounts do not have a single interest rate. Even a bank savings account or money market account has a rate which changes periodically. Vehicles like mutual funds have values whose variation cannot be predicted even from moment to moment.

Yet we’d often like to make a retrospective assessment of the growth of such an account producing a single *yield* or nominal interest rate which summarizes how well it has performed and allowing us to compare it to other investments. Second, we often want to plan savings well into the future. Here we do not even have a history of values to work with. However, we’ll see how we can use the hunt-and-peck method to correlate goals and yields. This can help us decide

5.10 The hunt-and-peck method

whether the goals are realistic and what kind of investments may be needed to realize them.

The method is basically the same. Try to make as good a first guess as you can for the rate. Once again, the problem will often suggest a good first guess. Then try to make a second guess which brackets the desired yield. This involves understand which way to jump from comparing the trial savings from your first guess with the known amount in the problem. Use the information from these guesses to produce a third guess by one of the same three methods—**bisection**, **eyeball**, or **interpolation** and repeat this step if necessary.

The only real difference is that relation of rates and balances for savings is reversed from that for loans. In a savings account, the higher the rate of interest the higher the final balance will be. (It's as if interest were the current in a stream. In paying a loan, we are swimming upstream a faster current slows us down. In saving, we are swimming downstream and a faster current speeds us up.) So, if the trial balance is low, we need to guess a higher rate; if the balance is high, a lower rate. To convince you that there's nothing really new, let's jump right in and solve an example.

EXAMPLE 5.10.25: Here's a typical application. Suppose you are currently 35 and have been contributing \$150.00 a month in to a retirement account which is invested in a mixture of stocks and bonds for the past 8 years and the amount in the account is now \$20,523.00. You'd like to estimate how much you are likely to have in the account when you retire in 30 years. If you knew what yield you had been getting on the account for the past 8 years, you could make such an estimate by assuming that you'd continue to get this yield for the next 30 years and applying the **FUTURE AMORTIZATION FORMULA 5.6.8**—we'll do this in a moment. We can find the yield to date by applying the **HUNT AND PECK METHOD 5.10.18** to the **FUTURE AMORTIZATION FORMULA 5.6.8**.

Your account to date has a sum of $S = \$20,523.00$, a term of 8 years



5.10 The hunt-and-peck method

or $T = 8 \cdot 12 = 96$ months and a monthly payment of $D = \$150.00$. Since we have no guidance, we'll use our general knowledge about yields—bonds historically yield around 6% and stocks 8 – 10%—and make a conservative first guess of $r = 7.2\%$ or $p = 0.006$. (I added the 0.2% to the r so I'd have a simple guess for p : I might as well keep things as simple as I can to start). Using this p , the [FUTURE AMORTIZATION FORMULA 5.6.8](#) gives a trial sum S of

$$S = \$150.00 \left(\frac{(1 + 0.006)^{96} - 1}{0.006} \right) = \$19,396.23665.$$

This is low—but not wildly so—so our first guess p is also a bit low. As a second guess, let's try $p = 0.00800$ (or $r = 9.6\%$). This guess is at the high end of the historical range so I expect it will bracket the true value. Since this second guess need not be too careful, I again felt free to choose a “round” p . Now we find a trial sum of

$$S = \$150.00 \left(\frac{(1 + 0.008)^{96} - 1}{0.008} \right) = \$21,541.39970.$$

which is high as expected.

We're ready to **buff** our guess. I'll use this example to illustrate all three methods. To avoid a long string of calculations, I'll mix all three up in this problem. (You probably *won't* want to imitate this in working the problems which follow. Just pick your favorite method and stick with it.) Also, I won't show the [FUTURE AMORTIZATION FORMULA 5.6.8](#) anymore, just the trial savings S which come out.

First, I'll **bisect**. The average of my first two guesses—0.006 and 0.008—for p is 0.007-or $r = 8.4\%$ —which gives a trial balance $S = \$20,433.66174$ which is only off about \$100. In real life, I could probably stop here. After all, I'm only going to use this estimate to try to predict the likely *future* performance of my retirement account and no accurately I calculate the historical yield in the account, I can only expect to get a crude measure of the likely future yield.

But, just to see that we can get as much accuracy as we need let's continue. Next, I'll **eyeball**. When $p = 0.007$ my S is off by \$100 or so

5.10 The hunt-and-peck method

while when $p = 0.008$ it's off by about \$1,000. That suggest that the yield I'm after is about 10 times as close to 0.007 as to 0.008 so let's try $p = 0.0071$ next. This gives a trial savings of \$20,541.18837 off by about \$18.

Now to completely nail things down let's use the [LINEAR INTERPOLATION FORMULA 5.10.9](#). How can we use a formula which involves ps and Bs when our guesses involve ps and Ss ? Simple, wherever the formula calls for a B just use the corresponding S .

The reason this will work is that the idea behind the formula—plot two existing guesses as points, find the line joining them and use the formula to guess the right point on this line—is really very general. It works just as well whether the coordinates of the points plotted represent rates and balances, or rate and savings—for that matter populations and times, or soybean harvest and fertilizer, or any other two quantities related by a formula. If you're skeptical, just watch.

My two best guesses so far are $(p', S') = (0.007, \$20,433.66174)$ and $(p'', S'') = (0.0071, \$20,541.18837)$ and the savings amount I am trying to match is $S = \$20,523.00$ so the formula $p = \frac{p'(S-S'')-p''(S-S')}{S''-S'}$ gives

$$p = \frac{0.007(20523-20,541)-0.0071(20523-20,433.66174)}{20541-20,433.66174} = 0.007083084776$$

and this in turn produces a trial savings of $S = \$20,522.95130$. We're off by a nickel! Converting p we get a nominal rate of $r = 8.5\%$ for the yield on our account.

Now we can answer the problem we started with of estimate the likely amount in the account when you retire. We'll just assume that your deposits of \$150.00 a month continue for another 30 years and that the account continues to yield 8.5%—or rather a periodic rate of 0.007083. In other words, your account looks like a savings account with a 38 year term (the 8 years you already chipped in and the 30 to come) and we can estimate the final sum S in it using the [FUTURE AMORTIZATION FORMULA 5.6.8](#) as

5.10 The hunt-and-peck method

$$S = \$150.00 \left(\frac{(1 + 0.007083)^{38 \cdot 12} - 1}{0.007083} \right) = \$508,055.10 \approx \$5 \times 10^5.$$

I used that scientific notation to emphasize how rough this estimate probably is: you are likely to have about half a million when you retire but a few hundred thousand more or less is not out of the question.

PROBLEM 5.10.26: Show that a swing up or down of 1% in the final average yield (over the entire 38 years) on your account will cause the final sum in it to swing by *more* than \$100,000.00

That's it: there really was nothing new. The [HUNT AND PECK METHOD 5.10.18](#) worked just as well for amortizations involving future value as for present value. We'll see other applications in the next subsection. Here are some typical questions for you to answer. These contain a few wrinkles not found in the example so I have provided sample solutions to guide you.

PROBLEM 5.10.27: My daughter's college fund is a big priority. I am going to aim for a final sum of \$150,000.00 and what I'd like to estimate is how likely various monthly deposits (starting from her birth and continuing until age 18) are to get me to this goal. What will the fund have to yield—say to the nearest percent—for me to reach my goal if I put in,

- i) \$200 a month?

Solution

In [EXAMPLE 5.6.9](#), we saw that to build up \$120,000.00 over 18 years (so $T = 216$) at a yield of 3.9% I had to put in \$384.05 a month. Now, I want to halve the deposit and increase the goal by 25% so I will need a much higher yield. Let's try $r = 10\%$ or $p = 0.0083$. This gives a sum of

$$S = \$200.00 \left(\frac{(1 + 0.0083)^{216} - 1}{0.0083} \right) = \$119,565.52.$$

That's way too low so let's try $r = 12\%$ or $p = 0.0100$. Now if get a sum of $S = \$151,572.00$. We just found that raising r by 2%



5.10 The hunt-and-peck method

raised S by about \$20,000. Since I am less than \$2,000 off the desired S there is no point in adjusting r any further: I'll need a yield of about 12%. I'd have to be very optimistic to think I can achieve that over eighteen years.

- ii) \$300.00 a month?
- iii) \$400.00 a month?
- iv) What do *you* think my monthly deposit should be?

PROBLEM 5.10.28: I have been looking for an apartment in Manhattan in the \$200,000.00 range but none of the buildings I want to buy in will sell to me unless I have a 20% down payment. So I have decided to start putting \$300.00 a month into a mutual fund account until the sum in it reaches \$40,000.00. I'd like to understand how the yield on the account will affect the length of time it is likely to take me to reach this goal.

i) One way to get some insight into the answer is to turn the problem around, set various time limits, and ask what yield would let me reach my goal in that time. To the nearest percent, what yield will allow me to reach my goal in

- a. 5 years?
- b. 7 years?
- c. 9 years?

About how long, do you think I'll have to wait for that apartment?

ii) There is a second way to attack this problem. I can pick a yield and then ask, "How long will an account with that yield take to build up to \$40,000.00?" In other words, instead of solving for p and converting to get r , I try to solve for T and convert to get y . How many years will it take to reach my goal if I expect the account to yield

- a. 2%?

Solution

Here we know r and hence $p = \frac{0.01 \cdot r}{100} = 0.00166666666 \simeq 0.0017$ and need to guess T . The idea is always the same. Make a first guess for T , plug it in, stare at the sum and

5.10 The hunt-and-peck method

make a better guess. We don't have a formula for making first guesses here as we did above but this is of minor importance. Any reasonable guess will do.

For example, I can note that if \$300.00 a month amounts to \$3,600.00 a year which would give \$36,000.00 in 10 years with no interest. A small amount of interest might just get me up to \$40,000.00 So let's try $y = 10$ and $T = 12 \cdot y = 120$. We get a final sum of $D \left(\frac{(1+p)^T - 1}{p} \right) = \$300.00 \left(\frac{(1+0.0017)^{120} - 1}{0.0017} \right) = \$39,897$. That's so close it's got to be the answer. After all in another year another \$3,600.00 will come in even ignoring interest.

For practice, suppose my goal were \$50,000.00. Then, I need to accrue another \$10,000.00: in another three years (for a total of 13) I am sure to do this as the new deposits will total more than \$10,000.00 so let's see what happens in 12 years or $T = 144$ months. Using $B = D \left(\frac{(1+p)^T - 1}{p} \right)$ we find

$$B = \$300.00 \left(\frac{(1 + 0.0017)^{144} - 1}{0.0017} \right) = \$48,900.06.$$

So to reach \$50,000.00 will take just over 12 years. The moral is that any common sense guess is enough in a rough estimation problem like this.

b. 8%?

c. 14%?

Points

This subsection discusses a standard feature of mortgage loans where we can apply the [HUNT AND PECK METHOD 5.10.18](#). To begin with, let's ask, "Why do banks like to write mortgage loans?" I expect that most of you will answer, "Because their business is lending money and mortgages are a good way to lend money". In most cases, this is wrong. Most banks sell your mortgage before the ink is



5.10 The hunt-and-peck method

dry to a financial institution which specializes in mortgage finance. These institutions package large groups of mortgages and resell the packages (called *collateralized mortgage obligations* or CMOs) as investments to yet other financial institutions like pension plans, insurance companies and so on.

This wholesaling reduces various administrative costs like those involved in the collection of payments. Another advantage of CMO's are that they reduce risk through diversification: they are rather like mortgage mutual funds. If a homeowner defaults—can't pay his or her mortgage—this has only a tiny effect on the income and yield in a CMO. There is an opposite risk that a homeowner may decide to pay off his mortgage (by selling his house, refinancing etc.) Why, if the loan is repaid, is this a risk for the lender? The answer is that repayment happens much more often when interest rates are low. This leaves the lender holding cash at exactly the moment when it is impossible to relend it at the original rate: in effect, the lender's yield is suddenly reduced and that's a risk. Finally, CMOs can often be converted into so-called *derivative securities* each of which may be more attractive to certain buyers than a slice of the whole CMO. CMOs are a huge—trillion dollar—industry.

Back to the bank which sold you the mortgage. What's in it for them if they are just going to resell it immediately? Fees and points. The bank charges you a substantial fee to prepare and verify your mortgage application and makes a small profit on this. More importantly, when you close the mortgage you generally give a small fraction of the amount—usually 1% to 2%—back to the bank. These are called *points*. For example, if you take out a mortgage with a face value of \$100,000.00 and pay the bank 2 points, you receive only \$98,000.00 at closing to give to the seller. If you really need \$100,000.00 to buy the house, you'll need a mortgage with a face value of about \$102,040.00

POINT 5.10.29: *Each percent of the face value of a mortgage which*

5.10 The hunt-and-peck method

is returned to the lender at closing as a fee for originating the mortgage is called a point.

Basically, the bank wants to sell you the mortgage to get the points. Then it can happily resell the mortgage: it has no risk of default or repayment, its capital is not tied up in your house for 15 to 30 years. It's just ahead a few thousand dollars on the deal and ready to write the next mortgage. Actually, things are not quite this simple since the bank must discount the mortgage slightly when it resells it—in effect, the bank pays points to the CMO firm—but the bank still comes out substantially ahead on the deal. You might be asking why the CMO firm needs the banks at all. Basically, banks provide neighborhood points of sale and the ability to evaluate buyer finances much more cheaply than a specialist mortgage firm can (although mortgage web sites have been changing this quickly in recent years).

Points are a major factor in discouraging people from buying homes which they might have to resell in a few years. If you keep a home for 20 years, that couple of percent of the sale price becomes a very small fraction of your total mortgage costs, but if you keep the house for only a year or two, it can be a substantial component. The effect is as if you had paid a very high interest rate on your mortgage for this period.

For this reason, many banks offer what are called no-points mortgages where you do not pay the bank *any* points at closing. “Sounds good”, you say, “Why would anybody want to pay points if they do not have to?”. There are two reasons. First, no-point mortgages are only offered to the lowest risk buyers (e.g., those with a large down-payment who mortgage payment will be well below the 28% of income threshold and so on). Second, the interest rate on a no-points mortgage is substantially higher than on a standard two points mortgage, generally about a half a percent higher.

In other words, the banks also realize that points amount to a higher interest rate. Their rule of thumb is that you pay a quarter percent

5.10 The hunt-and-peck method

higher interest for each point you do not pay. In the rest of this subsection, we'll try to decide whether this is a good deal or not. The answer is sometimes yes and sometimes no. So a better question is, "When should I prefer a no-points mortgage and when should I prefer standard one?" I hope that, when you come to buy a home and have to make this choice, you'll at least remember that you *used* to be able to answer it.

Let's try to get a feel for this. Suppose that you take out a 30 year mortgage at a face value of \$100,000.00 and a nominal interest rate of 7.5% so your periodic rate is $p = 0.00625$ and your monthly payment is $D = \$699.21$: see [PROBLEM 5.8.4](#). You pay the bank two points and then use the remaining \$98,000.00 mortgage to purchase a house (also valued at \$98,000.00). One way to look at your situation is to say that you are paying $D = \$699.21$ a month for 30 years to have the use of \$98,000.00 now.

This is another example of an amortized loan in which we know everything except the interest rate and we can find this by hunting and pecking to solve

$$98,000.00 = \$699.21 \frac{(1 - (1 + p)^{-360})}{p}.$$

PROBLEM 5.10.30: Use the [HUNT AND PECK METHOD 5.10.18](#) to show that the nominal rate at which a 30 year mortgage with an initial balance of \$98,000.00 would have a monthly payment of \$699.21 is 7.7% to the nearest tenth of a percent. Make one more guess without rounding and show that the rate is 7.71% to the nearest hundredth of a percent.

Hint: This is one problem where we get a good first guess for free—just use the original nominal rate $r = 7.5\%$ which means $p = 0.00625$.

In other words, if you hold the mortgage for the entire 30 year term paying 2 points has the same effect as paying an extra 0.2% interest.



5.10 The hunt-and-peck method

This suggests a rule-of-thumb that each point you pay is like paying an extra tenth of a percent in interest. Let's check this.

PROBLEM 5.10.31:

i) Use the [HUNT AND PECK METHOD 5.10.18](#) to show that the nominal rate at which a 30 year mortgage with an initial balance of \$99,000.00 would have a monthly payment of \$699.21 is 7.6% to the nearest tenth of a percent.

Hint: One good first guess here is $r = 7.5\%$ or $p = 0.00625$ as in the preceding problem. But since we think we know what the answer is going to be we can make an even better guess: $r = 7.6\%$ or $p = 0.00633$.

ii) Use the [HUNT AND PECK METHOD 5.10.18](#) to show that the nominal rate at which a 30 year mortgage with an initial balance of \$97,000.00 would have a monthly payment of \$699.21 is 7.8% to the nearest tenth of a percent. You can use the hint above to make a very good first guess.

Unfortunately, this rule of thumb works only for 30 year mortgages. For shorter terms, each point paid corresponds to a larger bump in the interest rate. The next problem shows this: you'll need to find effective rates to within a hundredth of a percent so you should carry out all the calculations in these problems *without rounding*.

PROBLEM 5.10.32: In [PROBLEM 5.8.4](#), you should also have found that the monthly payment on a 30 year mortgage with a face value of \$100,000.00 and a nominal interest rate of 7.5% is $D = \$927.01$.

i) Use the [HUNT AND PECK METHOD 5.10.18](#) to show that the nominal rate at which a 15 year mortgage with an initial balance of \$98,000.00 would have a monthly payment of \$927.01 is 7.83% to the nearest hundredth of a percent.

ii) Use the [HUNT AND PECK METHOD 5.10.18](#) to show that the nominal rate at which a 15 year mortgage with an initial balance of \$96,000.00 would have a monthly payment of \$927.01 is 8.16% to the nearest hundredth of a percent.



5.10 The hunt-and-peck method

In other words, for a 15 year term each two points raises the rate by close to a third of a percent so each point raises the rate by a sixth of a percent, rather more than the tenth of a percent for a 30 year term. However, these rules-of-thumb apply only if you hold the mortgage through the full term. Since the rises for both terms are well below 0.25%, we can however draw one conclusion. If you expect to live in that house “forever”, you are better off to pay the points. You pay more with a no-points mortgage in the long run.

What if you sell your home before the end of the mortgage? Things are worse—much worse if you decide to sell after a short period. Suppose that you take out a 30 year mortgage at a face value of \$100,000.00 and a nominal interest rate of 7.5% so your periodic rate is $p = 0.00625$ and your monthly payment is $D = \$699.21$, as we saw in, [PROBLEM 5.8.4](#). You pay the bank two points and then use the remaining \$98,000.00 mortgage to purchase a house (also valued at \$98,000.00).

After two years, you decide to sell the house. The good news is that you'll get \$98,000.00 from the new buyer (as usual I've ignore changes in value, down payments etc.). The bad news is that you will have made $i = 24$ of the $T = 360$ payments and have $T - i = 336$ left to make. The [BALANCE AND EQUITY PRINCIPLE 5.9.2](#) says that your outstanding balance—what you owe the bank—is $B = D \left(\frac{1 - (1+p)^{-(T-i)}}{p} \right)$ or

$$B = \$699.21 \frac{(1 - (1 + 0.00625)^{-(336)})}{0.00625} = \$98,084.13.$$

That's \$84.13 more than the buyer gave you. You just acquired *negative* equity: you have to pay money to sell your home

I claim you paid what amounted to 8.6% interest on your mortgage. We can see this roughly by ignoring the small balance of \$84.13 and saying that basically all of the \$699.21 a month was interest. In other words, if p is the periodic (i.e., monthly) interest rate which you really

5.10 The hunt-and-peck method

paid, $p \cdot 98,000.00 = 699.21$ which gives $p = 0.007134795917$ and $r = 8.6\%$ to the nearest tenth.

Again this calculation applies only if you decide to sell after two years. It's possible to hunt-and-peck to work out what no-points rate corresponds to a mortgage with points which you close out after any number of payments but things get a bit complicated so I have made this [CHALLENGE 5.10.34](#). But it's clear that if your plan is to move soon after buying, then you should pick the no points mortgage. The extra half a percent of interest you'll pay is much less than the effective bump in interest which the points amount to over such a short span.

Let's summarize: we can think of paying points on a mortgage as equivalent to paying a higher interest rate. There is no simple formula for making the conversion but we can say that the shorter the term of the mortgage the more each point lifts the rate and the shorter the time you hold the property before selling, the more each point lifts the rate. If you plan to stay in the house until the mortgage is paid off, a standard mortgage with points is probably better. If you are likely to sell in a few years, a no-points mortgage is probably a better deal. The break even point is about 5 years for a 30 year term and $5\frac{1}{2}$ for a 15 year term.

RULE-OF-THUMB FOR POINTS 5.10.33: If you plan to keep paying the mortgage longer than five years, pay the points; if less, try to get a no-points mortgage.

CHALLENGE 5.10.34: In this project, we'll figure out how to compare standard and no-points mortgages under scenarios in which you sell the property after any number of payments. This will involve hunting and pecking in an equation more complicated than the [PRESENT AMORTIZATION FORMULA 5.8.1](#). I'll set the equation up—as I said above, it's a bit complicated—and let you deduce the consequences.

5.10 The hunt-and-peck method

Let's suppose that your mortgage has a face value B , payment D , periodic rate p and term T and that you decide to sell out after i months. We also introduce a second balance C , the amount you actually get from the bank after the points are deducted and a second periodic rate q which represents the interest which would leave you in the same financial position after you settle if we ignored the points. (I used C and q for the *second* balance and rate because, in the alphabet, these letters follow the letters B and p used to describe the *first* balance and rate.)

The first goal is to find an equation involving q . To do this we first ask two questions: What sums will you have received from the bank, and what sums will you have paid out to the bank? The first is easy: you got C at the start of the mortgage. The second is not much harder: there will have been i monthly payments of D each, and at the end of the i months, the [BALANCE AND EQUITY PRINCIPLE 5.9.2](#) says that your outstanding balance—what you owe the bank to close out the mortgage—is $B_i = D \frac{(1-(1+p)^{-(T-i)})}{p}$.

We'd now like to say that we get an equation by equating the money you got from the bank to the money you gave. However, we can't just crudely equate since these sums live at different points in time. The whole theme of this chapter is that money changes value when it travels in time. To equate sums, we have to first move them all to the same point in time. This poses two more questions.

First, what rate do we use to compound these sums when they travel in time? That's easy, q ! In fact, saying that q is the real rate you were paying exactly means that when you use it to compound the various amounts you exchanged with the bank they balance.

Second, which common point in time shall we move them to? The obvious choice is the start of the mortgage. Now let's tackle the sums one-by-one. Since your income C exists at this moment there is no need to adjust it. What about the outstanding balance B_i ? It lives i periods in the future from the start of the mortgage so we have to



5.10 The hunt-and-peck method

move it back i periods in time at a periodic rate q . By the **COMPOUND INTEREST FORMULA 5.2.4**, this means we multiply it by $(1 + q)^{-i}$ to get $(1 + q)^{-i} \cdot D \frac{(1 - (1 + p)^{-(T-i)})}{p}$. Finally, there are the i payments of D each. We can think of these as amortizing a loan with a payment D , a term of i periods and periodic rate q . (Since we want to move these payments to a common point in time, we again use the fair rate q to do it).

The **PRESENT AMORTIZATION FORMULA 5.8.1** tells us that at the start of the mortgage this loan will have a balance of $D \frac{(1 - (1 + q)^{-i})}{q}$. Now we can equate the total value at the start of the mortgage of the amounts you got from the bank with the total value at the start of the mortgage of the amounts you paid the bank to get

$$C = (1 + q)^{-i} \cdot D \frac{(1 - (1 + p)^{-(T-i)})}{p} + D \frac{(1 - (1 + q)^{-i})}{q}$$

That's one big ugly formula and the quantity q we're looking for now occurs three times. But don't be scared; you have nothing to fear but fear itself. You'll see that the hunt and peck method works like a charm. When we plug values into the right side of the formula to compute trial values of C , we have to hit a few extra keys on the calculator but that's all that changes. You can even adapt the **LINEAR INTERPOLATION FORMULA 5.10.9** if you wish: here we plug in qs and compute trial balances C so we just replace the ps and Bs in the formula by these qs and Cs .

i) Let's start by checking this formula against some values we computed above. The rates we worked out above should all be special cases and so we should get equalities if we plug in our earlier results on both sides. We will suppose that $B = \$100,000.00$, $r = 7.5\%$ and $p = 0.00625$ as usual.

- a. First, let's suppose that we have a 30 year term so $T = 360$ and that $i = 360$ too. In other words, we pay off the mortgage. Show that we get near equality when $C = \$99,000.00$ and $q = 0.00633$, when $C = \$98,000.00$ and $q = 0.00642$, and when

5.10 The hunt-and-peck method

$C = \$97,000.00$ and $q = 0.00651$. Compare this with the result of [PROBLEM 5.10.30](#).

b. Next, let's suppose that we have a 15 year term so $T = 180$ and that $i = 180$ too. In other words, we again pay off the mortgage. Show that we get near equality when $C = \$98,000.00$ and $q = 0.00653$ and when $C = \$96,000.00$ and $q = 0.00680$. Compare this with the result of [PROBLEM 5.10.31](#).

c. Finally, let's suppose that we have a 30 year term so $T = 360$ and that $i = 24$. In other words, we sell after two years. Show that we get near equality when $C = \$98,000.00$ and $q = 0.00715$. Which calculation above does this confirm?

ii) Now find a few solutions of your own using the hunt-and-peck method. Let's suppose that we have a 15 year term so $T = 180$ and that $C = \$98,000.00$ so we are paying two points. Find the interest rate q and the corresponding nominal rate s which we'd be paying if we sold out after $i = 24$ months and after $i = 120$ months.

iii) OK, let's go for the gusto! If we keep the 15 year term, we either pay two points ($C = \$98,000.00$) or the bank charges us 8% interest instead of 7.5%. How many months i do we have to hold on to the property to make these equivalent? That is, for if we plug in the values above for C , D , T and p and also set $q = \frac{0.01 \cdot 8.0}{12} = 0.00667$ *what value of i brings us nearest to equality.*

Hint: You can still hunt-and-peck: in some ways it's easier since i takes on only whole number values. The statement above that the break-even point between standard and no-points mortgages comes at about 5 years should give you a good first guess for i .

iv) Repeat part [iii](#)) for a 30 year term, again assuming you pay 2 points. What happens if you only pay 1 point?

Final comments

This section has a point somewhat broader than the problems we have attacked which all involved financial formulae. The [HUNT AND](#)



5.11 Saving to spend

PECK METHOD 5.10.18 can be used to find accurate approximate solutions to many kinds of equations. The next time you'd like to solve an equation for one of the variables and it's impossible or too complicated to isolate the variable and just plug in, *don't give up!*

Try guessing, plugging in and reguessing. With enough elbow grease this will do the job surprisingly often. If you need a lot of decimals, you'd probably be smart to interpolate between your guesses (one of the goals of this section was to give you a feel for doing this) but the slower methods can give whatever accuracy you need if you keep pecking a little longer. Likewise, if you have to solve lots of equations, or very complex ones, you'll probably find it worth your while to seek the advice of someone who knows some **calculus** or numerical analysis or even find it worthwhile to learn a bit about these subjects yourself.

5.11 Saving to spend

Now we are ready to pull out all the stops and combine savings and loans amortizations in a single problem. You might think that I'd have to make us some artificial situation to come up with such a problem but you'd be wrong. Your retirement planning is a perfect example. In considering retirement planning so far, we have always had as our goal *saving* a certain lump sum of money in the at retirement: this was always a future value amortization. That's OK if you are a Midas who just wants to count his gold but what most people want when they retire is an *income*. The real point of accumulating that lump sum in your retirement account is to be able to finance such an income.

Fixed term annuities

The mechanics of doing so are just like those of a mortgage or present value amortization. The only difference is that *you're the bank*: you hand over the lump sum (like the bank handing the homeowner that initial mortgage balance) when you retire and then sit back and receive a series of monthly payments (analogous to those the bank gets from the homeowner) which represent your retirement income. It's only rarely that simple. It is possible to use your retirement fund to buy what's called a **fixed term annuity** which promises you a fixed monthly payment for a fixed term. But most people are afraid—or better, hopeful—that they'll outlive that term and be left without any income. Such people usually buy a **life annuity** which promises a fixed monthly payment as long as you live.

This involves merging a whole series of amortizations with different terms—after all, the term of a life annuity is not fixed but depends on how long you live. Thus, the price of such an annuity has to take into consideration how likely you are to die at various ages—the fancy term for this is mortality—and somehow merge the various present values into a single price. Moreover, these are just the simplest wrinkles. You might want an income until both you and your spouse die (a **joint life annuity**) and you might want the annuity to take into account inflation or rises in the cost of living. The complications which arise have spawned an entire profession: actuaries specialize in the mathematics and statistics of this type of calculation.

Moreover, the yield on such annuities is generally fairly low. If you can tolerate a bit more risk, you can do better by continuing to invest your retirement fund yourself and paying yourself an income out of the fund. If you do this, you'd like to be sure that the checks are not going to suddenly start bouncing when you are 73. When the time comes, I'd suggest getting professional advice. But we already know enough to make calculations which can help you in planning when

5.11 Saving to spend

you are younger. I'd like to look at few examples of these to close the chapter.

EXAMPLE 5.11.1: Here's an example which will illustrate the basic idea. Right now, I am 35 and just starting my retirement planning. I have found an insurance company which offers fixed term annuity accounts which earn 5.1% a year (compounded monthly as usual) and I have decided that when I am 65 I'd like to be able to buy such an annuity with a term of 25 years and monthly payment of \$2,500.00 I plan to finance this purchase with the proceeds of a retirement account into which I will make monthly payments over the intervening 30 years. With this long term time horizon, I am prepared to invest this account fairly aggressively—that is, in higher yielding but riskier securities—and I think it is prudent to plan on an average return of 9% over this period. The basic question I need to answer now is, “How much do I need to deposit every month?”

What are we dealing with here? Two separate amortizations. My deposits into the retirement account are a savings or future amortization: I know the term is 30 years (so $T = 360$) and I am assuming the rate will be 9% (so $p = \frac{0.01 \cdot 9}{12}$) but I know neither the sum which is my goal nor the monthly deposit I will make. This means we are missing two ingredients in the **FUTURE AMORTIZATION FORMULA 5.6.8**: S and D and so I can't use it to find out either. So we'll put this aside for now.

The annuity I plan to buy represents the other amortization. It's a loan or present value amortization in which, as remarked above, I'm the bank and the insurance company plays the role of the mortgage holder who makes regular monthly payments. I know the term of this annuity is 25 years (so $T = 300$), the rate r is 5.1% (so $p = \frac{0.01 \cdot 5.1}{12}$) and the monthly payment D is \$2,500.00 Thus I can use the **PRESENT AMORTIZATION FORMULA 5.8.1** to find the balance B I'll need to pay the insurance company. Using $B = D \left(\frac{1 - (1 + p)^{-T}}{p} \right)$, we find,



$$B = \$2,500.00 \left(\frac{1 - \left(1 + \left(\frac{0.01 \cdot 5.1}{12}\right)\right)^{-300}}{\left(\frac{0.01 \cdot 5.1}{12}\right)} \right) = \$423,419.45.$$

Now comes the only new point. The balance B that I'll need to *purchase* the annuity is the same thing as the sum S that I'll want to have *saved* in my retirement account when I'm 65: so $B = S = \$423,419.45$. Now we *do* know enough to use the [FUTURE AMORTIZATION FORMULA 5.6.8](#) to find the balance to find the deposit D I'll need to make into the retirement account. Here, using , $D = S \left(\frac{p}{(1+p)^T - 1} \right)$, we find that

$$D = \$423,419.45 \left(\frac{\left(\frac{0.01 \cdot 9}{12}\right)}{\left(1 + \left(\frac{0.01 \cdot 9}{12}\right)\right)^{360} - 1} \right) = \$231.28.$$

I need to deposit \$231.28 a month.

To emphasize, let's restate the new idea here. We have a pair of related amortizations: a future or savings amortization in which we assemble a sum of money which we then use to finance a present or loan or annuity amortization. (Of course, both amortizations live in the future, the "present" one lives further in the future than the "future" one and the "loan" is really a purchase: we are using all these terms in the conventional sense established earlier in the chapter).

If we know everything but the final sum S of the savings amortization, the initial balance B of the loan amortization and one other quantity, then we can use a two step process like that above to completely describe both amortizations by identifying B and S . Either, as above, we know everything about the loan or annuity except the balance B and can solve for this using the [PRESENT AMORTIZATION FORMULA 5.8.1](#), or, as above we know everything about the savings amortization and can solve for S using the [FUTURE AMORTIZATION FORMULA 5.6.8](#). Then, we use the fact that $B = S$ to use the other formula to determine the remaining missing piece of information about

5.11 Saving to spend

the other amortization (this was the deposit of \$231.28 above). Here are some problems for you to try.

PROBLEM 5.11.2: Recalculate the payment I'll need to make in my retirement account keep all the values of [EXAMPLE 5.11.1](#) except that we assume that,

- i) my annuity has a yield of 3% and my retirement account has a yield of 5%. (Here I am asking, "What's the worst that can happen?".)
- ii) I am 45 years old and want to retire and buy a 30 year annuity when I am 60.
- iii) I want the annuity to pay \$3,000.00 a month.

PROBLEM 5.11.3:

- i) Suppose that I am 45 years old and starting a retirement account. Based on my current income, IRS will only let me put \$275.00 a month into this account tax free. If I think the account will have a yield of 10%, what monthly payment can I expect to get from a 25 year annuity yielding 4.5% which is purchased with the sum in the account when I am 65?
- ii) How does the answer to i) change if I am 25 today?
- iii) How do the answers to i) and ii) change if I expect a yield of 7.5% on both the retirement account and the annuity?

Life annuities

Let's conclude with a capstone problem in which we'll model a more realistic annuity that is to be paid until the holder's death. Such an annuity is called a **life annuity**. For any single individual, this annuity—like those above—will amount to a loan amortization. The big difference is that we will only discover what the term of this annuity is in hindsight, after the holder dies. The question we need to answer is how to price such an annuity *before* we start to pay it out. Our plan has four stages. First, we use a probability distribution to encode the chances that the holder will live to various ages. In other

5.11 Saving to spend

words, an outcome in our sample space is just an age from 65 on at which the holder might die. The probability that the holder will die at each such age is an empirical one. It depends on many things of which the most important is the holders current age: for example, a person who is 65 today has a very small chance of dying at age 100 (because he or she will almost certainly die before that age), but a person who is 99 today has quite a good chance of dying at age 100 (having avoided dying at any younger age).

Tables that give such probabilities are known as mortality tables. They are usually given by specifying how many of 1000 people who reach each age (e.g. turn 73) will die before reaching the next age (i.e. die before they turn 74) or how many of each 1000 persons born in the same year survive to reach each age. From such figures, some easy arithmetic lets us determine the values we want here, the percentage of people who reach age 65 who die at each subsequent age (e.g. 73 again). This is the easy second stage in our plan.

Let's start to look at [TABLE 5.11.5](#). I have just assumed our holder is a 65 year old male when the annuity begins. I fix the sex as well as the age because the mortalities for men and women are very different. In a more realistic study, we'd want to consider other factors (e.g. if you're a smoker, your annuity will be cheaper because you figure to die much younger). The data in the first column of [TABLE 5.11.5](#) were extracted from a much more detailed set of tables on the [website of the American Association of Actuaries](#) and I worked out the second as in shown in [PROBLEM 5.11.4](#).

PROBLEM 5.11.4: Use the values in the first column of [TABLE 5.11.5](#) to verify the chances that a person who reaches 65 will die at age 65 or at age 66. Hint: We want the number of people who died aged 66 or aged 67 as a percentage of the cohort that made it to 65. For example, the number of people who die aged 66 is the difference of the numbers of survivors in the 66 and 67 rows and the number of people who reach 65 is in the 64 row.



5.11 Saving to spend

Age	Survivors to end of year per 1000	Chance of dying at this age	Cost of annuity to this age	Products of Chance × Cost
64	834.0363	—	—	—
65	819.9828	1.69%	48076.92	810.10
66	804.8377	1.82%	94304.73	1712.46
67	788.6685	1.94%	138754.55	2689.99
68	771.4361	2.07%	181494.76	3749.95
69	753.1993	2.19%	222591.12	4867.10
70	733.7894	2.33%	262106.84	6099.83
71	713.1332	2.48%	300102.73	7432.50
72	690.7979	2.68%	336637.24	9015.08
73	666.8825	2.87%	371766.58	10660.15
74	641.4876	3.04%	405544.79	12348.10
75	614.6028	3.22%	438023.84	14119.48
76	586.2819	3.40%	469253.69	15934.18
77	556.4285	3.58%	499282.39	17871.30
78	524.9569	3.77%	528156.15	19929.49
79	491.8531	3.97%	555919.37	22065.03
80	457.3545	4.14%	582614.78	24098.93
81	421.5940	4.29%	608283.44	26081.06
82	385.1092	4.37%	632964.85	27688.91
83	348.3274	4.41%	656696.97	28960.95
84	311.6033	4.40%	679516.32	29920.36
85	275.2797	4.36%	701458.00	30549.60
86	239.7934	4.25%	722555.77	30743.07
87	205.6588	4.09%	742842.08	30402.28
88	173.4259	3.86%	762348.16	29462.38
89	143.6174	3.57%	781104.00	27916.64
90	116.6662	3.23%	799138.46	25823.55
91	93.0483	2.83%	816479.29	23120.73
92	72.7796	2.43%	833153.16	20247.25
93	55.7462	2.04%	849185.73	17342.73
94	41.7461	1.68%	864601.67	14513.18
95	30.5093	1.35%	879424.68	11848.32
96	21.7946	1.04%	893677.57	9337.87
97	15.1869	0.79%	907382.28	7188.79
98	10.2986	0.59%	920559.89	5395.50
99	6.7780	0.42%	933230.66	3939.28
100	4.3163	0.30%	945414.10	2790.44
101	2.6758	0.20%	957128.94	1882.66
102	1.6106	0.13%	968393.21	1236.70
103	0.9387	0.08%	979224.24	788.94
104	0.5280	0.05%	989638.69	487.36
105	0.2855	0.03%	999652.59	290.58
106	0.1478	0.02%	1009281.34	166.61
107	0.0729	0.01%	1018539.75	91.48
108	0.0341	0.00%	1027442.06	47.86
109	0.0150	0.00%	1036001.99	23.72
110	0.0061	0.00%	1044232.68	11.06
111	0.0023	0.00%	1052146.81	4.82
112	0.0008	0.00%	1059756.54	1.94
113	0.0003	0.00%	1067073.60	0.71
114	0.0001	0.00%	1074109.23	0.23
115	0.0000	0.00%	1080874.26	0.09
Value of Life Annuity			\$581,711.31	

TABLE 5.11.5: LIFE ANNUITY TABLE FOR 65 YEAR OLD MALE



5.11 Saving to spend

Next, we turn to the third stage of your plan and fill in the third column. I have assumed that the interest that will be paid on the annuity is 4%, that it will be paid annually and that the last payment is made at the end of the year in which policy holder dies. To fill in this column, we repeatedly use the [PRESENT AMORTIZATION FORMULA 5.8.1](#) with different terms to calculate what it will cost to provide the annuity to a holder who dies at each possible age.

PROBLEM 5.11.6: Since $m = 1$ (annual payments) here, the term T is just the number of years in which payments are made. This is the age at death of the policy holder minus 64 (and not 65, since someone who dies aged 65 still get a single payment). Use this to verify the cost the annuity when the holder dies age 65 and the cost when the holder dies aged 80.

Now we are almost home. Our table is perfectly set up to be used to compute an expected value. It has one row for each outcome x in our sample space—that is, each possible age at death of a man who had reached age 65—and the second entry of the x^{th} row is exactly the probability $\Pr(x)$ of that outcome. The entries in the third column are the values of the random variable Y that assigned to each outcome or age of death the cost of providing an annuity of \$50,000 a year to that age.

Conceptually, the expected value $\mathcal{E}(Y)$ is just what we're after: the expected cost of providing an annuity of \$50,000 a year to a person aged 65 until his death. In this expected value, each of the different dollar amounts (values of $Y(x)$) is weighted by the probability that a policyholder will die at the corresponding age x and collect $Y(x)$. So all we have to do, by [OUTCOMES EXPECTED VALUE FORMULA 4.8.2](#), is first to form the products $Y(x)\Pr(x)$ and then take their total. The products form the right column of [TABLE 5.11.5](#) and the total \$581,711.31 is shown at the bottom.

A few closing comments on this example. First, note that, rather paradoxically, it's *cheaper* to purchase a life annuity than one with

5.11 Saving to spend

a fixed term. For example, the \$581,711.31 cost of our life annuity is less than the cost of an annuity with a 16 year term. This latter would make annual payments to you each year through age 80, so from the table or from [PROBLEM 5.11.6](#) would cost 528,156.15.

PROBLEM 5.11.7: Show that you'd be pretty foolish to opt, at age 65, for an annuity with a 16 year term option. Hint: Show that more than half of all of men who reach age 65 will survive past age 80.

In other words, you need to “bid high” if you are trying to set the term of an fixed term annuity since you need to retain your income to any age that there's a reasonable chance you might hit. Even a 25 year term (which takes you though age 89 and cost 200,000 more than our life annuity) still leaves you with a one in seven chance of surviving uncovered beyond the end of the annuity. But much of the value of an annuity with such a long term is likely to be wasted. For example, you have a roughly 50 chance of dying by age 80. If you do, the last 10 years of an annuity with a 25 year term will benefit not you but your heirs.

The fixed term option is more expensive because much of what you are buying will not come back as retirement income. The life annuity is cheaper because it guarantees to apply the full purchase price to your retirement needs. The seller can afford to carry the fewer than 1 in 5000 policy holders who live to 106 and cost it over \$1,000,000 each because it benefits from the more than 25% who die by age 75 and cost it less than 450,000. Of course, there is a bit of double jeopardy. When you buy a life annuity you either hit the jackpot (live to a ripe age and get a “discount” on your retirement income) or get slimed (die you and “overpay”). But those air quotes are needed because in both cases the life annuity made it possible for you to contract for what you want, lifetime income, as cheaply as possible.

Finally, a project for those who want to start to understand how real-world complications affect the kind of financial lessons we have learnt in this chapter.

PROJECT 5.11.8: What fundamental reality has been totally ignored in this entire chapter? Hint: it's not death. Right, taxes! We have discussed all kinds of financial planning decisions which in real life are critically affected by tax considerations without ever mentioning this issue. In this project, I'd like you to pick a few of the examples we have looked at and to try to understand how tax issues should affect your thinking about them. Here are a three suggestions but feel free to pick others if they interest you.

- i) (Mortgage interest) The *interest* you pay on a mortgage on your primary residence may generally be deducted from your income for federal tax purposes. How does this affect the relative desirability of owning versus renting? Illustrate your general discussion with a concrete scenarios involving low, middle and upper income taxpayers.
- ii) (Retirement income) The Federal Government encourages you to save for retirement by allowing you to deduct many forms of retirement savings from your income for federal tax purposes. However, when you spend these savings at retirement, they—and any gains realized from them—are subject to taxes. How might this affect your planning for retirement? Are there situations when you might want to pay taxes on income before investing it for retirement?
- iii) (Capital gains) Your profits on many investments like stocks and bonds are subject to tax *when realized*. For example, if you bought 100 shares of a stock for \$10.00 a share and it is now selling at \$20.00 a share you have a paper profit of \$1,000.00. If you sell the stock and realize the profit, it will be subject to capital gains tax. You will not, however, be taxed if you continue to hold the stock. In effect, the government is allowing you to invest that profit without paying taxes on it but only in the stock you already own. Discuss how this affects the real yield of investment strategies. Illustrate your discussion with comparative scenarios. For example, how much higher a before-tax yield must you be able to achieve via a strategy which involves selling assets every 6 months to match the after-tax yield of a strategy in which you hold assets for an average of a decade?

Index

- A, 512
- A_0 , 531
- A_T , 531
- B, 545
- B_i , 616
- D, 615
- I, 512
- S, 545
- S_i , 615
- T, 512
- e, 68, 556
- exp, 71
- ln, 56
- m, 515
- p, 513
- (), 150
- r, 514
- r_a , 563
- r_c , 568
- { }, 141
- [], 180
- abomination, 197, 217, 284, 296
- abomination formula, 296
- allele, 372
- alphabet, 151
- American Idol, 167
- amortization, 614
- amount, 512
- Ancre, pique et soleil, 107
- and, 218
 - and also or restrictive and, 221
 - and then or extensive and, 218
 - in probability problems, 329
- and also, 221
- and then, 218
- and also, 329
 - in probability problems, 329
- and also, 221
- and then, 329
 - in probability problems, 329
- annualized yield, 563, 564, 572
 - formula, 564
- annuity
 - fixed term, 716
 - joint life, 716
 - life, 719
- anthropogenic, 4
- antonyms
 - describe complements, 238
- approximations
 - continuous, 551►, 556, ◀560
 - future amortization, 623
 - present amortization, 647
 - interest, 649
 - simple interest, 548►, 550, ◀551
 - future amortization, 625
 - present amortization, 648
- arithmetic summation, 36
- at random, 326, 326
 - making multiple choices, 326
- atom, 138
- Bürigi, Joost, 50
- Bầu cua cá cốp, 107
- balance, 616, 643
 - intermediate, xvi, 616
- ballot box, 353
- ballota, 353
- bang zoom principle, 161
- bang zoom rule, 160, 161
- baseball
 - 1981 National League playoffs, 96
- base e, 556
- Bayes' "Theorem", 393, 395

- Bayes' Theorem, 392, 393
- Bayes' theorem, 394
- Bayes, Thomas, 392
- Bayesian probability, 392, 395
 - posterior, 395
 - prior, 395
- bell curve, 490
- Bernoulli
 - trial, 420
- Bernoulli trial, 420
- Bernoulli trials
 - equivalent, 420, 423
- Bernoulli's limit
 - for exp, 74
 - for $l = \ln$, 66
- Bernoulli, Jakob, 81, 354
- better guess, 690
- binary sequence, 154
- binomial
 - distribution, 426
 - formula, 426
 - event, 421
 - experiment, 420
 - sample space, 420
- binomial coefficient, 172, 296
- birdcage, 107
- birthday problem, 94
- bisecting, 701
- bisection, 686, 692, 692, 700
- bisection method, 679, 682
- bit, 154
- border, 373, 374
- bovine spongiform encephalopathy, 3, 122
- bracket, 678, 690, 691
- Brahe, Tycho, 51
- branch, 375
- branches, 375
- Brigg, Henry, 51
- BSE, 122
- bttdltr, 19
- buffing, 678, 690, 691, 701
- calculator parentheses rule, 17
- calculators
 - Murphy's law, 554
- calculus, 557, 715
- Cantor, Georg, 167
- Cartesian product, 157
- Center for Disease Control, 369
- character
 - of a population, 365
- checks
 - compound interest, 550
 - continuous approximation, 556
 - future amortization, 623
 - simple interest approximation, 625
 - present amortization, 647, 649
 - simple interest approximation, 648
 - simple interest, 548
- chuck-a-luck, 107, 272
- CJD, 122
- closed form, 38, 619
- coin
 - fair, 305
- combination, 195, 296
- combinations, 172
- come-out roll, 384, 385
- commute, 160
- complement, 236
 - of A in S , 236
- compound experiment, 353
- compound experiments, 374
- compound interest, 528
 - formula, 530►, 531, ◀546
 - method for, 543
 - provisional method for, 534
- compounding frequency
 - effect of, 546, 550, 551
- compounding period, 528
- conditional probabilities
 - distinguishing from intersection, 349, 350
- conditional probability, 337, 340, 360
- confidence interval, 476
- conjecture, 289
 - abomination, 289
- consecutive values, 279
- constants
 - $10^{\sqrt[3]{2}}$, 44
 - e , 68
 - discovery of, 81►, ◀82
 - $\sqrt{2}$, 43
 - $10^{\sqrt{2}}$, 43
- container, 139
- continuous, 490
- continuous approximation, 551►, 556, ◀560
 - future amortization, 623
 - present amortization, 647
- continuous yield, 568, 569, 572
 - formula, 570
- convergent evolution, 107
- conviction is not knowledge, 291



- Copernicus, 52
- counterexamples, 292
- Coyote, Wile E., 42
- craps, 301, 384
 - come-out roll, 384
 - craps total, 384
 - natural, 384
 - point, 384
 - shoot, 384
- Creutzfeldt-Jakob disease, 122
- Crown and anchor, 107
- de Bruijn, Nicolas, 297
- de Finetti diagram, 409
- de Finetti, Bruno, 409
- de Saint-Vincent, Grégoire, 80
- de Sarasa, Alphonse Antonio, 81
- deductive, 292
 - proof, 292
- deflation, 581
- dependent, 412, 414
- dependent events, 412
- deposit, 615, 620
- Diaconis, Persi, 325
- direct product, 157
- disjoint, 147, 222, 310
- divide and conquer, 156, 212
- divide and conquer counting strategy, 247
- dominant, 373, 373
- downer, 132
- effect of more frequent compounding, 546, 550, 551
- effective, 489
- effective rate, 566
- effective yield, 566
 - formula, 566
- element, 140, 299
- Elkies, Noam, 292
- empirical probability, 316
- empty set, 147
- equality of dollars, 513, 525, 530
- equality of periods, 513, 525
- equally likely, 327
- equally likely outcomes, 317, 317, 317, 318, 340, 366
- equity, 652
- equivalent, 420, 423
- Erdős, Paul, 167
- Euclid, 46
- Euler, Leonhard, 82, 291
- event, 299, 300, 302
- evidence
 - by proof or deduction, 292
 - inductive, 113, 289
- expected value, 272, 440, 441, 442, 443, 449, 449, 476
 - equally likely outcomes
 - formula by outcomes, 443
 - formula by values, 449
 - formula by outcomes, 442
 - formula by values, 449
- expected values
 - method for, 449
- experiment, 113, 300
 - compound, 353
- experimental probability, 316
- experiments
 - compound, 374
- exponential property, 75
- exponentials
 - Bernoulli's limit for, 74
 - definition of exp, 71
 - definition of b^x using $\exp(x)$, 78
 - definition of e^x as $\exp(x)$, 77
 - increasing property, 73
 - key property, 75
 - other properties of exp, 75
 - tend to ∞ fast, 73
 - using calculators to find values, 79
 - values of exp as, 76
- extensive "and", 218
- eyeball, 681, 692, 692, 700
- eyeballing, 701
- face card, 158
- failure, 420
- fair, 305, 325
- fair game, 109, 453
- false negative, 367
- false negatives, 90
- false positive, 367
- false positives, 90
- Fermat's Last Theorem, 291
- Fermat, Pierre, 291
 - Last Theorem, 291
- fixed term annuity, 716
- flush, 215, 279, 283
 - straight, 279
- formulae
 - binomial distribution, 426
- formulas

- abomination, **296**
- and-or formula for orders, **230**
- and-or formula for probabilities, **309**
- annualized yield, **564**
- complement formula for orders, **237**
- complement formula for probabilities, **309**
- compound interest, **530**►, **531**, ◀**546**
- conditional probability, **340**
 - order matters, **341**
- continuous yield, **570**
- effective yield, **566**
- equally likely outcomes formulae for probabilities, **320**
- equally likely outcomes probability of an event, **318**
 - of an outcome, **317**
- future amortization, **620**, **630**, **718**
- geometric series, **41**
- geometric summation, **39**
- interest rate conversion, **515**, **534**
- intersection probability, **340**
 - for independent events, **418**
- linear interpolation, **682**, **684**, **685**, **691**, **695**, **702**
- multiplication principle, **219**
- multiset, **296**
- orelse formula for orders, **232**
- orelse formula for probabilities, **311**
- power set counting, **164**
- present amortization, **643**, **717**
- probability of an event, **308**
- product set counting, **157**
- relations of complements, **236**
- sequence counting, **155**
- simple interest, **513**
- term conversion, **518**, **534**
- four-of-a-kind, **215**, **279**, **283**
- fractions
 - converting to decimals, **29**
- Frye, Roger, **292**
- full house, **29**, **279**, **283**
- future amortization
 - formula, **620**, **630**, **718**
 - method, **621**
- future value, **536**
 - ratio to present value, **545**
- Gaussian, **489**
- general compounded quantity, **586**
 - inflation, **581**►, ◀**591**
 - Moore's law, **598**►, ◀**604**
 - population, **591**►, ◀**598**
 - radioactivity, **604**►, ◀**612**, ◀**613**
- geometric series, **41**, **389**
 - formula, **41**
- geometric summation, **37**, **619**
 - formula, **39**
 - ratio, **37**
- given, **340**, **348**
 - event, **348**
- given event, **340**
- gold standard, **370**, **370**
- guessing
 - first rule of, **193**
 - second rule of, **193**
- ham, **395**
- Hardy, G. H., **409**
- Hardy-Weinberg principle, **409**
- hazard, **384**
- high card, **279**, **283**
- Hoo hey how, **107**
- house advantage, **453**, **454**
 - roulette, **454**
- hunt-and-peck method, **690**, **713**
- Huygens, Christian, **81**
- if, **348**, **348**
- increasing, **67**, **73**
- independence, **106**, **127**, **421**
- independent, **411**, **412**, **412**, **418**, **422**, **426**, **464**, **482**, **486**
 - random variables, **464**, **466**
- independent events, **412**
- index of summation, **36**
- indicator, **462**, **468**
- inductive, **113**, **289**
 - evidence, **113**, **289**
- infinite order, **144**
- inflation, **581**►, **581**, ◀**591**
- inflation rate, **582**
- interest
 - compound, **528**
 - simple, **513**
- interest approximation, **649**
 - present amortization, **649**
- interest rate, **511**, **525**
 - conversion formula, **515**, **534**
 - effective, **566**
 - nominal, **514**
 - periodic, **513**

- interpolating, **692**
- interpolation, **684**, **692**, **700**, **702**
- intersection, **221**, **226**, **330**
 - of two sets, **221**
 - probability
 - distinguishing from conditional, **349**, **350**
- intersection probability, **360**
- intersection probability formula, **340**
 - for independent events, **418**
- interval of uncertainty, **678**
- inverses, **72**, **295**
 - exp and ln, **72**
 - of combination to abomination, **295**
 - order of, **72**
- investment, **561**
 - buying, **561**
 - comparing, **561**
 - liquidating, **561**
 - making, **561**
 - selling, **561**
- irrational, **46**
- joint life annuity, **716**
- Kepler, Johannes, **51**
- kicker, **216**
- Knuth, Donald, **297**
- Lander, Leon J., **292**
- Laplace, Pierre Simon, **51**
- Laplace, Pierre-Simon, **354**
- law of large numbers, **121**
- leaf, **375**, **378**
 - leaves, **375**
 - vertex, **375**
- Leibniz, Gottfried, **81**
- length, **150**, **179**, **180**
 - of a sequence, **150**
- letter, **151**
- life annuity, **716**, **719**
- limit, **66**
- linear interpolation, **684**, **685**
- linear interpolation formula, **682**, **684**, **685**, **691**, **695**, **702**
- list, **142**, **179**, **180**, **181**, **195**, **197**
 - from a set, **181**
 - length of, **179**, **180**
- lists, **150**, **201**
- loan, **614**, **640**, **641**, **643**
- logarithm property, **52**
- logarithms, **52**
 - applications, **50**
 - area definition
 - discovery of, **80**►, **481**
 - area definition of, **56**
 - Bernoulli's limit for, **66**
 - defining property, **52**
 - exponent interpretation, **77**
 - horizontal lines, intersection, **69**
 - increasing property, **67**
 - log₁₀, **49**, **60**
 - natural, **56**, **569**
 - other properties, **54**
 - take on all real values, **69**
 - tend to ∞ slowly, **69**
 - uniqueness of ln, **60**
 - using calculators to find values, **63**
- lottery
 - pick3, **92**
 - pick4, **90**
 - win4, **93**
- lower limit, **36**
- mad cow disease, **122**
- Madoff, Bernie, **640**
- magic factor, **530**, **587**
- Malkiel, Burton, **640**
- martingale system, **454**
- mean, **477**
- Mercator, Nicolaus, **81**
- methods
 - bisection, **679**, **692**
 - compound interest, **543**
 - divide and conquer counting strategy, **247**
 - expected value, **449**
 - eyeball, **681**, **692**, **694**
 - future amortization, **621**
 - hunt-and-peck, **690**, **713**
 - linear interpolation, **685**
 - permutation, computing, **187**
 - present amortization, **644**
 - quadrant, **240**
 - simple interest, **518**
 - tree diagram, **380**
- Moore's law, **598**►, **598**, **4604**
- multiset, **296**, **296**, **297**
 - see also abomination, **296**
- multiset formula, **296**
- mutually exclusive, **299**, **300**, **310**, **311**
 - do not assume without *checking*, **311**



- Napier, John, 51
- national debt, 26
- natural logarithm, 56, 569
- natural selection, 410
- negative
 - for a tested character, 367
- negative predictive value, 367
- Newton, Isaac, b
- nominal rate, 514, 522
- normal, 489
- normalize, 493, 494
- not, 236
- null hypothesis, 476, 476, 477
- number
 - prime, 290
- one period principle, 530
- or, 227, 232
 - disjoint union, 232
 - eitherorboth, 227
 - orelse, 232
 - union, 227
- oracle, 140
- order, 144
- order of operations, 12
- ordered pair, 157
- outcome, 299, 300, 301, 305
- outlier, 119, 510
- overcounting, 266
- Pacioli, Fra Luca, 576
- pair, 279, 283
- parentheses, 12
- parentheses rule, 12
- Parkin, Thomas R., 292
- partition, 377, 381
 - of a sample space into events, 377
- payment, 615, 643
- period, 513, 517, 518, 528
 - compounding, 528
- periodic rate, xvi, 513, 522
- periods per year, 515
- permutation, 185, 195
- permutations, 185
 - method for computing, 187
 - of a set, 185
- phenotype, 373
- Phillip's Law, 256
- Phillips, Wendell, 256
- pogemdas, 18
- point, 384
- points, 707
 - rule-of-thumb for, 711
- poker, 29
- population, 591►, ◀598
- positive, 420
 - for a tested character, 367
- positive predictive value, 367
- Posterior knowledge, 395
- power set, 163, 164
- power set counting formula, 164
- precedence, 18
 - breaking ties, 19
 - rules, 18
- predictive value
 - negative, 366
 - positive, 366
- present amortization
 - formula, 643, 717
 - method, 644
- present value, 537
 - ratio to future value, 545
- prevalence, 127, 129, 132, 338, 366
- prevalences, 132
- prime, 290
- principle
 - bang zoom, 161
- prior expectation, 395
- probability
 - conditional, 340, 360
 - for equally likely outcomes, 340
 - equally likely outcomes
 - conditional, 340
 - intersection, 360
 - simple, 360
- probability distribution, *see* probability measure
- probability measure, 305, 307
- product, 157, 330
 - Cartesian, 157
 - direct, 157
- product set counting formula, 157
- proof, 292
- Pythagorean Theorem, 44
- Pythagorus, 44
- Q-diagram, 238
- Q-problem, 238
 - proportional, 239
- Q-set, 238
- quadrant method, 240
- quarter, 513

- radioactivity, 604►, ◀613
- random, 327
- random variable, 441, 441
 - normalized, 494
- random variables
 - independent, 466
- randomvariables
 - independent, 464
- range, 447
- rate
 - effective, 566
 - inflation, 582
 - nominal, 514
 - periodic, 513
- ratio, 37
- recessive, 373, 373
- reflection in the line $y = x$
 - definition of $\exp(x)$, 71
 - picture of, 70
 - swaps x and y coordinates, 71
- restrictive “and”, 221
- risk
 - moral, 635
- root, 375
- rounding
 - first rule of, 30, 33
- rules, laws and morals
 - AMOANS, 249
 - conviction is not knowledge, 291
 - Phillip’s Law, 256
 - preservation of shorthands, 251
 - see similarities, not differences, 258
 - Think twice, calculate once, 266
- rules, laws, morals and warnings
 - “The checks for ‘mutually exclusive’ and for ‘independent’ are *completely* different”, 311
- and morals
 - bang zoom, 160
 - bttdltr, 19
 - calculator parentheses, 17
 - fractions, converting to decimals, 29
 - law of large numbers, 121
 - of exponentials, 49
 - parentheses rule, 12
 - points, 711
 - predicting the future, 630
 - risk versus term, 635
 - rounding, first, 30, 33
 - trusting calculators, 554
 - givens are in the subject, 350
 - givens are known, 349
 - guessing
 - first rule of, 193
 - second rule of, 193
 - Hardy-Weinberg principle, 409
 - leave probabilities as fractions, 314, 319
 - periodic rate rule, 515
 - pogemdas, 18
 - precedence, 18
 - saving
 - first rule of, 633
 - shorthands remember, 248
- Runner™, Road, 42
- runs, 437
- sample space, 224, 236, 299, 300, 301, 305
- saving, 614, 615, 641
 - first rule of, 633
- scientific notation, 26
- seeing patterns, 541
- semester, 513
- sensitive, 396
- sensitivity, 366, 367
- sequence, 150, 194, 196, 327
 - binary, 154
 - length, 150
- sequence counting formula, 155
- sequences, 150
- set, 139
- set braces, 141
- “\$7’ll get you \$12”, 85, 398
- shoot, 384
- shorthand, 194, 195
 - for list count, 195
 - for sequence count, 195
 - for subset count, 195
- Sic bo, 108
- simple interest
 - formula, 513
 - method, 518
- simple interest approximation, 548►, 550, ◀551
 - future amortization, 625
 - present amortization, 648
- simple probability, 360
- Simpson’s paradox, 99, 101
- slide rules, 53
 - multiplication using, 53
- South Sea Bubble, b

- spam, **395**
- specific, **396**
- specificity, **366, 367, 370**
- spot card, **158**
- squared deviation, **478**
- stable, **407, 408, 409**
- standard deviation, **493**
- standard divisions, **247**
- straight, **279, 283**
 - flush, **279**
- straight flush, **279, 283**
- subset, **145, 145, 195, 197, 299**
- success, **420**
- suit, **158**
- sum, **615, 620**
 - intermediate, **615**
- summation
 - arithmetic, **36**
 - geometric, **37**
 - lower limit, **36**
 - upper limit, **36**
- summation notation, **36**
- tables
 - borders, **362, 363**
 - interior, **362, 363**
- term, **35, 512**
 - conversion formula, **518**
 - negative, **539**
 - positive, **539**
 - sign of, **539**
- term conversion formula, **534**
- test
 - for a character, **366**
- test negative, **366**
- test positive, **366**
- theorems
 - irrationality of $\sqrt{2}$, **46**
 - Pythagorean, **44**
- Think twice, calculate once, **266**
- three card Monty, **85, 87, 398**
- three-of-a-kind, **279, 283**
- tree diagram, **375**
 - method, **380**
- trial, **305**
- trichromacy, **363, 364, 364**
- two pair, **279, 283**
 - house advantage, **453**
- union, **227, 232**
 - of two sets, **227, 232**
- United States Senate, **356**
- universal, **489**
- universal set, **224, 299**
- universe, **224**
- upper limit, **36**
- value, **158, 279**
- values
 - consecutive, **279**
- variance, **478**
- Ward, E. M., **b**
- watch lists
 - National Security Strategy, **87**
 - No-Fly list, **87**
 - Terrorist Watch List, **88**
- Weinberg, Wilhelm, **410**
- Wiles, Andrew, **291**
- worse guess, **690**
- yield, **562**
 - annualized, **563, 564**
 - continuous, **568**
- zygosity, **372, 404**
- zygotic frequencies, **373, 373**
 - stable, **407, 409**
- U.S.D.A., **122**
- undercounting, **267**
- unfair game, **453**

