

## CONTENTS

Preface	xi
<b>1</b>	<b>Life in a Random Universe 1</b>
1.1	<i>Twenty-First-Century Science: An Example</i> 3
1.2	<i>A Small-Data Example</i> 5
1.3	<i>Newton's Perfect World, and Ours</i> 7
1.4	<i>Statistical Inference</i> 11
1.5	<i>Are Those Genuine Basquiats?</i> 12
1.6	<i>The 18 Baseball Players: Learning from the Experience of Others</i> 14
1.7	<i>The Statistical Skeptic</i> 20
<b>2</b>	<b>Probability, Fortune, and the Paradox Dice 23</b>
2.1	<i>Randomness and Probability</i> 24
2.2	<i>Probabilities and Proportions</i> 28
2.3	<i>Where Do Probabilities Come From?</i> 36
2.4	<i>The Paradox Dice</i> 38
2.5	<i>The Bell-Shaped Curve</i> 43
<b>3</b>	<b>Reverend Bayes and the Physicist's Twins 50</b>
3.1	<i>Bayes' Rule</i> 51
3.2	<i>Hume, Bayes, and the Existence of Miracles</i> 53
3.3	<i>Will the Sun Rise Tomorrow?</i> 55
3.4	<i>Bayes' Rule: Its Fun Side and Dark Side</i> 59
3.5	<i>Why Isn't Everyone a Bayesian?</i> 62
3.6	<i>Bayesians and Frequentists</i> 67

viii CONTENTS

- 4 The Gold Standard of Comparisons 69
  - 4.1 *The Most Accurate Comparison Ever Made* 70
  - 4.2 *Statistical Significance: A Breast Cancer Study* 72
  - 4.3 *Are There More Red Cars or Blue Cars?* 77
  - 4.4 *Bias, Blinding, and Placebos* 80
  - 4.5 *Permutations and p-Values* 84
  - 4.6 *Observational Studies* 89
- 5 A Statistician's Guide to the Medical News 91
  - 5.1 *Is Anything Good for You?* 93
  - 5.2 *Is Anything Bad For You?* 96
  - 5.3 *The Health Mountain and the Health Plateau* 101
  - 5.4 *A Scorecard* 105
  - 5.5 *Are All Published Articles Wrong?* 114
- 6 Correlation, Causality, and Regression to the Mean 119
  - 6.1 *Correlation and Heredity* 120
  - 6.2 *Regression and the Method of Least Squares* 125
  - 6.3 *The Bell-Shaped Mountain* 130
  - 6.4 *Causality and Randomization* 134
  - 6.5 *Causation without Randomization* 137
- 7 Estimation, Accuracy, and the Bootstrap 147
  - 7.1 *Learning from Noisy Data* 148
  - 7.2 *Surveys, Samples, and the Bootstrap* 152
  - 7.3 *Sample Size  $n$  and Its Slower Friend the Square Root of  $n$*  160
  - 7.4 *The Square Root Law and the Bell-Shaped Curve* 165
- 8 Decisions and Regrets 170
  - 8.1 *Errors of the First and Second Kind* 172
  - 8.2 *Making the Best Possible Decisions* 178
  - 8.3 *Twenty-First-Century Significance Testing* 183
  - 8.4 *The Drunkard's Walk* 191

<b>9</b>	<b>Survival, Prediction, and the Perverse Law of Waiting</b>	<b>195</b>
9.1	<i>Hazard Rates and Risky Behavior</i>	197
9.2	<i>Censored Data and Waiting Times</i>	202
9.3	<i>Prediction</i>	207
9.4	<i>Your Guess Is Not as Good as Mine</i>	210
<b>10</b>	<b>Thinking Like a Statistician</b>	<b>217</b>
<b>Appendix</b>	<b>223</b>	
A.1	<i>Chapter 1</i>	223
A.2	<i>Chapter 2</i>	227
A.3	<i>Chapter 3</i>	232
A.4	<i>Chapter 4</i>	239
A.5	<i>Chapter 5</i>	246
A.6	<i>Chapter 6</i>	251
A.7	<i>Chapter 7</i>	258
A.8	<i>Chapter 8</i>	263
A.9	<i>Chapter 9</i>	271
	<b>Bibliography</b>	<b>275</b>
	<b>Index</b>	<b>277</b>

# 1

## Life in a Random Universe

I don't believe we know anything with complete certainty, but everything probably and to different degrees of probability.

—CHRISTIAAN HUYGENS

June 8, 2004, was the first transit of Venus visible on Earth since October 6, 1882. Not wanting to miss the big event, I consulted a chart by Richard Proctor giving the precise minute of onset—the time when the small dot of Venus would begin its trip across the face of the sun. The chart was exactly right. Mr. Proctor drew it in 1874.<sup>1</sup>

Impressive! If this kind of steel-trap precision is your idea of “science,” then you would be mostly right in 1874, but considerably less so in the twenty-first century. What has become clear in the past hundred years is that steel-trap science is a luxury enterprise, a high-end shop inside a much bigger scientific emporium. A less precise but more flexible brand of science now shares the stage, helping us make sense of the world we live in, a random universe where truth most often lies concealed behind a foggy curtain of noise.<sup>2</sup>

1. The † mark in the margin indicates topics that are discussed further, along with references, located in the chapter's appendix.

2. “Noise” is shorthand for the kind of random disturbances that afflict observational data. Look ahead to figure 1.4 for an example.

Statistics was the first information science, preceding its upstart nephew, computer science, by two centuries. Astronomers have stars, geologists rocks, chemists elements, and statisticians have information. There is a big difference: stars, rocks, and elements are part of the natural world, but information is man-made, which means that statisticians don't have nature to fall back on as a check on their theories.

*Statistical inference*, what I'll be talking about in this book, is a collection of powerful ideas designed to see the truth behind the foggy curtain, and no more than the truth. The book's title, *To Think Like a Statistician*, presupposes that this is a good thing. It is! If sciences are judged by how often they appear in public media, then probability and statistics are superstars. My hope is to give you, the reader, a pleasant familiarity with the tricks of the trade, without equations, theorems, or long-winded tables.

"Tricks" isn't quite the right word. There's some underlying philosophy involved, but mainly I want to talk about a collection of intriguing topics, partly familiar in the news but still mysterious, and the methods statisticians have devised to understand them:

- Correlation and causation (Do tall children cause tall parents?)
- Bayes' Rule and its contentious history
- Reading the medical news (Is anything good for you? Bad for you?)
- Prediction (Why will your championship team probably do worse next year?)
- Accuracy and the bootstrap (What does "plus or minus" or "margin of error" really mean?)
- Crucial comparisons (How do we know the new drug is better than the old?)
- Survival and the perverse law of waiting
- Learning from the experience of others (I didn't have any accidents, so why did my car insurance premium go up?)

I have in mind a grand tour of the statistical point of view, hitting the high points while avoiding textbook prose. This chapter offers a preview of the tour, touching on several examples of statistical thinking in action in science, sociology, medicine, art, and sports.

These are powerful ideas, of 250 years' gestation. They enable scientists, and everyone else, to navigate more confidently through a random universe. The current era of Big Data (sometimes *really big*, as in the

example of the next section), is testing traditional statistical methods to the limit, and adding some new ones that we'll visit later in the book. All the ideas, new and old, reflect the statistician's approach to making sense of our random universe.

## 1.1 Twenty-First-Century Science: An Example

Watching the construction of a new data science building across from my office impressed on me the nonchalance of girder walkers and how unsuited I'd be for that line of work. Statisticians aren't in danger of falling, but some nonchalance is called for when faced with a data set like the genomic study described next. It's so darned big! A trillion numbers, enough to fill a million thick volumes if ever printed out. The technology that produced twenty-first-century computers has also given us the twenty-first-century experimental equipment necessary for genomic research.

Figure 1.1 is what is called a *Manhattan plot* because of its supposed skyline appearance. It illustrates the results of a study concerning the genetic basis of high LDL levels (the "bad cholesterol" that is thought to cause heart attacks). There are 2.4 million dots on the diagram, not all of which are visible, corresponding to 2.4 million positions on the human genome. That sounds like a lot, but it's a tiny fraction of the more than three billion "base pairs" on the total genome, each denoted by one of the four letters of the genetic code A-T-C-G. †

But these are a special 2.4 million positions. Called *single nucleotide polymorphisms* or SNPs (pronounced "snips"), they are positions where there is substantial variation in which letter occurs: maybe G 90% of the time but C the other 10%. Neucleotide variation can cause dramatic differences in disease rates—perhaps the C people are more prone to an illness of interest—and is just the kind of thing the researchers were hoping to find. *Genome-wide association study* (GWAS) is the popular name for this kind of wide-ranging investigation ("fishing expedition" being a less kind term).

The study included 190,000 subjects, each of whom had all of their SNPs determined and their LDL measured. The basic idea was simple enough: look for SNPs where an unusual pair of letters (one from each parent), say, Cs where there are usually Gs, was associated with high LDL. A Nobel Prize kind of result would have uncovered one or two killer SNPs that perfectly predicted high LDL, but that's not at all

4 CHAPTER 1




FIGURE 1.1. Manhattan plot of a study concerning the genetic basis of high LDL levels. Heights of the points measure effects on high LDL of an individual SNP. Of the 2.4 million SNPs measured, 3078 were significantly associated with high LDL, seen rising above the light horizontal line, and further concentrated in 58 *loci* to form the observed “skyscrapers.” For a more in-depth discussion of figure 1.1 and the Manhattan plot, see appendix A.1.

how things turned out, as you can see from all those skyscrapers in the Manhattan plot.

This is where statistical inference comes in. The height of each dot is a measure of correlation between the letter of that SNP (A, T, C, or G) and the subject’s LDL levels. High points (those above the horizontal line) indicate important SNPs, ones that increase LDL. A great deal of caution is needed here. We are searching through needles-in-haystacks land: with 2.4 million SNPs to consider, the overwhelming majority of which are not predictive, random variation alone could produce a torrent of “false positives.”

The height of the horizontal line is set high enough to avoid most of those (the discussion in appendix A.1 gives some idea of how this is done); 3078 SNPs cleared the bar. These were concentrated at 58 positions, or *loci*, along the genome, shown as skyscrapers in the Manhattan plot.

There is no question that the 3078 SNPs are related to high LDL, but as with many GWAS investigations, the results were predictively

underwhelming: combined with the 100 or so previously identified loci, only 12% of the observed variability of LDL levels could be explained. (The universe seems particularly noisy at the genomic level.) The GWAS can be thought of as an opening salvo in science's attack on the LDL question, helping map out the landscape for future investigation.

A GWAS is obviously "scientific," but it is not Mr. Proctor's kind of science. Transits of Venus are predicted from Newton's laws of motion, beautiful in their simplicity and precision. Medical science has made, and will continue to make, progress in understanding disease, though without any prospect of a precise Newtonian formulation. The human body is a miraculous assemblage, but compared with celestial mechanics it is a soft machine that follows its rules more statistically than punctiliously, and requires statistical analysis (that is, an analysis allowing for variability and contrary events) for its interpretation.<sup>3</sup>

## 1.2 A Small-Data Example

With 190,000 subjects and 2.4 million SNPs per subject, the LDL study makes it into anyone's definition of Big Data. Not only has data gotten big in the twenty-first century, but the technology that produces all those numbers has mushroomed, too.

Small data hasn't gone away, though. It's been a substantial part of my consulting work in the Stanford School of Medicine and is featured in many, maybe most, of the examples I use to illustrate statistical thinking in action.

Table 1.1 concerns a small-data inference problem. A European car insurance company has tabulated the number of claims per customer made in 1950: 7840 customers made zero claims each (ideal from the company's point of view), 1317 made one claim each, and so on, up to one customer, not the most careful driver, who made seven claims. Now we only have eight numbers to look at, instead of a trillion, but they still manage to conceal a surprise.

3. Celestial mechanics has its limits too. Mr. Proctor could not have accurately predicted the next 1000 transits. Henri Poincaré in the early 1900s showed that tiny disturbances would eventually propagate out of control (the "butterfly effect"). Perhaps I should say that the randomness component is much smaller for orbits than for disease occurrence, but smaller isn't zero.

TABLE 1.1. Number of claims per customer made in 1950 for a European car insurance company; 7840 customers made no claims in 1950, 1317 customers made 1 claim each, 239 made 2 each, etc.

Claims	0	1	2	3	4	5	6	7	Total
Customers	7840	1317	239	42	14	4	4	1	9461

The company wants to set premiums for 1951. It would surely help to know how many claims each customer *will* make in the coming year. One guess is that customer  $x$  will make as many claims again in 1951 as in 1950, which would suggest charging the 7840 zero-claimers next to nothing. That can't be right. . . . Or, the company could assume that everyone would make 0.214 claims in 1951, the 1950 average, but that would treat the zero-claimers no better than the seven-claimer.

Amazingly, a much better answer was discovered in 1956. *Robbins' Rule* says that the expected number of claims in 1951 for a 1950 zero-claimer is the number of 1950 one-claimers divided by the number of zero-claimers, or

$$1317/7840 = 0.168.$$

That's only about one-sixth of a claim per person, not much but definitely more than zero. In the same vein, the expected number of 1951 claims for a 1950 one-claimer is

$$2 \times 239/1317 = 0.363,$$

or a little more than twice the 0.168 of the zero-claimers, with similar formulas for 1950 two-claimers, three-claimers, etc., as explained in appendix A.1.

Table 1.2 shows Robbins' Rule for 0, 1, 2, 3, and 4 claims in 1950. (The number of 1950 claims greater than four in table 1.1 gets too small for the Rule to be accurate.) Following Robbins' Rule lets the company set their 1951 rates according to the number of expected claims—less for 1950 zero-claimers than others, but not nothing.

At this point, the zero-claimers might launch a class action lawsuit on the grounds that they are being forced to pay for the carelessness of the one-claimers. This gets us to a central precept of statistical thinking: that we can learn from the experience of others. Doing so requires us to subsume our individual experience (“I had a perfect driving record last year!”) in the pooled information from a relevant group. How to

TABLE 1.2. Robbins' Rule for the expected number of 1951 claims given the customer's number of 1950 claims.

1950 claims	0	1	2	3	4
Robbins' Rule	0.168	0.363	0.527	1.33	1.43

do so in a rational and effective way is what it means to think like a statistician.

The LDL GWAS involved many billion times more data than the insurance example, but its analysis also depended on learning from the experience of others: the other 2.4 million SNPs. To be deemed “predictive,” a SNP needed to clear the horizontal bar in figure 1.1, and the height of the bar was determined by the results from *all* the SNPs, which is where the learning came in. Both the LDL and insurance examples depend on Bayesian thinking (chapter 3), one of the linchpins of statistical inference. The 18 baseball players seen in section 1.6 will make the same point again, in more dramatic fashion.

### 1.3 Newton's Perfect World, and Ours

Isaac Newton (1642–1727) was a two-fold genius: as what we would now call a scientist (but then a “natural philosopher”) and as a mathematician of the first rank. In combination, these talents produced a view of the world that has dominated science ever since. Newton believed that

1. There are inflexible laws that underlie the workings of nature.
2. These laws can be expressed in mathematical terms.

What I called “steel-trap science” is in fact Newtonian science, a symbiosis of mathematical logic and careful observation.

The success of Newton's laws had a bracing effect on humanity's intellectual confidence, as if a strong set of spectacles had snapped the world into sharp focus. Consider his Second Law of Motion:

$$\text{acceleration} = \text{force}/\text{mass}.$$

Our day-to-day world is filled with first things pushing on second things, and if you push harder they usually move faster. The second




FIGURE 1.2. In Newton’s perfect world, acceleration (plotted vertically) exactly equals force divided by mass.

law says precisely how this happens, depending on the force of the first thing and the mass of the second, and not on their colors or temperatures or heights above sea level or anything else. Its more usual statement, “ $F = MA$ ” (force equals mass times acceleration), was the world’s most famous equation before Einstein’s “ $E = mc^2$ .”

Figure 1.2 shows Newton’s Second Law in graphical form. This is a 3D graph: you’re meant to imagine the two “predictor variables” of force and mass, located on the surface of your desk, with the “response variable” of acceleration plotted vertically above them. Low mass and high force, at the far corner of the desk, produces the greatest acceleration. In Newton’s hands, and that of successors like Laplace, the laws of motion enabled humans to land softly on the moon, without missing it or being smashed to bits.

How did Newton discover his laws? First of all, he was a prodigious observer (apples falling from trees, etc.). Second, he had a formidable imagination and could theorize general truths that had eluded millennia of earlier savants. Of course, he might also have done




FIGURE 1.3. What Newton might have seen in the real world.

experiments, trying different forces on different masses and measuring the resulting accelerations. Perhaps not though: acceleration requires careful measurement of time, just what *wasn't* available in the mid-1600s.<sup>4</sup>

Figure 1.2 portrays a perfect world, one of absolute precision. This was the world Richard Proctor was working in, and much of contemporary science still thrives there. It isn't, however, the world of our book. Figure 1.3 pictures, rather fancifully, what Newton might have seen if he'd actually tried to learn the Second Law by experiment. Now we are in an imperfect world. Noise has corrupted a smooth underlying reality into what looks like a dangerous version of Disney's Space Mountain.

Newton was able to look at figure 1.3 and visualize figure 1.2. Physicists are still trying to uncover a perfect underlying reality through

4. Galileo, a half century earlier, did attempt such measurements, using marbles rolling down inclined planes, with his own pulse and a water clock as timing devices. Evidently, the effort was unfruitful.




FIGURE 1.4. Kidney function of 159 healthy volunteers plotted against their age; a smooth curve has been fitted to the points. From the laboratory of Dr. Brian Myers.

keen observation and mathematical logic, string theory being a particularly heroic example, but there aren't that many Newtons to go around and reality itself may not be all that perfect in a GWAS kind of environment.

Most often, statisticians are called on to help scientists during the imperfect-world portion of a study. Figure 1.4 reports on a study of 159 healthy volunteers enrolled in a nephrology (kidney) investigation. There are 159 points, each showing a volunteer's kidney function (plotted vertically) and age (plotted horizontally). Learning how kidney function decreases with age was the goal of the study, but the term † *scatterplot* seems all too appropriate here. In the absence of Newton, I've fit a smooth curve to the points. Does this capture the hidden reality of kidney function versus age? Certainly not exactly, but the curve is evocative of a slow early decline getting steeper past age 50. Younger donors are preferred for kidney transplants, as the figure suggests, but they can be in short supply. Chapter 7 says a little more about the kidney study.

The traditional statistical worldview is typically one of a smooth and orderly underlying reality hidden from sight by corrupting layers

of random noise. The job of the statistician is to strip away the noise from the picture. How this is done, and to what degree it's possible, is what I try to describe in this book.

## 1.4 Statistical Inference

Inference is one of those miraculous human abilities that we use, and have always used, frequently and automatically:

“There’s something rustling in the grass: maybe it’s a saber-toothed tiger!”

or, more recently,

“That car is a block away: I can cross the street safely.”

*Statistical inference* is what statisticians do for a living. A powerful way of thinking about evidence, it depends upon hard-won theories of learning from experience, especially experience that arrives a little bit at a time in small, noisy, and sometimes contradictory pieces (as in the kidney study). We’ve already made a few statistical inferences, in the LDL GWAS, the insurance company’s prediction problem, and the kidney study of figure 1.4.

Inferential theories—there are two main ones we’ll be meeting, Bayesian and frequentist—are amalgams of mathematics, philosophy, and applied science. Here we avoid most of the math, relying on examples, charts, graphs, and figures to get across the big ideas, along with some aspirational hand-waving. The theories have been, by necessity, adapted to the computational limits of the day. In our day, computation has sped up by factors of millions. Statisticians now stride in seven-league boots, using massive numerical resources to execute statistical inference on Brobdingnagian scales (think GWAS or “Deep Learning”), but the basic inferential ideas have mostly stayed the same.

Deduction is the process of drawing specific conclusions from given axioms. *Induction* is the opposite process: arguing from observed cases back to an underlying mechanism. For instance, one inference from the 159 cases in figure 1.4 was quicker decline of kidney function after age 50. Philosophers have disagreed on whether induction is actually possible—

“That car looks blue but I can only see one side; maybe the other side is red?”

— but life would be hopelessly muddled without it.

Statistical inference is the weaponized version of induction, putting it in action in the cause of science, medicine, society, and politics. There are obvious pitfalls lurking here. Induction can be a dangerous game compared to slow-but-safe deduction. A large part of statistical theory is defensive, aimed at avoiding inferential errors, as I’ll discuss later.

### 1.5 Are Those Genuine Basquiats?

Not all inference problems are numerical in nature. Here is the beginning of an article from the *New York Times* of May 30, 2022:

**FBI Investigating Basquiat Painting Show at Orlando Museum of Art**—The ongoing fascination with the life and work of Jean-Michel Basquiat shows little sign of dimming whether in the form of brisk sales for \$29.99 Basquiat-theme t-shirts at the Gap, large crowds for Basquiat’s latest art exhibitions, or an actual canvas by the painter auctioned last week for \$85 million.

Jean-Michel Basquiat was an African American painter, a friend and protégé of Andy Warhol, who died at age 28 in 1988 during the height of international praise and fame, which has recently intensified. The same person who paid that \$85 million, Yusaku Maezawa, then bought one of Basquiat’s striking “Black Skull” paintings for \$110 million, a record for work by an American artist, which is to say that there is big money involved here. The article continues:

The FBI’s Art Crime Team is investigating the authenticity of 25 paintings that the Orlando Museum of Art says were created by Basquiat and are on exhibit there. . . . The paintings in the “Heroes and Monsters: Jean-Michel Basquiat” exhibition were said to have been recovered from a Los Angeles storage unit in 2012. . . . The paintings’ owners and the museum’s chief executive Aaron De Groft say the paintings are genuine Basquiats, citing statements from art world experts commissioned by the owners. . . . The paintings are set to leave the museum on June 30 for public exhibitions in Italy.

The FBI's interest reflects allegations of tax evasion and perhaps money laundering in the high-end art world and, in this case, the tangled provenance of the 25 paintings. Worked on cardboard, they were discovered in 1982 and said to have been sold by the artist for \$5000. The trove was

... bought for about \$15,000 by William Force, an art and antiques dealer, and Lee Mangin, a retired salesman. A third owner is the Los Angeles trial lawyer Pierce O'Donnell, who in 2016 represented Amber Heard in her divorce from Johnny Depp and Angelina Jolie in her divorce from Brad Pitt.

This is where the *Times* story gets more pointed:

Much of the backstory establishing the paintings' origins rests largely on the word of Mangin and Force, who have both served time in prison for felony drug trafficking under different names. . . . In 1996 the Securities and Exchange Commission arrested Mangin for securities fraud, alleging that Mangin was part of a criminal ring that forged documents. . . . O'Donnell also has a criminal record. . . . Richard LiPuma, a lawyer for Mangin, said that the provenance of the paintings was "airtight," and the fact that the owners were once in trouble with the law was irrelevant to the question of whether the works are genuine. . . .

Not so fast, LiPuma. "Innocent until proven guilty" doesn't apply to the world of inference. A thoughtful oil sheikh, in considering whether to lay out \$100 million for a Basquiat, would probably prefer non-criminal records for the sellers. On the other hand, the sheikh might give more credence to the owners' art experts, whom the rather one-sided article ignores: they and the museum must have had some reason for putting their reputations on the line in an obviously contentious situation.

Other evidence, more con than pro, is reported in this article, such as from a handwriting expert (Basquiat wrote a few words on his paintings) who may or may not have verified authenticity, and an "independent brand expert consulted by the *Times*," who said that the Federal Express imprint on the cardboard backs wasn't used by FedEx until 1994, six years after the artist's death.

At this point, the sheikh might wish to hire an "independent Bayesian expert." Bayes' Rule, historically the first word in statistical

inference, provides a rational way to update one's opinions as new evidence arrives. "New evidence" often takes the form of numerical data, but it can also be vague and qualitative, as with the maybe-  
† Basquiat paintings.

Chapter 3 applies the Rule to an important inferential question that arose in 2020: Was the drug hydroxychloroquine effective against COVID-19? The venue there was the interface between politics and medicine rather than art and money, but it was even more contentious. After the hydroxychloroquine discussion, I'll encourage you (with some hints) to try a Bayesian analysis of the Basquiat authenticity question.

Gathering information from different sources is a trademark of statistical thinking, Bayes' Rule being just one of the methods we'll discuss. The problem of the 18 baseball players in the next section shows how information can accrue in surprising ways.<sup>5</sup>

## 1.6 The 18 Baseball Players: Learning from the Experience of Others

Asked to free-associate on the word *statistician* (and told to avoid *nerd* and its cognates), one is likely to get "data" and "numbers" first out of the box. The Basquiat story shows that statistical evidence doesn't have to be numerical, but most often it is. Numbers are to statisticians as words are to poets: the raw material from which gold can be mined. What follows is an example of the mining process applied to the world of sports. The setting is baseball, but as will become clear, the reader who knows nothing about baseball will have an advantage. Prediction, not baseball, is the real subject here.

Early in the 1970 baseball season, Carl Morris<sup>6</sup> was looking through the sports pages of the *Los Angeles Times*. It listed the batting averages, the proportion of successful attempts to hit the ball, for all Major League players, several hundred of them. Wanting to deal with a smaller data set, Morris focused his attention on those players with

5. Here's a more personal inference story: I was walking downtown for lunch one day when I was approached by a young policeman who said, "Are you from the Veterans' Home?" "No," I said, "I teach statistics. Why do you ask?" "Someone has wandered away." He was gone before I realized the need for a snappy comeback.

6. Then at the RAND Corporation, later as Professor of Statistics at Harvard.

TABLE 1.3. Early batting averages for 18 Major League baseball players after 45 at-bats in the 1970 season. Two predictions of their final averages for the season are shown: the early average and the “James–Stein” prediction.

	Hits	Early average	James–Stein prediction	Final average
Clemente	18	.400	.294	?
F. Robinson	17	.378	.289	?
F. Howard	16	.356	.284	?
Johnstone	15	.333	.279	?
Berry	14	.311	.275	?
Spencer	14	.311	.275	?
Kessinger	13	.289	.270	?
L. Alvarado	12	.267	.265	?
Santo	11	.244	.261	?
Swoboda	11	.244	.261	?
Unser	10	.222	.256	?
Williams	10	.222	.256	?
Scott	10	.222	.256	?
Petrocelli	10	.222	.256	?
E. Rodriguez	10	.222	.256	?
Campaneris	9	.200	.251	?
Munson	8	.178	.246	?
Alvis	7	.156	.242	?

exactly 45 at-bats (that is, 45 attempts to hit the ball, regardless of success). There were 18 such players, whom Morris listed in decreasing order of their successful hits, as shown in table 1.3. At the top of the list was Roberto Clemente with 18 hits out of his 45 at-bats, for a batting average of

$$18/45 = 0.400.$$

(Morris chose 45 in order to include Clemente, a favorite of his.)

Morris then asked a classic inferential question: How well could he predict each player’s final batting average for the remainder of the season (about 400 more at-bats)? The players bat independently of each other, so Munson’s miserable performance could have no effect up or down on Clemente, and likewise for any of them. With this in

mind, the best predictor for future performance seems certain to be past performance: 0.400 for Clemente's season, 0.178 for Munson's, etc.

Among other things, Morris was a leading sports statistician, but it wasn't just baseball that led him to table 1.3. The recently invented *James–Stein Rule* offered a different if controversial prediction formula that he wanted to test out. The third column of table 1.3 gives the James–Stein predictions. These are based on the same early averages in the second column but with a twist: for the James–Stein predictions, each player's early average is shrunk a certain proportion of the way to the average of the early averages (equaling 0.265 in this case). The clever algorithm that says what "certain" means is discussed in appendix A.1, but its exact expression isn't needed to see its effect here.

Time went by. In October, after the end of the baseball season, Morris was able to fill in the question marks in the last column of table 1.3. Column 3 of table 1.4 shows the season's final averages for the 18 players. We have two possible predictors for the final average: column 1 shows the early averages while column 2 shows the James–Stein predictions. Which did the best job of predicting the final averages?

Table 1.5 makes the comparison. Column 1 lists the prediction errors made by the early averages. For example, Clemente's early average was .400 while his final average was .346, for a prediction error of .054 (i.e.,  $0.400 - 0.346$ ). His James–Stein prediction of .294 underestimated his final average of .346, for a prediction error of  $-0.052$  (i.e.,  $0.294 - 0.346$ ). For both methods, some of the prediction errors are positive, some are negative, but you can see that the James–Stein errors are usually smaller in size.

At the bottom of table 1.5, I've added up the squares of the prediction errors for each method (squares in order to count positive and negative errors equally). The sum of squared errors provides an overall assessment of prediction inaccuracy. It is more than three times worse,

$$0.075/0.021 = 3.6,$$

for the early averages than for the James–Stein predictors. This shows a stunning difference in predictive performance! Nobody cares much about 50-year-old baseball statistics, but if we were predicting cancer cure rates at 18 hospitals, a factor of three-times-less-error might be life-saving. This wasn't a fluke. James and Stein used inferential theory to prove that their method will always be expected to give better results.

TABLE 1.4. Now showing the players' final batting averages for the 1970 season.

	Early average	James–Stein prediction	Final average
Clemente	.400	.294	.346
F. Robinson	.378	.289	.298
F. Howard	.356	.284	.276
Johnstone	.333	.279	.222
Berry	.311	.275	.273
Spencer	.311	.275	.270
Kessinger	.289	.270	.263
L. Alvarado	.267	.265	.210
Santo	.244	.261	.269
Swoboda	.244	.261	.230
Unser	.222	.256	.264
Williams	.222	.256	.256
Scott	.222	.256	.303
Petrocelli	.222	.256	.264
E. Rodriguez	.222	.256	.226
Campaneris	.200	.251	.285
Munson	.178	.246	.316
Alvis	.156	.242	.200

What is going on here? We seem to have a paradox: the 18 players bat independently of each other and classical statistical theory shows that, individually, each player's batting average is the best predictor of his future performance. How can the James–Stein method do better?

The answer is simple but surprising: there is extra information available from the full set of 18 results. In this case, it is information about baseball. You don't have to be a baseball fan to learn from the data in table 1.3. The 18 batting averages fall between 0.156 and 0.400; there aren't any 0.000s or 0.999s, and in fact most of the batting averages are not far from the "average" average (that is, the average of the 18 averages) of 0.265. As explained in appendix A.1 (in particular figure A.2), the James–Stein method takes advantage of these facts to pull each estimate closer to the middle value of 0.265. But for our discussion here, the details aren't as important as the bigger picture of statistical

TABLE 1.5. Comparison of the prediction errors of the two methods.

	Early avg minus final avg	JS prediction minus final avg
Clemente	.054	-.052
F. Robinson	.080	-.009
F. Howard	.080	.008
Johnstone	.111	.057
Berry	.038	.002
Spencer	.041	.005
Kessinger	.026	.007
L. Alvarado	.057	.055
Santo	-.025	-.008
Swoboda	.014	.031
Unser	-.042	-.008
Williams	-.034	.000
Scott	-.081	-.047
Petrocelli	-.042	-.008
E. Rodriguez	-.004	.030
Campaneris	-.085	-.034
Munson	-.138	-.070
Alvis	-.044	.042
Sum of squared errors	.075	.021

methods scooping up information from both likely and unexpected sources.

My definition of statistics as the science of learning from experience applies doubly to the 18 baseball players. At a first level, each player's early average was a statistical summary of his first 45 at-bats. Clemente, for example, might have had this early season history:

```

1 1 0 0 1 1 1 0 0 0 0 0 0 0
1 0 1 0 0 1 0 0 0 1 0 1 1 0 1
0 0 0 1 0 0 1 0 1 0 1 1 0 1 0

```

with 1 indicating a hit and 0 an out. His early-season average of .400 was the average, or mean, of the 1s and 0s.<sup>7</sup>

Clemente's early season .400 average was based entirely on his own experience, and likewise for each of the other 17 players. The James–Stein predictions were based on a second level of statistical learning where the data was the 18 early season averages. In other words, each player was learning from the experience of others.

In my experience helping medical school researchers analyze their data, there is often a wide choice of possible “others”, such as previous studies, related diseases, different age groups, etc. Statistical gems like the James–Stein estimator help point the way for the data analyst, but there is still a lot of art mixed in with the science.

If the James–Stein phenomenon seems bizarre to you, you're in good company. The appearance of their jointly authored paper in 1961 flummoxed much of the statistics community. In the James–Stein world, Munson's bad performance *does* pull down Clemente's prediction (by lowering the average of averages)! How can this be if they are performing independently? The answer has much to do with *Bayesian inference*, which we'll begin investigating in chapter 3, “Reverend Bayes and the Physicist's Twins.”

Bayesian methods bring past knowledge to bear on current problems. A baseball fan would naturally use Bayesian thinking in answering Morris's prediction question. They would know that Clemente's early 0.400 was almost certainly too high for a whole season—there hasn't been a season-long .400 batter in 75 years—and that Munson's 0.178 was too low (he was one of the best hitters of his time). Sports statistics are a friendly venue for Bayesian inference, making sports fans into natural Bayesians.

Not all statistical problems have useful Bayesian knowledge at hand. My hypothetical “cancer cure rates at 18 hospitals” example very well might not, and even if it did, that knowledge might be controversial or undependable. Frequentist inference, the kind we've been tacitly following here, aims to exclude prior knowledge from its methods. We'll

7. It can be shown that the mean is the *best* single-number summary of the full history, but in fact scientific acceptance of the mean is comparatively recent. Given a set of varying observations, eighteenth-century astronomers argued over using the mean versus using only the one observation considered most reliable. Stigler gives an insightful discussion of the debate in the first chapter of *The Seven Pillars of Statistical Wisdom*.




FIGURE 1.5. The northern view in a world where stellar locations and brightnesses have been selected at random. Each star has equal probability of being at any point in the circle.

be going back and forth between frequentist and Bayesian thinking in the chapters ahead.

One definition of statistics is *the science of learning from experience*, especially experience that arrives in small, sometimes contradictory pieces. In this sense, inferential theory concerns the efficient extraction of information. There were literally billions of pieces of data that went into the GWAS investigation of figure 1.1, a library’s worth of hefty volumes, all of which were boiled down to the discovery of a few new groups of significant SNPs.

## 1.7 The Statistical Skeptic

My “peering through the fog” metaphor for statistical inference has two faces: seeing what’s hidden beneath the fog and *not* seeing what’s not there. We’re happy if the saber-toothed tiger in the grass turns out to be a false alarm, but it’s a different matter when a new cancer drug is mistakenly claimed to be effective. Some of the key ideas we’ll be talking

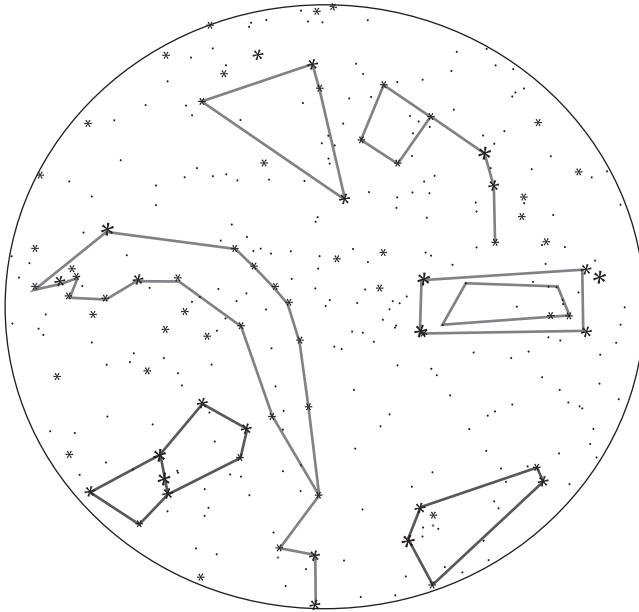


FIGURE 1.6. My constellations in the sky of figure 1.5. Clockwise from top: the Right Triangle, the Dipper, the iPhone, the Coffin, the Serpent, O'Brian and O'Brian's Belt.

about are defensive in nature, designed to protect overenthusiastic scientists from seeing only what they want to see.

Here's an example of seeing what one wants to see. Figure 1.5 pictures the northern sky in a world where the stars have been placed completely at random. I spent half an hour drawing constellations in that sky with the somewhat childish results shown in figure 1.6. Someone with a keener artistic eye could do better, but that's not the point. Our actual constellations tell thrilling stories, and people love stories, but stories get less desirable when we're trying to really understand an intricate system. A newspaper headline like "President Blamed for Inflation" means someone has done some constellation-drawing concerning national politics.

Chapter 4, "The Gold Standard of Comparisons," talks about the statistical guardrails that keep medical science on track: randomization, blinding, hypothesis testing, the famous 0.05 rejection level, statistical significance. . . . When the Hippocratic oath is quoted in the news, it's usually for its directive "do no harm." That's excellent defensive advice even 2500 years later. Human health and sickness are complicated topics, maybe the *most* complicated that science confronts.

That doesn't stop, and in fact seems to encourage, a high volume of health advice in the news: for instance, from the *New York Times* science section:

“An Hour of Running Lengthens Life Expectancy by 7 Hours”

(seeming to offer the prospect of eternal life, but the authors put in a time limit). Chapter 5 presents a statistician's guide to reading the medical news in terms of two competing viewpoints on personal well-being: the Health Mountain and the Health Plateau. It concludes with a health news scorecard you can use to rate the plausibility of stories like the one under the headline quoted above.

(continued...)

## INDEX

- .05 criterion, 86, 187, 240, 251  
.05 threshold, 72, 74–76, 178, 186, 187, 242
- A-T-C-G, 3
- A/B tests, 80
- accuracy, 2, 68, 103, 128, 130, 138, 147–149, 157, 158, 164, 183, 218, 220, 221, 244, 259
- aducanumab (Aduhelm), 82, 87, 174, 264
- algorithm, 16, 66, 67, 85, 87, 129, 130, 158, 187, 209, 211, 220, 225, 228–230, 245, 274  
big hat, 274  
Deep Learning, 11, 146, 168, 211, 213, 216, 221, 274  
least squares, 130, 221, 254  
lowess, 126, 158–160, 166, 263  
pure prediction, 211, 215, 216, 219, 221, 274  
Random Forests, 213–215, 221, 274
- Alzheimer’s disease, 82, 174, 196
- Anodyne, 63, 64, 74, 84–87, 115–117, 241, 245, 252
- area under the curve, 182
- Arianism, 54
- Ars Conjectandi*, 228
- assortative mating, 253
- asthma, 112–114, 250
- astronomy, 65, 128, 206, 207, 217
- AUC, *see* area under the curve
- Average Man, 43, 121
- base jumping, 201, 202
- baseball players, 7, 14, 15, 18, 24, 117, 190, 220, 226, 227, 251
- Basquiat, Jean-Michel, 12–14, 61, 225, 238
- Bayes Rule, 2, 13, 14, 24, 35, 51–60, 62, 64–67, 171, 180, 218, 233–236, 239, 265, 266, 269
- Bayes, Thomas, 19, 50, 51, 53–56, 58, 235, 237
- Bayesian, 7, 11, 13, 14, 19, 20, 51–53, 56, 58–68, 76, 77, 96, 150, 181, 183, 187, 189, 191, 218, 220, 221, 235, 238, 239, 241, 251, 268–270  
inference, 19, 64, 66, 218  
*The Bayesian*, 64
- bell-shaped curve, 25, 29, 35, 39, 43–45, 47–49, 76, 84, 121, 126, 128, 132–134, 148, 154, 155, 157, 165, 168, 169, 178, 185, 245, 246, 260, 262
- bell-shaped mountain, 121, 130–134, 253, 255–257
- Bellhouse, David, 235
- Benjamini–Hochberg Rule, 186–189, 221, 267–270
- Bernoulli, Jacob, 27, 54, 57, 162, 228, 231, 234, 261  
*Ars Conjectandi*, 228
- beta carotene, 90, 94, 97, 247
- BH, *see* Benjamini–Hochberg Rule
- bias, 27, 59, 61, 62, 65, 78–81, 92, 139, 140, 142, 146, 153, 164, 173, 227, 251  
confirmation, 80, 82, 243, 244  
selection, 137, 176, 189, 190, 258
- biased coin, 63, 235  
design, 137, 258
- Big Data, 2, 5, 216
- big hat, 154, 274
- billiard table, 58, 237
- bivariate normal distribution, 132, 253, 255, 256
- black holes, 70
- blind audition, 81
- blinding, 21, 33, 70, 75, 80–82, 84, 90, 105, 219, 222, 243  
double-blinded, 81, 82, 96
- blue cars, 77, 78
- Bond, James, 23
- Bonferroni, Carlo, 186, 187, 189, 224, 267, 268
- bootstrap, 2, 33, 47, 138, 147, 148, 152–160, 165, 166, 168–170, 194, 213, 218, 220, 221, 240, 242, 260, 261, 263, 264

- bootstrap (*continued*)  
  mean, 155–157, 260  
  replication, 158, 159, 169, 260, 263, 264  
  sample, 154, 156, 158, 259, 261, 274
- Bru, Bernard, 28, 228
- Bru, Marie-France, 28, 228
- Buffon's needle, 263
- Buffon, Comte de, 131, 263
- butterfly effect, 5
- California, 33, 34, 142, 209  
  San Francisco, 33, 34, 209, 210, 273
- cancer  
  clusters, 100, 101, 248, 249  
  head and neck, 195–197, 202–205, 207, 271  
  leukemia, 99–101, 248, 249  
  liver, 93, 105, 106, 108, 111, 249  
  prostate, 184–189, 211–213, 219, 258, 266–270, 274
- car insurance, 2, 5–7, 11
- Cardano, Gerolamo, 162
- Cardiac Arrhythmia Suppression Trial, 95, 96, 174
- cars  
  blue, 77, 78  
  red, 77, 78
- case control study, 99, 105
- case fatality rate, 34
- CAST, *see* Cardiac Arrhythmia Suppression Trial
- Cato Institute, 176, 264
- Cauchy, Augustin-Louis, 54
- causality, 97, 98, 105, 112, 119, 134, 136, 138, 139, 144, 146, 244
- causation, 2, 68, 119, 136, 137, 146, 258
- celestial mechanics, 5
- censored data, 202, 203, 206, 207, 272
- Central Limit Theorem, 44–47, 148, 168, 218, 262
- Ceres, 128, 129
- ChatGPT, 215, 274
- chemotherapy, 72, 77, 195, 203, 240
- Chernobyl, 99, 248
- chi-squared test, 245
- cholesterol, 3, 82–84, 244, 254–256  
  LDL, 3–5, 7, 11, 184, 223, 224, 254–256
- cholesterol data, 83
- cholestyramine, 82–84, 244, 246
- Christmas, 38, 261
- Clemente, Roberto, 15–19, 227
- climbing Mt. Everest, 201
- coffee, 93, 96, 108–111, 113, 114, 142, 246, 249
- coin  
  biased, 63, 235  
  biased design, 137, 258  
  fair, 137, 191, 192, 227, 228, 230, 258
- comparison, 2, 16, 18, 32, 41, 68–72, 78, 79, 84, 87, 105–108, 110, 136, 143–145, 176, 178, 194–196, 204, 245, 258, 272  
  gold standard, 21, 60, 69
- compliance, 83, 84
- computer science, 2, 68, 225, 227
- conditional probability, 233, 234
- cone of uncertainty, 238
- confidence interval, 67, 88, 156, 168, 169, 183, 205, 218, 219, 241, 260, 263–265, 272
- confirmation bias, 80, 82, 243, 244
- constellations, 21, 221
- correlation, 2, 4, 68, 69, 104, 112, 119–124, 129, 131, 136, 137, 146, 195, 215, 216, 221, 252, 253, 255, 256, 258, 263  
  Pearson's coefficient, 252
- COVID-19, 14, 23, 24, 26, 28, 33, 34, 53, 59, 60, 94, 95, 104, 140, 143, 176, 198, 202, 229, 237, 247, 257, 264  
  deaths, 33, 34, 139, 140, 144, 177, 202, 237, 238, 258, 272
- dagger mark (†), xi, 1
- dart thrower, 36, 37, 230
- Darwin, Charles, 117, 124, 131, 253  
  *Origin of Species*, 120, 131
- data  
  Big, 2, 5, 216  
  censored, 202, 203, 206, 207, 272  
  cholesterol, 83  
  genomic, 127  
  kidney, 51, 125, 126, 130, 158–160, 165, 168, 209, 210, 217, 225  
  noisy, 24, 65, 148, 221  
  obituary, 150, 151, 153–159, 161, 168, 220, 260, 263  
  prostate cancer, 187–189, 212, 213, 267, 268  
  Social Security, 197, 198, 200, 208
- de Finetti, Bruno, 66, 230
- de Tocqueville, Alexis, 69
- “De Vetula”, 31, 228
- death rate, 59, 111, 139, 140, 176, 177, 198, 202, 205, 237, 258, 264, 272
- decision boundary, 179–181, 265, 266
- decision theory, 171, 172, 178, 189, 194
- deduction, 11, 12, 55
- Deep Learning, 11, 146, 168, 211, 213, 216, 221, 274
- denominator, 33, 34, 269

- diabetes, 94–96, 98, 99, 247–249, 254–256, 258
- Diaconis, Persi, 227
- dice  
     nontransitive, 41  
     paradox, 23, 38, 39, 41, 230  
     shaved, 32, 45, 46
- discoveries  
     False, 74, 186, 187, 189, 251, 268  
     True, 74, 187–189, 251, 267, 270
- distribution, 47, 48, 140, 235, 259, 262, 263, 265, 266  
     bivariate normal, 132, 253, 255, 256  
     multivariate normal, 256  
     normal, 49, 169, 184, 186, 226, 231, 242, 256, 266  
     Pareto, 206, 272, 273  
     posterior, 58, 65, 66, 221, 236, 239  
     prior, 56–59, 62–64, 66, 238, 241, 265, 268, 269  
     probability, 38, 39, 41, 49, 230, 231, 237, 241, 259, 261
- dose response, 110, 111
- double-blinded, 81, 82, 96
- drunkard's walk, 191–193
- Dutch book, 35
- $E = mc^2$ , 8, 53, 160, 216
- EB, *see* Epstein–Barr
- ecological fallacy, 140, 258
- educated guess, 59
- effect size, 111, 113, 117
- Einstein, Albert, 8, 70, 71, 125  
     theory of relativity, 24, 240
- empirical Bayes, 170, 191, 220, 224, 269, 270
- Epstein–Barr, 134–136, 138–140, 257
- error  
     margin of, 2, 147, 148  
     of the first kind, 75, 170, 172, 173, 175, 177–179, 186, 187, 193, 221, 242, 265, 267  
     of the second kind, 75, 170, 172, 173, 175, 177–179, 193, 221, 242, 265  
     type I, 177, 179, 189, 267  
     type II, 177, 179
- estimate, 17, 47, 67, 123, 141, 142, 148, 149, 151, 153, 156, 157, 161, 162, 183, 189, 190, 195, 204, 210, 218, 220, 269, 270, 272  
     empirical, 32, 270  
     James–Stein, 226, 227  
     least squares, 129, 130  
     lowess, 160, 169, 264  
     regression discontinuity, 142  
     Tweedie's, 189, 190, 270
- estimation, 66, 138, 147, 148, 218, 220–222, 241  
     Lynden–Bell, 207  
     maximum likelihood, 67, 88, 170, 182, 238
- eugenics, 124, 125
- experience, 6, 19, 56, 62, 64, 67, 77, 184, 206, 218, 233  
     learning from, 2, 6, 7, 11, 14, 18–20, 140, 164, 184, 188, 222, 223, 270  
     of others, 2, 6, 7, 14, 19, 140, 184, 188, 222, 223, 270
- $F = MA$ , 8
- fair coin, 137, 191, 192, 227, 228, 230, 258
- fallacy  
     ecological, 140, 258  
     single-cause, 139
- False discoveries, 74, 186, 187, 189, 251, 268
- false positive, 4, 224
- fatality rate  
     case, 34  
     population, 33, 34
- FBI, 12, 13
- FDA, *see* Food and Drug Administration
- Feminine Forever*, 92, 247
- Fermat, Pierre de, 27, 28, 30, 228
- Feynman, Richard, 243
- fiducial inference, 183
- fish oil, 90, 112, 114, 248, 250  
     omega-3, 94, 95, 248
- Fisher, Ronald A., 46, 65–67, 75, 76, 78, 79, 84–89, 125, 131, 172, 181–183, 186, 187, 194, 225, 235, 238, 243, 245, 253  
     information, 88
- Flynn effect, 48
- foggy curtain, 1, 2
- Food and Drug Administration, 64, 67, 69, 75–77, 82, 86, 87, 117, 171, 174, 175, 181, 221, 222, 241
- Ford, Henry, 59, 60, 237, 238
- framing, 34, 162
- fraternal, 50–53, 233, 234, 236
- frequentism, 63, 67, 218
- frequentist, 11, 20, 63, 66–68, 77, 183, 191, 218, 220, 241  
     inference, 19, 218
- Galilei, Galileo, 9, 92
- Galton, Francis, 119–125, 127–135, 146, 217, 252, 253, 255–257
- Galvani, Luigi, 80, 221
- Gauss, Carl Friedrich, 43, 54, 55, 128–130, 254, 262

280 INDEX

- genome-wide association study, 3–5, 7, 10, 11, 20, 46, 216, 232
- genomic, 3, 5, 127
  - data, 127
  - study, 3
- genotype, 223, 224
- GFR, *see* glomerular filtration rate
- giving birth in the US, 201
- glomerular filtration rate, 126, 158, 168, 254
- gold standard, 21, 60, 69, 75, 79–81, 90, 92, 135, 136
- Gosset, William Sealey, 245
- green tea, 258
- GWAS, *see* genome-wide association study
  
- HAL, 216
- Hall, Monty, 171, 172, 189
  - problem, 171, 172
- hazard rate, 197, 199–202, 204–207, 272
- HCQ, *see* hydroxychloroquine
- head and neck cancer, 195–197, 202–205, 207, 271
- Health Mountain, 22, 91, 101–104, 114
- Health Plateau, 22, 91, 101, 103–105, 114
- heredity, 120–122, 124, 125, 131, 133, 134
- hidden factor, 89, 97, 98, 136–138, 140–142, 144, 145, 257
- Hippocratic oath, 21
- histogram, 43, 45, 63, 74, 86, 116, 121, 133, 148, 151, 155, 157, 168, 169, 185, 186, 264, 268, 269
- HIV, 136, 206
- hokum, 59, 108
- Holmes, Sherlock, 219
- Holmes, Susan, 239
- Holy Grail, 137
- hormone replacement therapy, 92, 247
- HRT, *see* hormone replacement therapy
- Hume, David, 53, 54, 56, 57, 66, 134, 136, 235, 237
- Huygens, Christiaan, 1
- hydroxychloroquine, 14, 59–61, 237, 238
- hypothesis test, 21, 73, 76, 181, 183, 184, 219, 220, 222, 245, 251
  
- identical, 50–53, 64, 232–234, 236
- impossible zone, 180, 182
- independence, 218, 231, 253, 268, 270
- induction, 11, 12
- inference, 5, 11–14, 51, 64, 66, 138, 170–172, 183, 184, 191, 218–220, 235
  - Bayesian, 19, 64, 66, 218
  - fiducial, 183
  - frequentist, 19, 218
  - statistical, 2, 4, 7, 11, 12, 14, 20, 26, 28, 49, 51, 53, 55, 59, 62, 66, 88, 170, 182, 194, 218, 222
- inflection point, 44, 47, 48, 262
- information, 2, 6, 14, 17, 18, 20, 29, 52, 53, 62, 103, 114, 180, 184, 225
  - Fisher, 88
  - new, 52, 53
  - prior, 65, 66
  - science, xi, 2, 218
- instrumental variable, 142–145
- interferometer
  - LIGO, 70, 71, 240
  - Michelson, 70
- Ioannidis, John, 114, 115, 117, 240, 250–252
- IQ, 48, 49, 122, 172
  
- James–Stein, 17, 19, 117, 184, 190, 221, 225, 226
  - estimate, 226, 227
  - prediction, 15, 16, 19
  - Rule, 16, 220, 252, 270
- Jeffreys, Harold, 65, 66, 238
  
- Kaplan–Meier curve, 196, 206, 207, 221, 272
- kidney, 10, 11, 51, 92, 126, 158, 170, 210, 263, 264
  - data, 51, 125, 126, 130, 158–160, 165, 168, 209, 210, 217, 225
  - function, 10, 11, 126, 146, 158, 169, 210
- Krebiozen, 221
  
- Lagrange, Joseph-Louis, 54
- Laplace, Pierre-Simon, 8, 43, 54, 57, 65, 129, 236, 262
  - Rule of Succession, 57, 58, 65, 235–238
- Large Language Models, 215, 216, 274
- law
  - Large Numbers, 36, 148, 161, 162, 168, 218, 261
  - Newton’s, 5, 7–9, 117
  - perverse, 2, 195, 205–207, 272
  - Small Numbers, 164
  - square root, 165, 167, 168, 183, 184, 218, 262, 263
- Law of Large Numbers, 36, 148, 161, 162, 168, 218, 261
- Law of Small Numbers, 164
- LDL, 3–5, 7, 11, 184, 223, 224, 254–256
- learning from experience, 2, 6, 7, 11, 14, 18–20, 140, 164, 184, 188, 222, 223, 270
- least squares, 129, 130, 225, 254
  - algorithm, 130, 221, 254
  - curve-fitting, 129, 253

- fit, 83, 130
- line, 83, 129
- method, 125, 128
- Legendre, Adrien-Marie, 128–130, 254
- leukemia, 99–101, 248, 249
- LIGO, 70, 71, 240
- Literary Digest*, 137, 153, 259
- liver cancer, 93, 105, 106, 108, 111, 249
- LLM, *see* Large Language Model
- LLN, *see* Law of Large Numbers
- lowess, 126, 130, 158–160, 165, 166, 169, 254, 263, 264
- Lynden–Bell estimator, 207
  
- MammaPrint, 72, 73, 75–77, 81, 86, 240, 245
- Manhattan plot, 3, 4, 223
- margin of error, 2, 147, 148
- Markov chain Monte Carlo, 66, 218, 220, 221
- Martingale Theorem, 192, 193
- matched sample, 105, 136
- matching, 139, 140, 143, 146, 257
- math stat, 182
- maximum likelihood estimation, 67, 88, 170, 182, 238
- MCMC, *see* Markov chain Monte Carlo
- medical news, 2, 22, 89, 91, 177, 221
- Mendel, Gregor, 46, 130, 131
- Mendelian randomization, 144, 145
- meta-analysis, 94, 105, 223, 247
- Michelson and Morley, 240, 243
- Michelson interferometer, 70
- micromort, 201, 202, 204, 272
- midmean, 158, 159, 166
- MINDACT, 240
- Moderna, 23, 24, 26, 33, 229
- Monopoly, 80
- Montmort, Pierre Remond de, 28, 31
- Monty Hall problem, 171, 172
- Mormons, 142
- Morris, Carl, 14–16, 19, 34
- motorcycling, 201
- MS, *see* multiple sclerosis
- multiple sclerosis, 134–136, 138–140, 257
- multivariate normal distribution, 256
  
- natural experiment, 142–144
- needles-in-haystacks, 4
- neural net, 225, 274
- Newium, 63, 64, 67, 84–87, 115–117, 171, 178, 181, 193, 194, 241, 242, 245, 246, 259
- Newton, Isaac, 7–10, 128, 210, 262
  - law, 5, 7–9, 117
  
- Neyman, Jerzy, 67, 88, 172, 181–183, 194, 265
  - construction, 183, 265
  - Neyman–Pearson lemma, 67, 180, 182, 194, 221, 264, 271
- noise, 1, 9, 11, 25, 51, 72, 127, 128, 130, 137, 160, 164, 168, 189, 210, 214, 217, 218, 226
- noisy data, 24, 65, 148, 221
- nonconformist, 54–56, 68
- nontransitive dice, 41
- normal distribution
  - bivariate, 132, 253, 255, 256
  - multivariate, 256
- normal family, 49
- Null Hypothesis, 67, 72–77, 85–87, 178, 181, 184, 186, 194, 195, 220, 240–243, 245, 246, 250, 251, 266–272
- numerator, 33, 34
  - records, 34
  
- obesity, 93, 96–99, 105, 106, 111, 113, 114, 140, 248, 249
- obituary data, 150, 151, 153–159, 161, 168, 220, 260, 263
- objectivity, 64, 65, 67, 90, 239
- observational study, 135, 140, 257
- odds, 27, 35, 41, 51–53, 59, 60, 238
  - posterior, 35, 52, 53, 59, 60, 251
  - prior, 35, 52, 53, 59, 60
  - ratio, 61, 234
- old pony joke, 160, 261
- Old Testament, 43
- Oldium, 63, 64, 67, 84–87, 116, 117, 171, 178, 181, 193, 194, 241, 242, 245, 246, 259
- omega-3, 94, 95, 248
- On the Origin of Species*, 120, 131
- Orlando Museum of Art, 12, 61, 238
  
- p*-value, 72–75, 77, 84–87, 108, 116, 178, 184–187, 195, 218, 219, 221, 241–243, 245, 246, 250, 251, 267–269, 272
- pandemic, 176
  - COVID-19, 14, 23, 24, 26, 28, 33, 34, 53, 59, 60, 94, 95, 104, 139, 140, 143, 144, 176, 177, 198, 202, 229, 237, 238, 247, 257, 258, 264, 272
  - HIV, 206
- pangensis, 131
- paradox dice, 23, 38, 39, 41, 230
- parameter, 149–151, 160, 169, 220, 226, 238, 239
- Pareto, Vilfredo, 206, 272, 273
- Pascal, Blaise, 27, 28, 30, 228
- Pearson, Karl, 88, 181, 182, 245, 253
  - correlation coefficient, 252

282 INDEX

- peer review, 108
- permutation, 84, 86, 87, 154, 170, 221, 245, 246, 259
  - test, 67, 84–86, 138, 245, 246
- perverse law of waiting, 2, 195, 205–207, 272
- Phase 3 trial, 64
- philosophy, 2, 11, 38, 68, 105, 219
- physicist's twins, 19, 50–53, 64, 233
- placebo, 33, 59, 78, 80–84, 92, 174, 237, 244, 246–248, 250
- Plato, 36, 230
- plus-or-minus, 2, 148, 150–153, 155, 157, 159, 161, 163–166, 168, 169, 261–263
- Poincaré, Henri, 5, 44, 45, 91, 114, 117
- point cloud, 120, 122, 123, 126, 132, 146, 254, 256
- Pollock, Jackson, 225
- population fatality rate, 33, 34
- posterior odds, 35, 52, 53, 59, 60, 251
- potatoes, 76, 78, 89, 250
- power, 75, 177, 178, 181, 182, 241, 246, 267
- prediction, 2, 11, 14–16, 19, 34, 35, 62, 66, 126, 127, 130, 146, 157, 195, 207, 209–212, 214, 215, 221, 224, 225, 240, 252, 253, 255, 256, 274, 275
  - error, 16, 18
  - James–Stein, 15, 16, 19
  - pure, 211, 215, 216, 219, 221, 274
  - rule, 212–215, 224
  - tree, 213, 214
  - variable, 8, 125, 129, 254, 274
- Price, Richard, 54, 56–58, 66, 235, 237
- Principle of Insufficient Reason, 57, 58, 65, 66, 234
- prior odds, 35, 52, 53, 59, 60
- probability, 1, 2, 20, 23, 24, 26–31, 34–38, 49–51, 54–57, 66, 73–76, 100, 101, 126, 132, 145, 156, 161, 162, 167, 168, 171, 173, 183, 186, 192, 193, 199–201, 209, 218, 220, 224, 228–231, 235, 236, 238, 240, 246, 248, 258, 262, 263, 265–267, 269, 270, 273
  - conditional, 233, 234
  - distribution, 38, 39, 41, 49, 230, 231, 237, 241, 259, 261
  - model, 28, 76, 84, 132, 153, 218
  - personal, 35, 66, 229
  - posterior, 53, 57, 236, 269, 270
  - prior, 56, 57, 64, 66, 269
  - sequential ratio test, 271
  - subjective, 35, 66
  - theory, 25, 26, 28, 31, 32, 35, 55, 66, 218, 235, 236
- profiling, 62
- proportions, 14, 16, 27, 28, 31–36, 38, 48, 58, 80, 83, 86, 129, 130, 140, 148, 150, 161, 167, 169, 183, 186, 187, 196, 220, 229, 233, 246, 262, 268, 270
- prostate cancer, 184–189, 211–213, 219, 258, 266–270, 274
- Pseudo-Ovid, 31, 32, 229
- pseudorandom numbers, 230
- pure prediction algorithm, 211, 215, 216, 219, 221, 274
- quantum theory, 36, 91, 96
- Quetelet, Adolph, 43, 44, 48, 121, 257
- r* squared, 129, 130, 254, 255
- radiation, 72, 99, 100, 195–197, 203, 248, 272
- rainfall, 209–211
- random
  - digit dialing, 259
  - numbers, 37, 86, 189, 230, 246
  - pseudo, 230
  - quantities, 256, 270
  - sample, 146, 147, 152, 156, 157, 261
  - selection, 26, 27, 146, 152, 259
  - universe, 1–3, 5, 36, 127, 138, 139, 151, 160, 162, 217
  - walk, 191–194
- Random Forests, 213–215, 221, 274
- randomization, 21, 27, 59, 70, 75, 78, 81, 82, 84, 89, 97, 134, 137, 138, 143, 144, 146, 182, 219, 222, 230, 250, 274
  - Mendelian, 144, 145
- randomized clinical trial, 33, 59, 60, 69, 70, 75, 79–82, 84, 85, 87, 89, 90, 92, 94–98, 105, 107, 109, 113, 135–138, 144–146, 195, 219, 222, 229, 241, 242, 246–248
- randomness, 5, 24, 28, 36, 48, 77, 128, 149, 151, 160, 162, 195, 228, 230, 268
  - amplifier, 36, 230
- rate
  - case fatality, 34
  - death, 59, 111, 139, 140, 176, 177, 198, 202, 205, 237, 258, 264, 272
  - glomerular filtration, 126, 158, 168, 254
  - hazard, 197, 199–202, 204–207, 272
  - population fatality, 33, 34
- RCT, *see* randomized clinical trial
- receiving general anesthetic, 201
- red cars, 77, 78
- regression, 124, 125, 130, 170, 241, 254, 258, 272
  - curve, 126, 127, 130, 146
  - discontinuity, 140–143

- line, 120, 122, 123, 125, 129, 130, 134, 141, 142, 253–255, 258
- to the mean, 41, 112, 119, 121–123, 125, 128, 131, 218
- rejection level, 21, 77, 186, 187, 251
- religion, 27, 53
- replication crisis, 115, 117, 251
- response variable, 8, 125, 129
- risk
  - base jumping, 201, 202
  - climbing Mt. Everest, 201
  - giving birth in the US, 201
  - motorcycling, 201
  - receiving general anesthetic, 201
  - scuba diving, 201
  - skydiving, 201
  - smoking, 201
  - wine drinking, 109, 201
- Robbins' Rule, 6, 7, 220, 224, 270
- robust, 158
- Roman soldiers, 25, 26, 29, 35
- Rothamstead, 76, 78
- roulette, 37, 38, 73
- Royal Society, 55, 56
- rule
  - Bayes, 2, 13, 14, 24, 35, 51–60, 62, 64–67, 171, 180, 218, 233–236, 239, 265, 266, 269
  - Benjamini–Hochberg, 186–189, 221, 267–270
  - James–Stein, 16, 220, 252, 270
  - Laplace's Succession, 57, 58, 65, 235–238
  - Robbins', 6, 7, 220, 224, 270
- Rumsfeld, Donald, 140
- Ruth, Babe, 34
- sample
  - matched, 105, 136
  - random, 146, 147, 152, 156, 157, 261
  - size, 84, 106–108, 110, 111, 117, 146, 151, 153, 154, 160–164, 166, 168, 181, 237, 241, 250, 261, 265, 271, 274
  - survey sampling, 152, 237
- San Francisco, 33, 34, 209, 210, 273
- SAT scores, 140–144
- saunas, 92, 106–108, 111, 113, 114, 249
- Savage, Leonard, 66, 229
- scatterplot, 10, 125, 126
- Scientific American*, 167, 230
- scorecard, 22, 93, 96, 98, 105, 113, 114, 221
- scuba diving, 201
- selection bias, 137, 176, 189, 190, 258
- sequential analysis, 193
- sequential probability ratio test, 271
- Set
  - Test, 213–215
  - Training, 213–215
- Shafer, Glenn, 28, 228
- shaved dice, 32, 45, 46
- shrinkage, 189, 226, 252
- significance, 82, 87, 116, 186, 219, 222–224, 240, 241, 251
  - statistical, 21, 69, 70, 72, 73, 75, 86, 90, 108, 241, 242, 251, 272
  - test, 71–76, 78, 82, 84, 183, 184, 186, 193, 240–243, 259
- significant, 97, 138, 181, 186, 187, 219, 240, 241
- simulation, 37, 219, 242
- single nucleotide polymorphism, 3–5, 7, 20, 145, 146, 184, 223, 224, 232
- single-cause fallacy, 139
- size
  - and power, 177, 181, 182, 241
  - effect, 111, 113, 117
  - sample, 84, 106–108, 110, 111, 117, 146, 151, 153, 154, 160–164, 166, 168, 181, 237, 241, 250, 261, 265, 271, 274
- skydiving, 201
- smoking, 89, 98, 99, 103, 136, 142, 144–146, 201, 209, 246, 249, 250, 257
  - cessation, 246
- snake-eyes, 24, 28, 30
- SNP, *see* single nucleotide polymorphism
- Social Security, 207, 208
  - data, 197, 198, 200, 208
- sonogram, 50, 52, 233, 234
- square root
  - law, 165, 167, 168, 183, 184, 218, 262, 263
  - of  $n$ , 160, 162–164, 166, 168
- standard deviation, 47–49, 121, 124, 127, 128, 133, 134, 155, 178, 184, 238, 255, 256, 259–263, 268
- statistical
  - inference, 2, 4, 7, 11, 12, 14, 20, 26, 28, 49, 51, 53, 55, 59, 62, 66, 88, 170, 182, 194, 218, 222
  - significance, 21, 69, 70, 72, 73, 75, 86, 90, 108, 241, 242, 251, 272
  - skeptic, 20, 96, 97, 111, 135, 136, 138, 151, 221, 222, 259
- Steinhaus, Hugo, 231
- Stigler, Stephen, 19, 25, 235
- Student's  $t$ -test, 245
- study
  - case control, 99, 105

- study (*continued*)  
  genome-wide association, 3–5, 7, 10, 11, 20, 46, 216, 232  
  genomic, 3  
  observational, 135, 140, 257  
subjective probability, 35, 66  
sugar, 96–99, 103, 248  
sun, 1, 55–58, 167, 235, 236, 240  
survey, 78, 147, 152, 153, 207, 259  
  sampling, 152, 237  
survival, 2, 38, 60, 68, 144, 195, 196, 198–201, 203–207  
  analysis, 195, 200, 204, 206, 207, 221  
  curve, 104, 195–198, 200, 204, 206, 207  
  time, 196–198, 203  
syllogism, 55  
symmetry, 30–32
- Tang, Hua, 223
- test  
  A/B, 80  
  chi-squared, 245  
  hypothesis, 21, 73, 76, 181, 183, 184, 219, 220, 222, 245, 251  
  permutation, 67, 84–86, 138, 245, 246  
  sequential probability ratio, 271  
  significance, 71–76, 78, 82, 84, 183, 184, 186, 193, 240–243, 259  
  Student's *t*, 245
- Test Set, 213–215  
Texas bull's-eye, 101  
*The Shadow*, 221
- theorem  
  Central Limit, 44–47, 148, 168, 218, 262  
  Martingale, 192, 193
- TIAA, 104  
Training Set, 213–215  
transit of Venus, 1, 5, 210, 217, 223  
Transportation Security Administration, 62
- trial  
  Cardiac Arrhythmia Suppression, 95, 96, 174  
  Phase 3, 64  
  randomized clinical, 33, 59, 60, 69, 70, 75, 79–82, 84, 85, 87, 89, 90, 92, 94–98, 105, 107, 109, 113, 135–138, 144–146, 195, 219, 222, 229, 241, 242, 246–248  
  UK Recovery, 59, 60, 237  
True discoveries, 74, 187–189, 251, 267, 270  
truncation, 207  
TSA, *see* Transportation Security Administration  
Turner, William, 51  
Twain, Mark, 39, 164  
Tweedie's formula, 189, 190, 268, 269  
twins  
  fraternal, 50–53, 233, 234, 236  
  identical, 50–53, 64, 232–234, 236  
  physicist's, 19, 50–53, 64, 233  
type I error, 177, 179, 189, 267  
type II error, 177, 179
- UK Recovery Trial, 59, 60, 237  
uninformative priors, 65, 66, 238, 239
- variability, 5, 47, 129, 130, 133, 134, 150, 160, 164, 232, 254, 255  
variable  
  instrumental, 142–145  
  prediction, 8, 125, 129, 254, 274  
  response, 8, 125, 129
- Venn, John, 65, 234  
Venus game, 25, 28, 29, 31, 32, 44–46  
vitamin C, 90, 93, 94, 97, 247  
vitamin D, 257  
von Mises, Richard, 230
- waiting times, 202, 206, 273  
weak learners, 215  
WHI, *see* Women's Health Initiative  
wine drinking, 109, 201  
winner's curse, 189, 190, 269  
wisdom of crowds, 124  
Wollstonecraft, Mary, 56  
Women's Health Initiative, 92, 247  
Zika infection, 236