FOUNDATIONS

The foundations are intended to help you intuitively grasp and visualize the consequences of uncertainty and risk. If you were learning to ride a bicycle, for example, the foundations phase would end as soon as you no longer required training wheels.

COPYRIGHTED MATERIAL

PART 1

THE BIG PICTURE

In Part 1, I will provide an overview of the Flaw of Averages, how it rose to prominence, and how technology and new business practices have the potential to provide a cure. I will finish with some general thoughts on the use and benefit of analytical management models.

CHAPTER 1

The Flaw of Averages

Our culture encodes a strong bias either to neglect or ignore variation. We tend to focus instead on measures of central tendency, and as a result we make some terrible mistakes, often with considerable practical import.

—Stephen Jay Gould, naturalist, 1941–2002

The measure of central tendency that Gould refers to is typically the *average,* also known as the *expected value,* and the mistakes he warns of result from a common fallacy as fundamental as the belief that the earth is flat. It permeates planning activities in business, government, and the military. It helped mask the recent subprime mortgage fiasco until it became a world crisis, and it will plague those trying to clean up the mess. It is even enshrined within our accounting codes. I call it the Flaw of Averages.[1,2] It states, in effect, that:

Plans based on *average* assumptions are wrong on *average.*

An apocryphal example concerns the statistician who drowned while fording a river that was, on average, only three feet deep, as depicted in the sensitive portrayal by cartoonist Jeff Danziger.

In everyday life, the Flaw of Averages ensures that plans based on *average* customer demand, *average* completion time, *average* interest rate, and other uncertainties are below projection, behind schedule, and beyond budget.

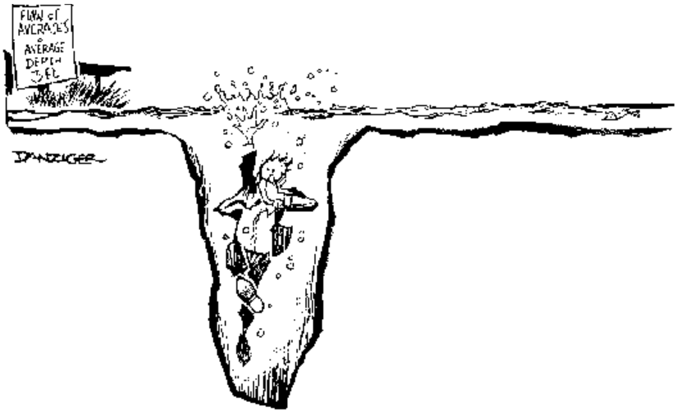

So people have been confused in the face of uncertainty for 2,000 years. What else is new? Plenty! What's new are several dramatic advances in software, data structures, and managerial outlook. Together, they form the bases of Probability Management, which brings a new transparency to the communication of risk and uncertainty. It is changing our perception of these concepts as profoundly as the light bulb changed our perception of darkness.

Give Me a Number

To understand how pervasive the Flaw of Averages is, consider the hypothetical case of a marketing manager who has just been asked by his boss to forecast demand for a new-generation microchip.

> *"That's difficult for a new product," responds the manager, "but I'm confident that annual demand will be between 50,000 and 150,000 units."*
>
> *"Give me a* number *to take to my production people," barks the boss. "I can't tell them to build a production line with a capacity between 50,000 and 150,000 units!"*

The phrase "Give me a number" is a dependable leading indicator of an encounter with the Flaw of Averages, but the marketing manager dutifully replies: "If you need a single number, I suggest you use the average of 100,000."

The boss plugs the average demand, along with the cost of a 100,000-unit capacity production line, into a spreadsheet model of the business. The bottom line is a healthy $10 million, which he reports as the projected profit. Assuming that demand is the only uncertainty and that 100,000 is the correct average (or expected) demand, then $10 million must be the average (or expected) profit. Right?

Wrong! The Flaw of Averages ensures that on *average*, profit will be less than the profit associated with the *average* demand. Why? If the *actual* demand is only 90,000, the boss won't make the projection of $10 million. If demand is 80,000, the results will be even worse. That's the downside. On the other hand, what if demand is 110,000 or 120,000? Then you exceed your capacity and can still sell only 100,000 units. So profit is capped at $10 million. There is no upside to balance the downside, as shown in Figure 1.1, which helps explain why, on average, everything is below projection.

But why are things behind schedule on average? Remember the chance of making it to the VIP reception on time, as described in the Introduction? When this occurs on an industrial scale, it can

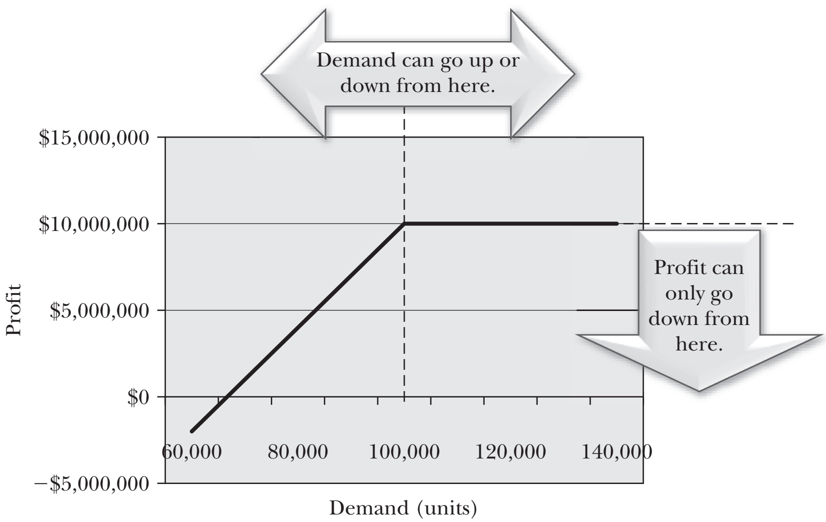

Figure 1.1 Average profit is less than the profit associated with average demand.

be much worse. Consider an idealized software project that will require ten separate subroutines to be developed in parallel. The boss asks the programming manager of the first subroutine how long development will take.

> *"I'm confident it will take somewhere between three and nine months" replies the programming manager.*
>
> *"Give me a number," says the boss. "I have to tell the chief operating officer when we'll be operational!"*
>
> *"Well," says the programming manager, "on average, programs like this take about six months. Use that if you need a single number."*

For simplicity of argument, assume that the boss has similar conversations with each of the nine remaining programming managers. The durations of all the subroutines are uncertain and independent, and they are expected to range between three and nine months with an average of six months. Because the ten subroutines are being developed in parallel, the boss now goes to the COO and happily reports that the software is expected to be operational in six months.

Assuming the durations of the ten subroutines are the only uncertainties and that each one has an average of six months, then the average or expected duration of the entire software project should be six months. Right?

Wrong! If you read the Introduction you know why. All ten projects coming in under six months is analogous to having your friends all show up on time. But now you must flip ten heads in a row instead of seven, and the odds are less than one in a thousand. Figure 1.2 displays a possible outcome in which many tasks take less than six months, yet the project takes 10.4 months.

And why is everything over budget on average?

Consider a pharmaceutical firm that distributes a perishable antibiotic. Although demand fluctuates, the long-term average is a steady five cartons of the drug per month. A new VP of operations has taken over the distribution center. He asks the product manager for a forecast of next month's demand. "Demand varies," responds the product manager, "but I can give you an accurate distribution, that is, the probabilities that demand will be 0, 1, 2, and so on." The product manager, who was apprehensive about his new

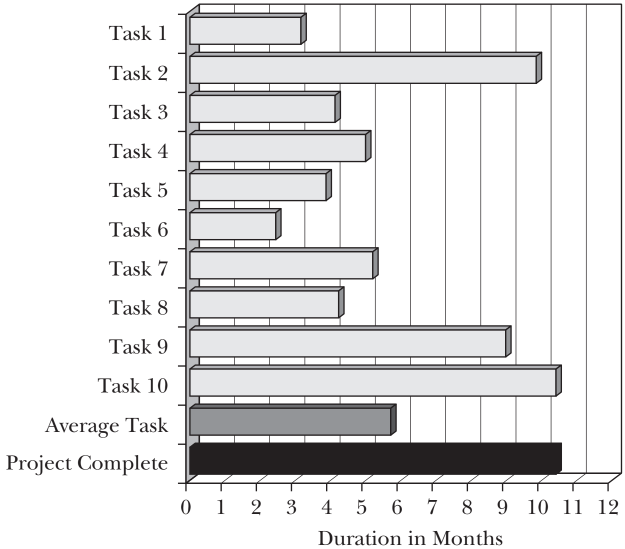

Figure 1.2 Many tasks come in under six months, but the longest is 10.4 months.

boss, is relieved that he could provide such complete information in his first professional interaction.

"If I had wanted a distribution, I would have asked for a distribution," snaps the boss, "give me a *number* so that I can calculate our operating costs." Eventually they settle on the time-honored tradition of representing the uncertainty by its average.

Armed with the accurate average demand of five cartons per month, the boss now proceeds to estimate inventory operating costs, calculated as follows:

- If monthly demand is less than the amount stocked, the firm incurs a spoilage cost of $50 per unsold carton of the perishable drug.
- On the other hand, if demand is greater than the amount stocked, the firm must air freight the extra cartons at an increased cost of $150 each.

A quick calculation indicates that if five cartons are stocked, and the demand happens to come in right at its average of five, then there will be neither spoilage nor air freight costs. Thus, the boss reasons, the average inventory operating cost will be zero. Right?

Wrong! If demand is below average, the firm gets whupped upside the head with spoilage costs, whereas if demand is above average, the firm gets whupped up the other side of the head with air freight costs. No negative costs exist to cancel out the positive ones; so, on *average*, the cost will exceed the cost associated with the *average* demand.

Later in the book I will distinguish between the strong form and weak form of the Flaw of Averages, as well as many subcategories. In the remainder of this chapter, I will present a few actual occurrences in various walks of life.

Statisticians to the Rescue?

Where do all those averages come from that people erroneously plug into their business plans? You guessed it. They come from statisticians and other analysts, whose so-called sophisticated models often perpetuate the Flaw of Averages.

Consider the graph in Figure 1.3 depicting economic growth data. By running regression analysis on the historical data represented by the solid black line, a statistician could estimate the average growth in the future, a single number represented by the slope of

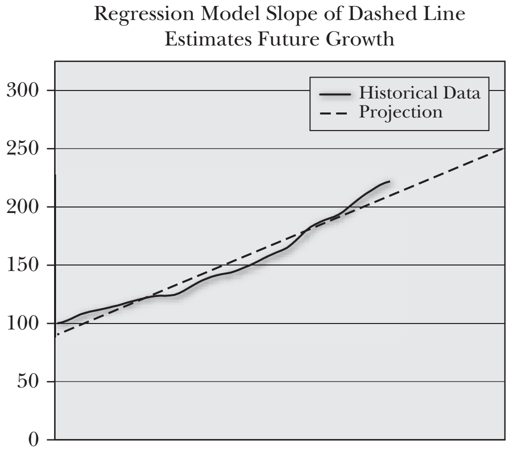

Figure 1.3 A statistical regression model provides an estimate of future growth as a single number, the slope of the dashed line.

the dashed line. Because this was derived from a "mathematical model" developed by an "expert," people are likely to believe it.

This nearly irresistible tendency to fixate on a single number has been well documented by Patrick Leach in a book entitled *Why Can't You Just Give Me the Number?*[3] He points out that "once a value is generated, put down on paper, and incorporated into the business plan, it becomes gospel."

In actuality, the data in Figure 1.3 represents the national growth in housing values from January 2000 to December 2005, taken from the S&P/Case-Shiller Home Price Index.[4] And numerous models analogous to the dashed regression line were used as justification for devastatingly bad investments.

What really happened to housing values is shown in Figure 1.4. Astonishingly, this possibility was not even considered by some of the risk models monitoring the economy. In the December 2008 issue of *Portfolio.com*, Michael Lewis chronicles the disaster in an article entitled "The End." According to Lewis, when someone inquired of Standard & Poor's what falling housing prices would do to default rates, they told him their model for price growth couldn't even accept negative numbers. This is like a model of coin flips that generates only heads!

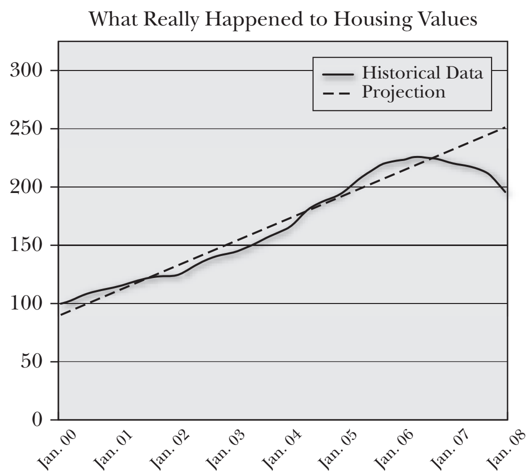

Figure 1.4 Actual data from January 2000 to January 2008.

Red Lobster

Summer 2003: Red Lobster seafood restaurants promote "Endless Crab: a celebration of all the hot, steaming snow crab legs you can eat." Shortly thereafter, the president of Red Lobster was replaced. According to the *St. Petersburg Times*,[5] "The move came after management vastly underestimated how many Alaskan crab legs customers would consume." Furthermore, "The chain was pinched by rising wholesale prices."

I suspect that during the planning of the ill-fated promotion, a high-level manager asked for the average number of customers expected to order crab. Further, the manager might have inquired about the average number of helpings per customer and the estimated price of crab. It would have been tempting to calculate the expected profit of the promotion based on these three numbers, but this approach would have been deeply flawed.

If the number of helpings exceeded expectations, then the chain was poised to lose money on each crab-eating customer. According to the *Times*, "'It wasn't the second helping, it was the third one that hurt,' company chairman Joe R. Lee said in a conference call with analysts." Worse, the uncertainties were linked: If demand exceeded expectations, the promotion itself had the potential to drive up the price of crab.

Thus estimated profit associated with the *average* demand, the *average* number of helpings, and the *average* price was higher than the *average* profit.

Red River

Spring 1997: The U.S. National Weather Service issues a forecast that the Red River is expected to crest at roughly 50 feet. *The New York Times* later quoted experts who said the problem was "that more precision was assigned to the forecast than was warranted."[6] The City of Grand Forks' communications officer, Tom Mulhern, said "[the National Weather Service] came down with this number and people fixated on it." According to *The Times*, "Actually, there was a wider range of probabilities," but the single number "forecast had lulled the town into a false sense of security." The article continued, "It was, they say, a case of what Alfred North Whitehead, the mathematician and philosopher, once termed 'misplaced concreteness.' And whether the problem is climate change, earthquakes, droughts or floods, they say the tendency to overlook uncertainties, margins of error and ranges of probability can lead to damaging misjudgments."

This was a classic case of the Flaw of Averages. Consider a hypothetical version of the Red River situation. Assume that, at the time of the forecast, the expected crest level is indeed 50 feet, but the actual level is still uncertain. In this version, Mother Nature determines the weather by flipping a coin. Heads creates torrential rains, which result in a 55-foot crest. Tails creates a mere drizzle, leading to a 45-foot crest. Because the dikes are designed to withstand a 50-foot crest, there is no damage when a tail occurs. But don't forget the 50 percent chance of a head, in which case flooding results in $2 billion in damage.

In short, the damage resulting from the average crest of 50 feet (the average of 45 and 55) is zero, whereas the average damage (the average of zero and $2 billion) is $1 billion.

In fact, what occurred in Grand Forks was a disastrous flood, forcing an estimated 50,000 people from their homes. *The New York Times* reported that "[i]t is difficult to know what might have happened had the uncertainty of the forecast been better communicated. But it is possible, said Mr. Mulhern, that the dikes might have been sufficiently enlarged and people might have taken more steps to preserve their possessions. As it was, he said, 'Some people didn't leave till the water was coming down the street.'" Figure 1.5 shows the difference between a flood slightly below and slightly above the average crest.

Now in case you are questioning the value of planning for above-average natural disasters, consider this scenario. If a Richter 7 earthquake hit a modern city with seismic building codes, it might kill a few hundred people, whereas similar quakes in a less developed part of the world regularly kill tens of thousands.

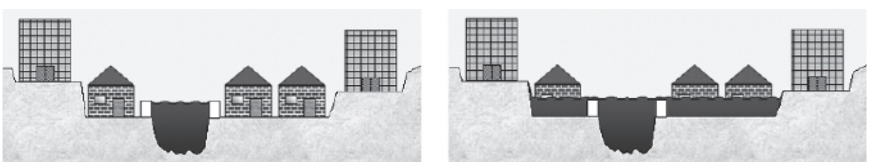

(a) (b)

Figure 1.5 (a) Flood slightly below average (no damage) versus (b) flood slightly above average (disaster).

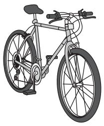

Visit FlawOfAverages.com for animations and simulations of several of the examples in this chapter.

Red Ink in Orange County

Summer 1994: Interest rates are low and are expected to remain so or fall even farther. Orange County, California, has created a financial portfolio to fund the pensions of its teachers and firefighters, based on the expected future behavior of interest rates. For several years, this fund, run by County Treasurer Robert Citron, has yielded much higher returns than comparable funds in similar municipalities. For the sophisticated investor, this is actually a red flag; there is no free lunch, as they say. John Moorlach, who unsuccessfully runs against Citron in 1994, argues in his campaign that "Mr. Citron believes he can accurately anticipate the market all the time, and also outperform everyone. That's impossible."[7] Nonetheless, so many people line up with their money that the county has to turn investors away. In fact, the fund has naively leveraged itself into a very risky position and goes bankrupt in December of 1994.

In 1995, Professor Philippe Jorion of the University of California at Irvine showed that if the county officials had explicitly considered the well documented range of interest rate uncertainties instead of a single *average* interest rate scenario, they would easily have detected the likelihood of the looming train wreck.[8]

There was absolutely no need for such a pension fund to shoot for the moon. Moreover, had the county's government members understood the increased risk they faced as a result, they would no doubt have adopted a more conservative investment strategy in time to prevent the debacle.

The Red Coats

Spring 1775: The colonists are concerned about British plans to raid Lexington and Concord, Massachusetts. Patriots in Boston (my friends in the United Kingdom use a less flattering name) develop a plan that explicitly takes a range of uncertainties into account: The British will come either by land or by sea. These unsung pioneers of modern decision analysis did it just right by explicitly planning for both contingencies. Had Paul Revere and the Minutemen planned for the single average scenario of the British walking up the beach with one foot on the land and one in the sea, the citizens of North America might speak with different accents today.

Why Forecasts Are Always Wrong: A Problem of Dilbertian Proportion

When managers ask for a forecast, they are *really* asking for a number, which involves the Flaw of Averages. For example, the product manager of the new microchip provided the *correct* forecast of average demand to his boss. But the boss turned around and used that single number to *incorrectly* forecast average profit. Each of the ten programming managers gave their boss the *correct* average completion time of six months for their subroutines. But the boss used those single numbers to *incorrectly* forecast the completion of the entire software project. The product manager at the pharmaceutical firm gave the VP of operations the *correct* forecast for average demand. But the VP used that number to *incorrectly* forecast operating costs.

In these cases, the bosses will ultimately claim that they got bad forecasts from their subordinates and will end up punishing them for providing what was in fact the *correct* average. This is indeed a problem of Dilbertian proportion.

In his book, *A Whole New Mind*,[9] author Daniel H. Pink predicts the ascendance in the economy of right-brained, big-picture thinking relative to left-brained analysis. The subjects of statistics and probability have traditionally been the domain of the left. But in fact, the Flaw of Averages often arises due to the left brain's stubborn insistence on a single precise answer. If anything, the right side of brain is better equipped to interpret the patterns inherent in uncertainty.

In subsequent chapters I will discuss new technologies that are making these patterns visible to the naked eye, allowing the right brain to get back in the game.