CHAPTER 1

What This Book Is and Why You Should Read It

Life is full of choices. At a job interview, you can give short, pleasant answers to questions. Or you can burst into an impassioned rant about how you will add value to the enterprise. You can dress sedately and behave discretely at a party, or go for maximum drama in your clothes and demeanor. In a basketball game you can throw up a quick shot, or pass the ball so the team can work into position for a higher-percentage shot. You can walk on by an interesting-looking stranger, or throw out a remark or a wink. These choices all concern risk.

In the basketball example, you have a coach. When the team is ahead late in the game, the coach will give one kind of advice. On offense, take plenty of time and get a high-percentage shot. On defense, deny the opponents easy shots and do not foul. Why? Because this style of play minimizes the variance of outcome, which is to the advantage of the team in the lead. The trailing team will try to shoot three-point shots quickly and will play aggressively for steals and blocks on defense. They don't mind fouls because those can change the score without running time off the clock. They are trying to maximize variance of outcome.

If you're not familiar with basketball, the same idea applies in virtually every competitive sport. The player or team that is ahead wants to minimize risk, whereas the opposing player or team wants to maximize it. In baseball, a pitcher with a lead throws strikes; when his team is trailing he will work the corners and throw off-speed pitches.

In soccer with a lead you try to control the ball and keep your defense back; when behind you attack aggressively. In hockey, the trailing team will sometimes even pull the goalkeeper. In American football, the team with the lead will run the ball up the middle and play prevent defenses, while the other team blitzes and throws long passes.

In the job interview, the short, safe answers are indicated if you think you're likely to get the job and just don't want to blow it. But if you're a long shot to be hired, maybe it's time to dust off that rant. Going to an obligatory party for your job, one you know will be boring? Navy suit, say as little as possible and only about the weather, don't drink, and leave early. But if you want to be the life of the party, have a great time, and maybe change your life? Think hot pink. And before you wink at the stranger, ask yourself if you're a bit bored and looking for new adventures—or is your life exciting and complicated enough already and you need peace and quiet more than a new friend?

Risk is something you dial up or down in order to accomplish a goal. It is neither good nor bad in itself. This is the sense in which I always use the word *risk* in this book. Compare this to the "risk" of a basketball player getting injured. I will use the word *danger* for this, not *risk*. Dangers should be minimized, subject to constraints. For example, we don't want to require so much protective padding that a game is not fun, or the cost is too great. So we don't try to set danger of injury to zero, but we also don't "manage" it; we never increase it for its own sake.

The counterpart to a danger on the good side is an "opportunity," such as the opportunity for a pitcher in baseball to get a no-hitter. This is considered so valuable that a manager will almost always leave a pitcher with a chance at a no-hitter in the game, even if he is tiring and a relief pitcher would increase the probability of winning the game.

Risk, Danger, and Opportunity

There are three tests to determine if something is a risk rather than a danger or an opportunity:

1. Risks are two-sided; you can win or you can lose. Dangers and opportunities are one-sided. If you have a sudden change of

health while playing football, it is highly unlikely to be an improvement.

2. Dangers and opportunities are often not measurable, and if they are, they are measured in different units than we use for everyday decisions. We can't say how many points a broken collarbone is worth, or whether two sprained ankles are better or worse than a broken finger. There is no dollar figure to put on the glory of setting a record or winning a championship. Risks, however, are measurable. In order to manage an uncertainty, we need some way of assigning relative values to gains and losses.
3. Dangers and opportunities often come from nature, and we usually have only limited ability to control them. Risks always refer to human interactions, and their level must be under our control—if not, they may be risks to somebody else but they are facts of life to us.

The distinction is not inherent in the uncertainties themselves; it is our choice how to treat them. For example, NASCAR has been accused of manipulating its rules to get an optimal number of fatal crashes per year: enough to keep a dangerous, outlaw edge but not so many as to kill all the popular drivers or provoke safety legislation. I have no opinion on whether this charge is true or false. If true, it means NASCAR is treating as a risk something that most people consider a danger. That might be immoral, but it is not illogical or irrational.

Some job applicants treat every question as a danger, carefully probing for traps and giving minimal answers to avoid the chance of mistake. They seldom get hired. Others treat every question as an opportunity to posture or boast. They never get hired. Some people go to parties that should be fun, and dress and act more appropriately for a funeral, letting the danger of embarrassing themselves overwhelm rational consideration of risk. Other people treat funerals as parties, grasping for opportunities that do not exist.

Another example of mixing up risk and danger is a famous memorandum by Ford Motor Company concluding that the cost to the company of settling lawsuits for Pinto owners burned to death in low-speed rear collisions was less than the $10 per car it would cost to shield the gas tank. This story, although widely believed, is a distortion of the facts, and Ford is innocent of any such decision. I mention it only to emphasize that the distinction between risks and dangers is in the eye of the beholder.

There are also things we can choose to treat as risks or opportunities. In *On the Waterfront,* protagonist Terry Malloy makes the famous lament, "I coulda had class. I coulda been a contender. I coulda been somebody, instead of a bum, which is what I am," blaming his brother for persuading him to purposely lose a boxing match for the sure thing "short-end money." He is not complaining that there was not enough short-end money, but that he sold something that was literally priceless. His brother treated his opportunity like a risk, and managed it.

A coward treats risks as dangers, whereas a thrill seeker treats them as opportunities. We call them thin-blooded and hot-blooded, respectively. A cold-blooded person treats both dangers and opportunities as risks. *Red-blooded* refers to people who are excited by challenges, but not to the point of being blinded to dangers and opportunities. To keep this straight, think of the classic movie plot in which the red-blooded hero and his hot-blooded sidekick push aside the thin-blooded person in charge, to fight the cold-blooded villain. We admire the first two people in different ways, feel sorry for the third, and hate the fourth.

Red-Blooded Risk Management

In emotional terms, thin-blooded people are motivated mainly by fear, hot-blooded people by anger and other passions—or even merely thrills—and cold-blooded people by greed. Red-blooded people feel anger and fear and greed like anyone else, but understand successful risk taking is a matter of calculation, not instinct.

This is not a self-help book. I do not have any advice for how to achieve this psychological state, if that is what you want to do. What I can tell you is how to compute the red-blooded action in risk situations. It's mathematics, not psychology. Red-blooded risk management consists of three specific mathematical techniques, which have been thoroughly tested in real-world applications. Although quantitative skills are required to implement them, the ideas are simple and will be explained in this book without math. The techniques are used to:

- Turn any situation into a system with clearly delineated risks, dangers, and opportunities.
- Optimize the risks for the best possible outcome.
- Arrange things so both dangers and opportunities make the maximum positive contributions.

This field was invented by a cohort of quantitatively trained risk takers born in the 1950s. In the 1970s, we rebelled against conventional academic and institutional ideas of risk. We sought wisdom from actual risk-takers, which took us to some disreputable places. In the 1980s, we found ourselves taking risks on Wall Street, and developed the ideas described in this book between 1987 and 1992, although of course most of the ideas can be traced to much earlier work. University of Chicago economics professor Frank Knight, for example, made a distinction between risk, with known probabilities and outcomes, and uncertainty, which is something akin to our dangers and opportunities. But he did this to emphasize the limits of mathematics in decision making under uncertainty. He did not appreciate the power of quantitative methods for separating risk from uncertainty, nor the tremendous benefit from applying mathematics to optimize risk taking. Most important, he failed to see that mathematics can be brought to bear just as fruitfully on nonquantifiable uncertainty as on risk. Knight was a deeper thinker than any of the Wall Street risk takers, but we had far more experience in making successful quantitative risk choices.

This group of risk-taking rebels became known as "rocket scientists." That was partly because several of us actually worked on rockets (I myself spent a summer on satellite positioning, which technically uses rockets, but not the big ones that lift payloads into space; anyway, my contribution was entirely mathematical. I never saw an actual rocket firing except on film, so the experience certainly doesn't make me a real rocket scientist.), but mostly to capture the combination of intense and rigorous mathematical analysis tied firmly to physical reality, exploration, and adventure. Recall that one of our generation's defining moments was the Apollo moon landing. We weren't astrophysicists and we weren't engineers. We didn't know exactly what we were, but we knew it was something in between. A more general term for people who use quantitative methods in finance is "quant," but that term also describes less rebellious researchers with quantitative training who came to Wall Street later and called themselves "financial engineers."

I am aware that "rocket scientist" is a stupid name, both boastful and inaccurate. I didn't make it up, and don't use it much. I describe myself as a "quant" with a lowercase q, unpretentious as in, "just a simple quant." I'm not humble, as you'll figure out if you keep reading, but I'm not given to overstatement. What I do isn't rocket science, most of it is trivially simple and the rest is more meticulous care than brilliance. But to be historically accurate, we're stuck with the term, and it does convey some of the spirit of the group.

We contrasted ourselves to people we called "Einsteins," an even stupider name. We had nothing against Albert Einstein, but we disagreed with people who thought risk was deeply complex and could be figured out by pure brainpower, without actually taking any risk or observing any risk takers. "Einstein" was rarely used as a noun. It was more common as an adjective. "He had a good insight, but went Einstein with it," or "He used to be a rocket scientist but got offered a tenure track position and went Einstein." Don't blame me. I don't defend the usages, I just report them.

The rocket scientists rebuilt the financial system from the ground up. I compare these changes to the differences between a modern digital camera and a point-and-shoot film camera from 1980. They look similar. They both have lenses and flashes and shutter buttons. They both run on batteries, in some cases the same batteries. They are used to take pictures of vacations and parties and family members. They cost about the same. From the standpoint of sellers and users, the difference seems to be just an improvement in technology for the same basic device.

But for someone making cameras, there is no similarity at all. The modern technology is built on entirely different principles from the old one. From 1982 to 1992 rocket scientists hollowed out the inside of Wall Street and rebuilt it. We didn't set out to do that; it just happened. Most people, including most people working on Wall Street, didn't notice the fundamental change. They saw some of the minor external design changes, and noticed one day there was no more film to develop, but missed that something unprecedented in history had been created.

At the same time, with even less intention, we figured out the 350-year-old riddle at the heart of probability theory. As has always been the case with probability, practitioners ran ahead of theory. No doubt we will someday have a coherent theoretical explanation of how modern financial risk management works. Until then, all I can do is show you how and why it came into being, and what it is doing to the world.

Risk and Life

Risk taking is not just a quantitative discipline, it is a philosophy of life. There are basically two sensible attitudes about risk. The first is to avoid it whenever possible, unless there is some potential payoff worth the risk. The second is to embrace risk taking opportunities that appear to offer a positive edge. The advantage of the second course is that you take enough gambles that the outcome of any one, or any ten or hundred, doesn't matter. In the long run, you will end up near your expected outcome, like someone flipping a coin a million times.

In my experience, people incline to one of these two strategies early in life. Perhaps it's in our genes. In this context, I always think of a highway sign you can see if you drive from Nice to Monte Carlo. There is a fork, and the sign points right to "Nice Gene" and left to "Monte Carlo Gene." On that choice, I'm a leftist. That doesn't mean I take huge risks; it means I take lots of risks. I have learned from others and invented myself ways to balance these to ensure a good outcome, insomuch as mathematics and human efforts can ensure anything.

There are three iron rules for risk takers. Since your plan is to arrive at an outcome near expectation, you must be sure that expectation is positive. In other words, you must have an edge in all your bets. Expectation is only an abstraction for risk-avoiders. If you buy a single $1 lottery ticket, it makes no practical difference whether your expected payout is $0.90 or $1.10. You'll either hit a prize or you won't. But if you buy a million tickets, it makes all the difference in the world.

Second, you need to be sure you're not making the same bets over and over. Your bets must be as independent as possible. That means you cannot rely on systems or superstitions, not even on logic and rationality. These things will lead you to make correlated bets. You must search hard for new things to bet on, unrelated to prior bets, and you must avoid any habits. In many cases you find it advantageous to make random decisions, to flip coins. For risk avoiders taking only a few big chances, correlation is a secondary concern and flipping a coin for a decision makes no sense.

Finally, risk takers must size their bets properly. You can never lose so much that you're taken out of the game; but you have to be willing to bet very big when the right gambles come along. For a risk avoider, being taken out of the game is no tragedy, as risk taking was never a major part of the life plan anyway. And there's no need to bet larger than necessary, as you are pursuing plans that should work out if nothing bad happens, you're not counting on risky payoffs to succeed.

While moderation is often a good strategy, I don't think you can choose a middle way between risk avoiding and risk taking. Consider an investment portfolio. You can invest in high-quality bonds with payoffs selected near the times you expect to need the money, and possibly hedge your bets further by buying hard assets. Or you can buy stocks and hope for the best. If you choose the latter route, the risk-taking approach, you should seek out as many sources of investment risk as you think the market compensates—that is, all the securities for which there is a positive edge. Both strategies make sense, but it's crazy to split the difference by buying only one stock. You either avoid risk as much as practical, or you try to find as many risks as you can.

You could, of course, put half your portfolio in bonds and the other half in diversified risky assets, but this still makes you a risk taker, seeking out as many risks as possible. You just run a low risk version of the strategy. There's nothing that says a risk taker has to have a high-risk life. In practice, however, once investors take all the trouble to create a broadly diversified portfolio, or individuals learn to embrace risk, they tend to exploit the investment.

It's good that people make this choice young, because each route requires skills and life attitudes that would be fatal to acquire playing for adult stakes. Risk takers must enjoy the volatility of the ride, because that's all there is. There is no destination. You never stop gambling. Risk avoiders must learn to endure volatility in order to get to the planned destination. The world needs both kinds of people.

If you are a risk taker, you need the material in this book to survive, assuming you haven't already figured it out for yourself. We know a lot about the mathematics of risk taking that no one in the world knew a quarter century ago. If you are not a risk taker, you should still understand the mathematics of risk due to its effect on the world.

Quantitative risk models from Wall Street are in considerable disrepute at the moment. I hope to convince you that attitude is wrong. Whether or not I do, I can tell you that these models have changed the world completely, and the pace of that change will only accelerate. So even if you think they are worthless or harmful, it's worth understanding them.

Play and Money

I'm going to cover some topics you might not expect in a book on risk. First is play. One of the characteristics of play is that it takes place within a delineated area—physical or mental—which is not allowed to interact with the rest of the world. Basketball, for example, takes place on a court with clearly defined physical boundaries—and has people to blow whistles if the ball goes beyond those boundaries, stopping play until the situation is rectified. You are not allowed to buy a basket for money or any other consideration outside the perimeter of the game. Whether two players like or dislike each other is supposed to be irrelevant; their actions depend only on whether they're on the same team or on opposing teams. This is what allows us to treat the in-game events as risks. When the outside world intrudes, as with an injury or an equipment failure, those events cannot be managed as risks, because by rule they are incommensurate with baskets.

Although the world is not supposed to intrude on play, play can have enormous effect on the world. Elections, trials, and some wars are contests governed by rules that occur in designated times and places. Market competition can be considered a game, and game theory is a major part of the study of economics. Less serious games constitute a large portion of the economy: sports, gambling, video games, hobbies, and many other activities represent sizable aggregate demand for products and services. We will look deeply into these matters because risk management depends on the kind of delineation and isolation required by play. In a deep sense, risk is play and play is risk.

We're also going to discuss money. When economists consider risk, they usually assume that the types of stakes don't matter—gambling for money is no different from gambling for anything else of value. That turns out not to be true. Optimizing requires goals and constraints. Optimizing risk requires that the two be interchangeable. One way that can happen is if both are measured in money. It also turns out that any time you set up a risk-taking activity with the same units used for goals and constraints, you create a form of money.

One of the major schools of mathematical probability makes betting the fundamental definition of probability. It is called Bayesian theory. Bruno de Finetti's famous example concerns the probability that life existed on Mars one billion years ago. It seems difficult to put a number on that, or even to know what a number would mean. But suppose there is an expedition that will determine the answer tomorrow. There is a security that pays one dollar tomorrow if life existed on Mars one billion years ago, and nothing otherwise. There is some price at which you will buy or sell this security. According to de Finetti, that price *is* the probability that life existed on Mars a billion years ago. It's subjective to you; someone else could have a completely different price. But there is always a definable probability for any event, because you can always be forced to name a price at which you would buy or sell. Saying you don't know the probability of something is saying you don't know what you think.

Rocket scientists were the first group to see the implications of that formulation and ask some obvious questions. We noticed that the bet involved money and asked, "What currency are you betting with?" For example, suppose you would buy or sell the security that pays $10 for 10 cents, suggesting that the probability that life existed on Mars one billion years ago is 1 percent. But this expedition to Mars financed itself by selling bonds denominated in Mars Expeditionary Currency, or mecs. Mecs are the currency colonists will use. Each mec sells for $1 today. But if the expedition discovers there was life on Mars one billion years ago, the value of each mec will soar to $10 because of the potential value of artifacts and scientific discoveries, and because it makes it more likely that Mars can be made hospitable to life today. If you would pay 10 cents for a security that pays $10 if there was life on Mars, you would pay 10 centimecs for a security that pays one mec in the same circumstance. That has to be true, because the 10 centimecs you pay is worth 10 cents today, and the one mec you collect if you win will be worth $10 in that circumstance. So priced in mec, the probability is 10 percent that life existed on Mars one billion years ago. How can the probability depend on what you are betting?

It might seem you can get around this by using a currency that has the same value in all futures states of the world. But no such thing exists, any more than there is an absolute frame of reference in physics. Real risk can only be analyzed using real probabilities, which require some kind of real money in their definitions. Rocket scientists grew up in an era in which the value of money was highly uncertain. We were acutely aware that not everything can be bought or sold for dollars, and that the value of a dollar was highly dependent on future states of the world. We witnessed uncontrollable inflation and hyperinflation. Tax laws were complex and changed frequently, and the marginal rates were often very high. Governments were imposing wage and price controls and rationing many commodities—or forbidding buying or selling altogether. There were alternative currencies and abstract numeraires (a numeraire is a unit of account that assigns relative values to a set of items without necessarily being a medium of exchange or a store of value; an example is inflation-adjusted dollars), of course, but none were perfect. Therefore, we rejected the idea of a fully defined probability distribution that covered all possible future events. Our probability distributions might cover 95 percent or 99 percent of possible events, but would leave 5 percent or 1 percent as undefined outcomes, states of the world in which money was worthless, or in which outcomes were dominated by considerations that could not be priced.

Frequentism

Frequentism is the second major branch of probability theory. It uses long-term frequency as the fundamental definition of probability. This does not require money to define. Unfortunately, frequentism can't tell us the probabilities we want to know, like the probability that if I take a certain drug it will help me, or the probability that I will make money buying a certain stock. It can only tell us about probabilities created by the experimenter, and not even about specific probabilities, just average probabilities of groups of predictions. In a frequentist interpretation of a drug trial, there is no estimate of the probability that the drug works, only of the probability that the randomization scheme for assigning subjects to treatment or control groups—randomness the experimenter created—produced the observed result under the assumption the drug had no effect. Things are actually worse for observational studies where the researcher does not create randomness, such as an econometric study of the effect of monetary policy on inflation. For these, the researcher makes a statement about the probability of randomness she pretends she created.

A frequentist might test hypotheses at the 5 percent level. She can tell us that in the long run, fewer than 5 percent of the hypotheses she rejects will turn out to be true. That's mathematically true (at least if her other assumptions are correct) without reference to a numeraire. But why would we care? What if the 95 percent she's right about are trivial things we knew anyway and the 5 percent she's wrong about are crucial? Only if we can somehow add up right and wrong predictions to get a net gain or loss will her probability statement be useful for decision making. Moreover, the statements must have equal stakes or, as we'll see later, we must be in control of the stakes.

Both Bayesian and frequentist textbooks often obscure this issue by treating only problems in which only one kind of thing is at stake, or by assuming some perfect numeraire. But real problems almost always combine lots of different considerations, which means we need a numeraire to relate many different kinds of things, in other words, a form of money. Since no numeraire is perfect, we need to separate out the dangers and opportunities that cannot be measured in the money we are using for the probability calculation. To do otherwise is to be cold-blooded, to treat all dangers and opportunities as risks. This does not work in any human setting. It may be theoretically possible to imagine a perfect numeraire that puts a price on everything from God, honor, winning a game, and human life; to iPods, toilet paper, sex, and cocaine; to excitement, boredom, pain, and love; but if you make decisions based on probabilities stated in this numeraire, you will come to disaster. This is an empirical observation that I believe strongly. There is a better way to compute probabilities, a better way to manage risk.

If someone says, "Given my study of river height variation, there is a 1 percent chance this levee will be breached sometime in the next year," it sounds like a statement of physical reality, that might be right or wrong, but either way has objective meaning. That is not in fact true. The statement contains implicit assumptions about the value of human life versus property damage, since both are at stake. To a Bayesian, that assumption is implicit in the definition of the probability. Someone with different values would set the betting odds at a different number. To a frequentist, the statement doesn't make sense in the first place. The analyst should say, "I reject the hypothesis that the levee will be breached sometime next year at the 1 percent level." That statement is perfectly consistent with the knowledge that this levee is certain to be breached, but 99 other levees whose breach was also rejected at the 1 percent level are certain not to be breached. Only if I don't care about the difference between 100 levees each having a 1 percent probability of being breached versus 1 levee certain to be breached and 99 levees certain not to be breached, is the original statement a reasonable guide to action. That, in turn, requires that I regard each levee breach as having the same fixed cost that can be added up and that I care only about the expected number of breaks, not variation around that number. In a sense, it requires that I don't care about risk.

Looking at it another way, the original statement seems to imply that the researcher is indifferent between paying $1 for sure versus paying $100 if the levee is breached next year. But it also has to imply the researcher is indifferent between killing one person for sure versus having 100 people die if the levee is breached. There is no logical reason why a person has to accept the same stake ratio in both cases, and evidence from both behavioral and neuroscientific studies show that people do not, in fact, make the same answer. We call the person who pays a dollar for sure "a prudent insurance buyer" and the person who kills one person for sure "a murderer." We treat them very differently. We have not considered the more difficult case of how many dollars the statistician would pay for sure to save 100 lives if the levee is breached. And the probability could be different still if used for species extinction, votes, or excitement as numeraires.

Rationality

This is a deep insight into the nature of risk, money, and rationality. Suppose I observe that you will bet one apple against one orange on some event. I don't know what probability you assign to the event, because I can't divide apples by oranges. But then suppose I see you trade one apple for two oranges. Now I know you were giving two to one odds, meaning you think the event has at least two chances in three of occurring. I have separated your decisions into preferences—how much you like apples versus oranges—and beliefs—how likely you think the event is. This is the basic separation required for the modern idea of rationality, the assumption underlying most modern economic utility theory. It depends crucially on both gambling and exchange—on randomness and money.

A little reflection will show that this separation is entirely arbitrary. It's not how you think about risk. Suppose you're driving along an unfamiliar road and see that you're low on gas. You see a gas station charging 15 cents a gallon more than you usually pay. You have to decide whether to stop and pay the extra for at least a partial tank, or to drive on hoping to find a cheaper station before you run out of gas. In conventional theory, you estimate the probability and cost of running out of gas before finding another station, and also the probability distribution of gas prices at stations up the road. You have to also weigh the value of money versus the inconvenience of running out of gas. But no one does anything remotely like this, at least not consciously. You weigh probabilities and preferences simultaneously, without clearly separating between them. And people often act in ways inconsistent with any reasonable separation into beliefs and preferences.

Of course, the way you think about how you think can be misleading, whether compared to findings of neuroscientists and cognitive psychologists, or to actual behavior. Research does show that many risk decisions can be modeled as separation into beliefs and preferences, that is, to a probability distribution and a utility function. However, there are many different distributions and utility functions that are equally good at explaining brain activity and behavior for individual decisions, and no distribution and utility function that explains all decisions, even for one individual at one time. Intelligent risk management has to begin with a numeraire, plus the awareness that the numeraire does not cover all possible outcomes of the decision.

One of my favorite statistics stories illustrating the essential importance of getting the numeraire right occurred during World War II. The Allied Air Force was trying to decide the optimal amount of armor to add to bombers. This seems to be a problem in which the numeraire is obvious. Each pound of armor means one less pound of bombs, which means more bombing runs to deliver the same payload. Armor has to protect more airplanes than are lost on the additional runs. Wartime examples usually teach bad statistics, because war forces people to treat most dangers and opportunities as risks. Problems get brutally simple.

Anyway, the Air Force collected statistics on what parts of bombers suffered the most flak and shrapnel damage: leading edges took more than trailing edges, for example, and the underside took more hits than the rest of the plane. Obviously the places with the most frequent damage would benefit the most from armor. It sent the data to the great statistician Abraham Wald, asking him to indicate the areas where the armor would do the most good. Wald sent back a diagram that shocked the analysts. He put armor everywhere no damage had been recorded, and no armor on the places with most frequent damage. When Wald was asked why he put armor in places the bombers never took damage, he replied, "The bombers hit in those places never came back."

In this problem, using the obvious numeraire led to exactly the wrong conclusion. Putting armor where it seemed to do the most good meant protecting bombers against damage that was rarely fatal. You have to reverse the numeraire in this case, from protecting against recorded damage to protecting against everything but the recorded damage. After hearing the story, most people laugh at the foolish Air Force analysts. But the identical mistake is made frequently both by professional statisticians and nonstatistical professionals working probability judgments. It may be the single most common error in quantitative decision making.

Bets

Rocket scientists asked two more questions that occur immediately to anyone making an actual bet. Who are you betting with, and who is doing the betting? For the first question, are you going to name the same price for the Martian-life security when betting with a loudmouth idiot in a bar as with a scientific expert—or with a little green man from a flying saucer who lands in your back yard? If the other person knows less than you, you will set the price somewhere between what you think and what you think he thinks, in order to maximize your expected profit. If the other bettor knows more than you, you'll set the price at what you think he thinks, in order to minimize your expected loss. Traditional theorists usually have you betting with yourself, which is entirely pointless.

From a frequentist point of view, suppose someone tells you he rejects the hypothesis that it will snow in New York City tomorrow at the 5 percent level. He can say this every day with perfect accuracy, since it snows in New York City on fewer than 5 percent of days. But the statement is useless for practical decision making. Some days have essentially zero chance of snow; on other days snow is virtually certain. You could make money betting on this probability only if you are betting against an idiot, perhaps someone who thinks New York City is in Australia. If someone can make money predicting snowfall accepting odds set by the National Oceanic and Atmospheric Administration's National Weather Service, I'm willing to call that a useful probability. I'm even more impressed if the person can show a profit posting New York City snow odds on BetFair and taking all comers, or best of all, if the person can make money trading weather futures.

More generally, a frequentist probability claim says nothing about the strength of evidence backing up the computation. The person saying there is a 1 percent chance of the levee being breached might have simply looked at a list of historical levee breaches and noticed they happen on average once per hundred years. He might not know anything about this particular levee. He is violating no standard of statistical practice by making the statement without examining the levee, knowing which river it contains, searching out contrary opinions, or testing any assumptions. He could even put 99 true statements in a hat, along with a piece of paper reading, "The levee will not be breached this year," and pick one at random. If he picks the levee paper, he can say with complete accuracy that the chance of drawing an untrue statement out of a hat with at least 99 percent true statements is 1 percent or less, so he can reject the null hypothesis that the levee will be breached at the 1 percent level. It is not the significance level that tells you the reliability of a frequentist statistical claim; it is the vigor and sincerity of the falsification efforts undertaken in measuring the significance. But academic papers always report the former. In too many cases the latter is either omitted or amounts to the authors betting against themselves.

It's just as important to know who is doing the betting. There may be lots of people making profits betting on something, quoting different odds. This is easiest to explain in thinking about financial markets. Suppose I study all the people making markets (that is, setting prices at which they will buy from or sell to anyone) in, say, oil futures. They set different prices, implying different betting odds on oil prices in the future. I'm only interested in the market makers generating consistent profits. But even within this group there is a variety of prices and also differences in the positions they have built up.

You might argue that the differences in prices are going to be pretty small, and some kind of average or market-clearing price is the best estimate of probability. One issue with this is none of the prices directly measure probability; all of them blend in utility to some degree. A person who locks in a price for future oil may not believe the price of oil is going up. She may be unable to afford higher prices if that does happen, and is willing to take an expected loss in order to ensure survival of her business.

Two other issues have greater practical importance, at least in risk management. Some of these market makers may just be lucky, pursuing strategies that generate small profits most of the time but occasional disasters that more than offset any gains. Following their probabilities leads to disaster. Other market makers may actually lose money on their oil future bets, but make money overall because they use the oil bets to hedge other bets in an overall profitable strategy. Following their probabilities is also bad. The probability we care about is that of a hypothetical risk-neutral oil market maker who makes consistent money in stand-alone oil futures trading averaging over all future scenarios. That turns out to give significantly different probabilities than our subjective estimates, or long-term frequency, or market-clearing prices; even if we agree on numeraire and the identity of the bettors on the other side.

The result of all this is rocket scientists invented their own notion of probability. A probability distribution can only be defined with respect to a numeraire, and therefore cannot be defined for all possible outcomes. If you flip a coin, heads you win a dollar and tails you lose a dollar, the coin could stick to the ceiling or land on its edge, you could win the bet and not get paid, or you could get paid but dollars might be worthless at the time—or a 100 percent tax on gambling winnings could be passed while the coin is in the air, or you could die before the coin lands. You might try to list all possible events and assign probabilities to each, but it's a hopeless task for practical risk decisions. On the other hand, we made the empirical discovery that estimating the sum of the probabilities you could define reliably was highly illuminating, often a more valuable quantity to know than the statistic you were estimating in the first place. Meaningful probabilities also might not exist because no active betting market or reasonable hypothetical betting market existed to define whom to bet with, or because no one could be trusted to make a profit in that market.

That may seem to be a defective definition of probability, but consider the alternatives. Bayesians claim there is always a probability, but there must be one for every individual. In practice, Bayesians often find they have to resort to "improper priors" in which probabilities do not add up to one or choose probabilities for mathematical convenience rather than subjective belief. Bayesians who refuse to do this on principle must live in ivory towers, because they cannot tackle real decision problems. Frequentists often find no probability is defined, and the same hypothesis will have a different probability for every experiment. There is no rigorous way to combine probabilities from different experiments into a single number. Frequentist probabilities also do not add up to one. When a frequentist rejects a hypothesis at the 5 percent level, that does not mean the negation of the hypothesis has a 95 percent chance of being true. It's possible to have a set of mutually exclusive hypotheses, all of which can be individually rejected at the 5 percent level.

Rocket scientists also believed that probabilities could not be defined exactly, only up to a bid/ask spread. Unless someone can make a profit taking bets from everyone, there is no one competent to define the probability of an event.

Exponentials and Culture

But we're not going to talk only about play, probability, and money. There will be an entire chapter on exponentials. The mathematical definition of an exponential is something with a rate of growth proportional to its level. The bigger it gets, the faster it grows. These relate to risk for three reasons.

The first reason is that if you examine a sudden, dramatic change, it usually turns out to be an exponential. It was small and growing slowly for a long time, and was unnoticed as a result. Exponentials work both ways: The smaller it is, the more slowly it grows. Once it starts getting big, it grows so fast that it seems to come out of nowhere.

By that time it has lost its exponential character, as nothing physical can grow forever. It hits up against some limit. People describe it as a Black Swan, an unanticipated event—in fact, one that was impossible to anticipate—and focus on the sudden growth and spectacular collision with its limit. Anyone serious about risk has to concentrate on the exponential nature instead. Once the thing becomes obvious, it's usually too late either to avoid its danger or to exploit its opportunity. Nonexponentials are much easier to deal with. If they are big or fast growing, you notice them. If they are small or slow growing, they don't cause a lot of problems or offer a lot of opportunities.

The second reason to discuss exponentials goes back to the 1956 discovery by physicist John Kelly that exponentials trump risk. If you can organize your risk taking to get the optimal level of exponential growth, you end up better off than you can possibly be using any other strategy. Mathematician and hedge fund innovator Edward Thorp named the strategy "fortune's formula." In a sense you conquer risk since your outcome is guaranteed to be better than that of someone who avoids risk. It's not risk if you can't lose. Kelly's result was theoretical, and we do not know how to conquer risk completely. But his work has led to sophisticated practical techniques for harnessing the power of exponentials to exploit risk.

The last reason to study exponentials is that risk-avoiding people often use them recklessly. Exponentials are powerful and dangerous, and once they're big enough to matter they never last long. When a CEO targets a compound average growth rate for earnings, she's trying to build an exponential. She will probably fail, but if she succeeds the company will soon hit a limit and the fallout will be unpredictable. When an economist justifies a government policy using projected future growth rates, he's relying on exponentials to bail out an idea that cannot float on its own. The opposite error is possible as well. Alarmists often use exponential growth rates to conjure sky-is-falling scenarios that would be laughed at without the mathematical camouflage.

Finally, we're going to discuss how risk is embedded in culture. One of the most difficult aspects of managing risk is competition from older belief systems, just as science sometimes finds itself in conflict with superstition or religion. A lot of power in the world is distributed according to claimed ability to make good decisions under uncertainty, through either superior prediction skills or a talent for managing events as they arise. This includes people from shamans to mathematical modelers, from priests to statesmen and generals. This is nice work if you can get it, since it's difficult to tell the legitimate practitioners from the charlatans. If you claim to be strong, or fast, or a good chess player, it's easy to establish the truth. But if you claim to be able to interpret the will of the gods, or to be a wise policy maker, or that your patent medicine helps ill people, or that the best chance of winning the battle is for everyone to do as you say, it takes a long time to compile evidence one way or the other. Being clever at explaining away errors and taking credit for accidental successes leads to more acclaim and power than making good risk decisions in the first place. In fact, good risk decisions usually lead to the appearance of alternating complacency and erratic actions. They are hard to defend even after the fact to people who were not involved in the decision making.

While few people would disagree with that last paragraph, I'm going to argue that it runs much deeper than is usually supposed. Bad risk management is ingrained into social institutions and popular theories. Among other things, that helps explain why it took so long after the major discoveries in the field for a book to be published that covers them thoroughly, and why so much nonsense about risk is written every day. Half of good risk management is just identifying and eliminating the bad risk management. That exercise can be extremely disruptive and can generate strong reactions because it challenges a major traditional base of power.

Payoff

What is the payoff for working through the four previous topics, as well as some more conventional risk management material? I will give you simple and logical answers to a variety of questions about risk. You'll have to decide for yourself whether the answers are true, but none of them will be airy generalities. I will not ask you to take anything on faith. The logic and evidence will be presented clearly, as will the historical development of the ideas. I believe everything in here is true, and I have tested it over many years of actual risk taking, plus observation of others. I cannot claim it is accepted widely, as it is not even known widely. But it does represent the consensus of successful modern quantitative risk takers in finance. It's how the global financial system works—and the global financial system is increasingly determining how everything works.

I have presented the material in this book over the years in articles and speeches, with mixed success. I find it easiest to communicate to professional risk takers who are good at mathematics. I hope this book will help broaden the audience to people who are not particularly fond of mathematics, and who take risk but do not focus their profession on risk. The group I have the hardest time with is risk avoiders who are good at mathematics. They seldom disagree with me and they claim to understand, but we talk at cross-purposes. I say "risk is good," and they agree, thinking I mean that risk must be accepted in order to improve expected outcomes. That makes risk a cost, something bad that you accept in order to get something good, which is not at all what I mean. In my terms, they treat all risks as dangers. I talk about making decisions and they agree, mentally imagining that means giving advice to others.

To avoid misunderstandings, I have reinforced the main points of the book with graphic material—comic strips. These are an important part of the book. If you find yourself agreeing with the text but not understanding the comics, you're probably missing the point. If you see eye to eye with me on the comics, you've absorbed the important ideas, even if you read nothing else.

I will ask you for a fair amount of trust. I have a big story to tell, with a lot of apparently disparate elements. We'll cover all of human history and the global economy and even bigger stuff. Unless you work in finance, and possibly even if you do, some of the ideas likely will be completely new to you, and strange. Some may contradict things you have accepted in the past. It may not all fit together until the last chapter. I've tried to make it interesting enough for each part to stand on its own, but this is not a collection of essays. If you will give me your attention for a few hours, I undertake to reward it.