CHAPTER 1

Now It's Personal

The future ain't what it used to be.

—Yogi Berra (1925–)

Over the last few decades, there has been a profound shift in investment responsibility – away from the professional institutions that traditionally manage financial assets and right into your lap. Today, tens of millions of people are faced with making personal investment decisions that will have a big impact on their future quality of life – but only a few feel like they know what they are doing. This shift in responsibility is not simply the result of some cyclical government policy change or a conspiracy among big corporations to avoid funding retirement for their employees, but rather it is a necessary response to a global population that is rapidly aging. In the future, whether you like it or not, you are going to be responsible for managing your retirement investments. The good news is that within the last 10 years, changes in technology have opened the door to the widespread use of modern financial economic methods in personal investing. For the first time, it is possible to apply the same rigorous techniques used by the largest institutional investors to personal investment choices. In this chapter I explore these techniques, how they became available to individual investors, and how they enable a powerful new way of looking at personal investment decisions—outcomes-based investing.

How the heck did we end up here?

Since when did it become accepted that truck drivers, marketing analysts, lawyers, bakers, computer programmers, dentists, and teachers should all become expert investment managers? Who decided that we all need to become well-versed in mutual fund selection and dealing with the impact of investment risk? I don't recall voting on this, do you?

If you are reading this book, chances are that you have at least a passing interest in becoming a better investor. For some people, the subject of investing is a passion or a hobby, for others it is more of a reluctant obligation – like cleaning the house. The reality is that we now live in a world that expects everyday people to understand a great deal about investing. Expectations go far beyond how to balance a checkbook or apply for a credit card. Today, virtually everyone over the age of 18 is expected to understand concepts such as the value of compounding, progressive tax rates, the differences between equities and bonds, and the benefits of tax-deferred savings. Those who do not bother to acquire this knowledge are roundly criticized by TV commentators, newspaper editors, and personal finance gurus for being dangerously out of touch. The era of the self-directed investor has truly arrived – even if most of us were not ready for it.

Unfortunately, investing is a sprawling and complicated subject, rich with intimidating jargon and numerous conflicts of interest thrown in for good measure. It is no surprise that many people approach investing decisions with a sense of dread, or at least acute anxiety. Why do *I* have to make all these investment decisions? Generally speaking, I am not expected to diagnose the strange knocking sound coming from the hood of my car on my own, nor am I presumed to figure out the origin of a blotchy rash on my son's back without expert help. Somehow, when it comes to investing, the prevailing opinion seems to be that you need to be educated enough to make good decisions on your own. But most people have better things to do with their lives than to become self-educated investment experts. And so, here we are, with tens of millions of individual investors directing how trillions of dollars of retirement money is invested, and bearing the consequences of those decisions for better or worse.

But why? Is this really the best way to make sure that we retire comfortably? It turns out that the dramatic movement towards more individual responsibility in retirement investing has its roots in a more fundamental and irreversible trend – we are all living quite a bit longer than we used to.

A CHANGING WORLD

Sixty years ago in the United States, things were a bit different. Back then, there were fewer decisions left to individuals. Mostly such choices were out of your hands. In the past, retirement was based on the concept of a three-legged stool, represented by family, government, and the employer. For much of recorded history, the most important leg of the stool (and often the only one) was your family. It was generally accepted that when you got old, your kids or extended family would take care of you (assuming you had any and that you had been nice to them). This was generally not an undue burden, as families were larger and most relatives tended to live in close proximity to one another. Moreover, retired people did not typically live that long after they stopped working. You might find it hard to believe, but for a male born in 1900, the National Center for Health Statistics reports that the average expected life span was only 48 years. Often, people would simply work until they were no longer able to continue on. Of course, things are a bit different today. Families are smaller and often spread out across large geographic distances. We are also living much longer. By 2003, the average 65 year-old woman in the United States could expect to live an *additional* 20 years. Moreover, steady advances in medical technology continue to increase average life spans. With earlier retirement dates, and longer life spans, retirement is now often measured in terms of several decades, not years.

In the summer of 1935, another leg of the stool was added—*Social Security*. With Franklin Roosevelt's signing of the *Social Security Act*, the U.S. government would provide a safety net for those in retirement of a lifetime guarantee of income. Social Security was a monumental change in public policy, and is one of the enduring legacies of the modern welfare state. When the system was created, however, there was a large working population paying into the system to support a relatively small retired population. Now, with the aging of the baby boomer population, the system's finances are under strain. Under current course and speed, either benefits will have to be reduced or payroll taxes will have to be significantly increased (or both) in order to fund the promised benefits, due to the growing population of retirees. Reaching such a political compromise has proven to be a difficult task in recent years, and so the magnitude of the funding problem continues to grow. Of course, this stalemate has potential serious consequences for taxpayers like you and me.

The third leg of the retirement security stool (at least in the last 50 years or so) was the employer. For those lucky enough to work for large companies, *defined benefit* (DB) plans often provided generous guarantees of retirement income, if you stuck with your employer sufficiently long. A DB plan provided employees with a guarantee of future retirement income funded from a pool of assets managed by the employer (usually without any input from the employees). The money was professionally managed on the employees' behalf, and if the markets went down, the company made contributions to the plan to insure the benefits were appropriately funded. If you were part of such a plan, all you had to do was show up to work each day for 30 years or so, and you were often comfortably set for retirement.

However, even in their heyday, defined benefit plans were not available to everyone. It is estimated that only about a third of people over the age of 55 receive any form of pension income.[1] In many situations, employees failed to qualify for any DB income because they changed jobs before they became eligible for the plan, or the companies they worked for simply did not offer such pension benefits. In the most recent generations, the availability of DB pension plans for new employees has declined dramatically.

Why are DB plans going away? One significant factor is the fact that people in today's economy tend to change jobs frequently. Under a traditional defined benefit plan (which provides guaranteed benefits), an employee must stay with a single employer for more than 10 years or more to accumulate significant retirement benefits. When the concept of "employment for life" was more popular, these restrictions were not as big a problem. However, today, most people change jobs multiple times during their careers, and hence would not accumulate significant benefits under the traditional defined benefit structure. But the biggest reason for the demise of DB plans is cost. It is increasingly expensive to run a defined benefit plan.

Most companies invest their defined benefit plan assets in a broadly diversified portfolio of stocks and bonds that is used to pay the future income of retired workers. If the market goes down, the company is obligated to make additional payments into the plan to ensure that it can pay out its promised retirement benefits. However, this creates a potential problem for companies such as General Motors and Ford, with large retired populations relative to their current workforce. With a DB plan, the company may have to come up with more money to fund the plan at precisely the time when money is scarce. When stock market returns are poor, chances are that the company is also feeling the pinch from a downturn in its business. The need to make payments into the plan when markets are down and when the cost of those payments to the company is highest (more on this concept later) makes DB plans a potentially expensive and risky proposition.

The magnitude of this cost problem is further accentuated when we take into account the impact of an aging population. Because we are all living longer, on average, the number of older retired people has grown dramatically over time, while, at the same time, birth rates have declined significantly. The following figure is an example of a *population pyramid chart* using 1950 data from the United States Census Bureau.[2] Such graphics are called pyramid charts due to their distinctive shape, consistent with the distribution of age groups in a growing population. Historically, in most societies, there have been a large number of young people relative to the population of older persons. If the age of the population is plotted on the

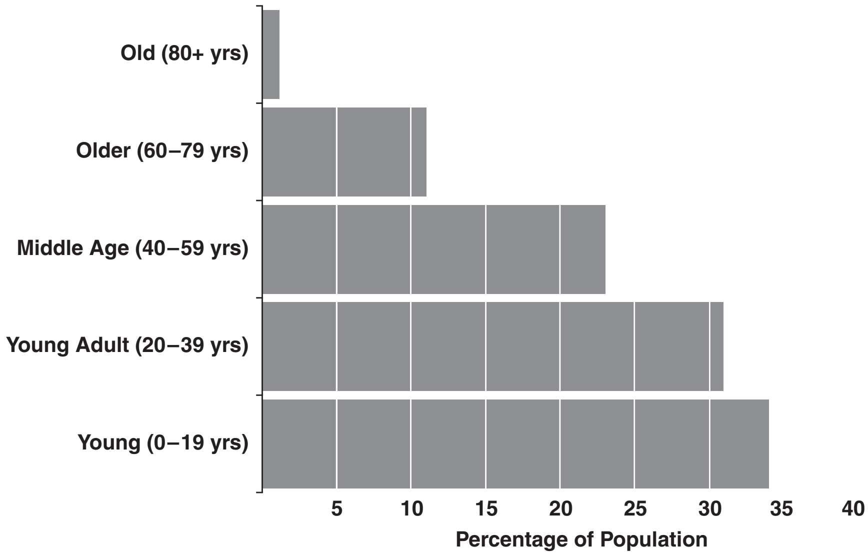

FIGURE 1.1 U.S. Population by Age Group in 1950
Source: U.S. Census Bureau.

vertical axis, and the length of the bars represent the proportion of the population in that age group, you get a distinctive triangle shape. In the United States of 1950, this pattern was clearly evident, as shown in Figure 1.1.

In 1950, the U.S. population was growing fast, and while life spans were gradually increasing, the proportion of older people had not yet begun to grow disproportionately. Now contrast this with the picture in 2006 (Figure 1.2).

The combination of lower average birth rates and greatly increased longevity has dramatically changed the relative proportions of active workers supporting retirees. For instance, in 1950, there were 47 people in their prime working years (between the ages of 20 and 59) for every person age 80 or older. By 2006, this ratio dropped to only 15—a decline of 68 percent. By 2050, there will be only 6 employees in their prime working years for every person age 80 or older. This implies a very different division of resources between generations than what has been historically observed.

The net effect of this demographic shift is that companies (and many governments) are increasingly unable to afford guaranteed income for life to the growing retiree population. If you guarantee income for retirees, it means working employees must fund this income with increasingly high taxes, or they must be willing to assume the financial risk of pension assets

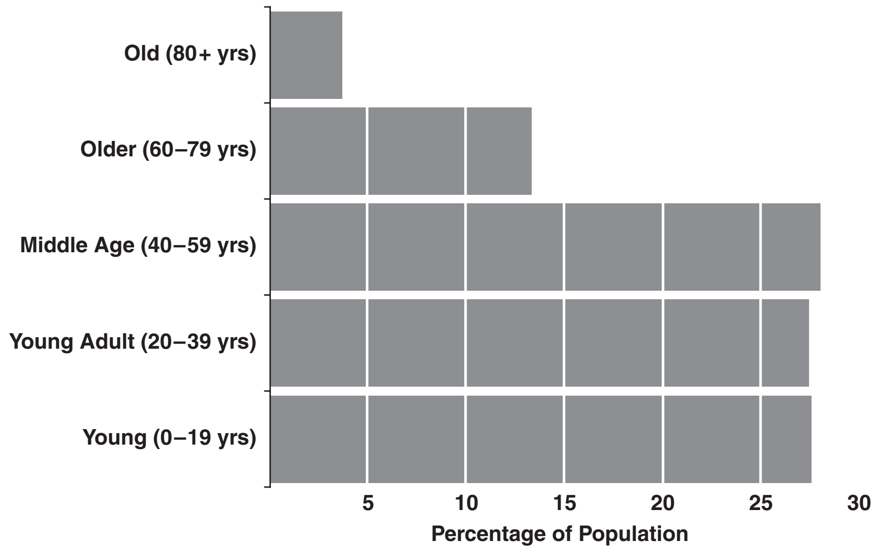

FIGURE 1.2 U.S. Population by Age Group in 2006
Source: U.S. Census Bureau.

falling short when markets turn sour. When there are only a small number of retirees for each worker, as was historically the case, this is not such a big problem. But when the retiree population gets large, look out. The burden on the working population becomes too onerous. Just as a company that takes on too much debt greatly increases the risk to stockholders and employees, guaranteeing retirement income for an ever-increasing retiree population is not an economically sustainable proposition. If retirees don't have any risk in their retirement income, then existing workers must assume all the risk. But now there are simply too many retirees to support—something has to give. The only way out of this predicament is to have the retired population share the risk of market downturns with the working population—that is, to make individual retirement income dependent on market performance. Enter the 401(k).

THE GRAND SOCIAL EXPERIMENT: 401(k)

In the last 20 years, the United States has seen enormous growth in *Defined Contribution* (DC) plans such as the *401(k)*. What was once an obscure part of the tax code, the 401(k) plan is now the primary retirement vehicle for most U.S. workers. In fact, DC plan assets now exceed the total value of DB plans by a wide margin. In 2006, according to the Federal Reserve, DB plans totaled \$2.3 trillion, while DC plan assets totaled more than \$3.3 trillion.[3] Defined contribution plans differ from DB plans in that the *individual* is responsible for making the investment decisions and for bearing investment risk, not the employer. With the rising prevalence of 401(k) plans, *Individual Retirement Accounts* (*IRA*s), *Health Savings Accounts* (*HSA*s), *529 Plans*, and other plans where individuals must manage their assets, more and more financial responsibility is being placed squarely on the shoulders of everyday people.

In many ways this is a grand social experiment—never before has so much financial responsibility been in the hands of so many nonexperts. We do not yet know what the results of the experiment will be for the United States. If financial markets perform well in the coming decades, people save aggressively and invest prudently, the experiment will likely be successful. However, if this rosy state of the world does not come to pass—there are going to be a lot of grumpy and cash-strapped retirees wondering how all this came to be.

The rise of defined contribution plans like the 401(k) has dramatically altered the landscape for individual investors. As of 2006, there are now more than 55 million individuals managing trillions of dollars of defined contribution assets.[4] For most employees hired in the last decade or so, the defined contribution plan is their principal (and often only) retirement benefit. In fact, DC assets now account for 61 percent of overall retirement assets in the United States. *Importantly, there is no turning back from this trend*. Since the economic need for defined contribution plans like the 401(k) is based on long-term demographic trends, there is no going back to the days of guaranteed retirement benefits. In the future, your retirement paycheck will depend to some extent on market performance and how you choose to allocate your investments. Put another way, if markets decline, you will be the one holding the bag—not your employer.

THE KNOWLEDGE GAP

Sadly, our educational system has been woefully behind the curve in preparing people for the heavy new financial responsibilities of a self-directed investment world. Unless you happen to work in financial services, what you know about these subjects is likely to have come from friends, family, and the media, not from high school or college coursework in economics and finance. Relatively few people have any significant exposure to academic finance or economics, and hence many of the important concepts that govern good decision-making are unknown to those who could benefit from them. Much of what we do learn is incomplete and reduced to three bullet-point sound bites on talk radio or cable TV shows. Most people rely on a haphazard mix of conventional wisdom, rules of thumb, and outright dart-throwing to manage their personal investments. The predictable result is that most investors are not doing a very good job of managing their assets. Unfortunately, most of us are flying by the seat of our pants.

To successfully navigate the waters of modern financial life, you have to make a basic choice. You can educate yourself so that you are able to make informed decisions on your own, or you can hire a competent and objective financial professional to help you (the third option of avoiding the subject, though immensely popular, is generally ill advised). The do-it-yourself approach has several advantages. You are in control, you do not have to delegate decisions to other parties, and you develop first-hand experience with the process of portfolio management. The downside is that it takes a lot of time and energy to become well educated on investments, and a comparable amount of time to monitor your portfolio on an ongoing basis. Also, if things go wrong, you have no one to blame but yourself (don't laugh; this is a major element of stress for many investors). In fact, most surveys on the subject, including those conducted by Financial Engines, estimate that only 10 to 20 percent of the population is truly comfortable with the do-it-yourself approach. The vast majority of people are looking for help, either to validate their own decisions, or to delegate the responsibilities to a professional who they can trust. Of course, finding an honest, qualified, and reasonably priced investment advisor can be a challenge.

THE TRADITIONAL ADVICE MODEL

It is true that many, if not most, individuals can use help with making investment decisions. What is less clear is how to get objective, personalized advice at a reasonable cost. If you are lucky enough to be wealthy—say, financial assets of a few million dollars or more—there are thousands of advisors eager to work with you. Of course, not all of these eager advisors know what they are doing, but nonetheless, there is a big pool to choose from. If your means are more modest, sources for high quality and objective investment advice can be particularly difficult to find. Like many aspects of life, the *caveat emptor* proverb applies to investment advice as well. As we will see, a big part of being an informed consumer of investment advice is being able to detect when your advisor may be making claims that don't stand up to scrutiny.

Up until a few decades ago, people with modest assets generally did not have a need for personalized investment advice. Most personal savings were simply kept in the bank. Few middle class individuals had reason to own mutual funds or even individual stocks until the early 1980s, with the rise of 401(k) plans. Since most clients for investment advice were wealthy, the advice model that developed was based on the concept of one or more expert advisors personally interacting with each client and developing a customized investment strategy. Generally advisors were expected to have high levels of expertise (often across multiple domains), and as a part of their services spent substantial time interacting with each client. The costs of providing such investment advisory services are considerable, typically paid as an annual percentage of client's assets. Fees for traditional investment advisory services generally range from 75 basis points (0.75 percent) to more than 200 basis points (2.00 percent) of assets under management, depending on the types of services provided and the size of the client account. For the process to be profitable for the advisor, the client account must generate sufficient fees to cover the overhead and compensation for the experts, who are typically paid handsomely. This model works well for high-net-worth individuals or families with big assets who can afford to pay for the personalized attention. But the reality of the advice business is that people with a limited number of zeros at the end of their account balance are simply left out. It is not economically feasible for well-paid experts to spend tens of hours putting together a financial plan for an individual investor with only $50,000 in financial assets.

To service the needs of investors who fall short of the high-net-worth benchmark, financial services firms have typically relied upon the structure of product-based compensation, or *commissions*. Rather than charging an explicit hourly or asset-based fee for advice, the compensation of the advisor or broker is structured in the form of commissions on the sales of investment products. However, this approach suffers from a big conflict of interest, as some products invariably result in larger commissions for the broker or advisor than others. The result is that the advisor may have a vested interest in selling certain products (such as the funds of their own firm), even if it is not necessarily in your best interests. This fundamental conflict of interest between commission-based brokers or advisors and the clients they serve is at the heart of many lawsuits and arbitrations that occur in the industry each year. Actually, it gets a little worse. Since commission-based compensation is triggered by transactions (buying and selling of products), the commission model also creates incentives for excessive trading. All of a sudden you find out from your broker that the great fund you bought 18 months ago no longer makes the cut.

Fundamentally, the commission-based structure does not alter the basic economics of providing advice—it merely changes (and conceals) the form of payment. Since the costs of the commissions are often built into the financial products (through *expense ratios* and *loads*), they are not as visible to the investor as in the case of an explicit advisory fee. However, the impact is the same—lower overall investment returns for the investor, and higher compensation for the broker or advisor. As we will see in Chapter 7, it is critically important to understand how fees impact the value of your investments.

INSTITUTIONAL TOOLS OF THE TRADE

Successfully addressing the needs of average investors requires a different type of advice model—one that leverages technology to change the economics of providing personalized investment advice. By using technology to automate much of the investment analysis process, it is now possible to provide high-quality advice at a much lower cost than was historically the case. The basic idea behind Financial Engines' advisory services is to emulate how the big guys manage investments using state of the art financial technology, but do it for millions of individual investors instead of a handful of giant funds. Want to know how to best invest your money? Apply the same techniques used by the largest institutional investors investing hundreds of billions of dollars.

For instance, consider the problems faced by pension funds. Putting together well-structured portfolios to fund future pension liabilities is not a particularly easy task, and usually requires significant expertise and analysis. Pension funds use advanced techniques to evaluate risk, forecast possible investment outcomes, value future liabilities, structure portfolios, and evaluate the performance of investment managers. While such techniques might seem intimidating to most investors, these same ideas are directly relevant for your personal retirement investing.

Albeit on a smaller scale, individuals face many of the same problems as institutional money managers like pension funds and university endowments. The biggest difference is that institutional money managers are able to apply much greater resources in evaluating these questions than most individuals (unless your last name happens to be Buffett or Gates). Large institutional investors such as the General Motors pension plan and the Stanford University endowment fund have benefited from the tools of modern financial economics for decades. But it has been only recently that comparable techniques have been applied to the needs of average individual investors. The reasons for this are myriad, but the biggest factor has been cost. Up until the last ten years or so, it has simply been too expensive to rigorously apply the techniques of modern financial economics to individual investment problems. In the mid-1990's, contemporaneous with the rise of the Internet, technology began to make possible a new approach to personal investment advice. A key development in this process was giving individual investors the ability to realistically view how their investments might perform in the future.

To understand how such techniques can be used to make better investment decisions, it is helpful to introduce some of the most important tools used by institutional investors. In later chapters, I will use these tools to illustrate important concepts with some real-world examples. It is not crucial that you understand the details of such methods, but having some basic understanding of what they are and how they work is very useful, whether you make your own decisions, or use the services of an investment advisor.

As a first step, when choosing an investment strategy, you want to develop a realistic view of the potential future outcomes associated with your investment choices. With knowledge of the possible range of outcomes, you can make an informed choice about the appropriate level of risk and savings to construct an investment strategy that will have a high probability of success. The methods for making such financial projections run the gamut from simplistic calculations based on a constant rate of return to sophisticated simulations that take into account the complex relationships among economic variables and asset returns. Among institutional investors, the most utilized method for understanding the range of possible investment outcomes is a technique called *Monte Carlo simulation*.

What Might the Future Hold?: Monte Carlo Simulation

Monte Carlo simulation provides a powerful way to analyze problems that involve uncertainty. The formal development of the Monte Carlo method dates from the Manhattan Project to develop the atom bomb during World War II and is based on the work of two famous mathematicians, Stanislaus

Ulam and John von Neumann. At the time, they were focused on random neutron diffusion in fissile material—how far little atomic particles, called *neutrons*, could penetrate into various types of metal. Solving this problem experimentally was expensive and dangerous (neutrons can give you cancer by damaging your DNA, not to mention they make your eyebrows fall out). While they had developed good theoretical models for the behavior of a single neutron, summarizing the interactions between large numbers of neutrons was so complex that it was analytically intractable (they could not even write down the equations, much less solve them). Their solution was to perform a large number of random hypothetical experiments simulating the paths of individual neutrons through various materials and then add them up. By averaging the behavior of individual particles over enough such simulations, the mathematicians could accurately estimate how far neutrons would penetrate different types of materials (they purportedly chose the name *Monte Carlo* after its namesake in Monaco in honor of Stanislaus's gambler uncle).

From its early use in nuclear physics, Monte Carlo simulation made its way into other fields of study and is now used in a wide range of domains, including weather forecasting, traffic analysis, chemistry, genetics, statistics, and investment analysis. The method is particularly useful in solving problems where the probabilistic behavior of a system is very complex. The probabilistic behavior of an investment portfolio consisting of a mix of different securities represents one such complex system. That is why institutional investors such as defined benefit pension managers have used Monte Carlo simulation for more than two decades to evaluate the likelihood of meeting future liabilities. More recently, Monte Carlo simulation has rapidly evolved from a mathematical curiosity into something of a marketing buzzword in the financial services industry. The purpose of Monte Carlo investment simulation is to provide a better understanding of the *range* of possible outcomes when the future values of the assets in the portfolio are uncertain. That is, Monte Carlo simulation can't help you predict the future (unfortunately no one can do that), but it can show you the range of possible future outcomes you might expect with different market performance.

These simulation tools provide important benefits for investors, including the ability to:

- Calculate the probability of reaching investment objectives
- Provide forward-looking measures of investment risk
- Test drive investment strategies prior to implementation

Virtually all types of financial assets have some degree of uncertainty in their future values. For instance, the future value of a portfolio of stocks may be worth more or less than its current value based on whether the stock market moves up or down. Even guaranteed securities like bank CDs, which are advertised as being very low (or no) risk, have some uncertainty in their future values. While the rate of interest may be guaranteed at the end of the period (say, one year), the actual value of the assets is dependent on how inflation behaves over the holding period. If inflation is low, the value of the CD account will be worth correspondingly more; if inflation spikes up unexpectedly, the value of the CD will be lower (in today's dollars), since the value of the accumulated returns will have been eroded by inflation.

The simulation of investment portfolios provides a powerful tool for investors. It allows the evaluation of different decisions (such as how much risk to take, how much to save, etc.) and how they affect the range of possible investment outcomes. Rather than evaluating such choices in the dark, we can directly observe how different decisions impact the outcomes that you ultimately care about. For instance, you can calculate how additional savings of $100 per month influences the probability of reaching your desired retirement income goal. Such simulations can be quite realistic, taking into consideration the impact of market performance, fees, specific security risks, taxes, account distribution rules, and even uncertainty in how long you might live.

The mechanics of Monte Carlo simulation are straightforward in principle but often complex in implementation. The Monte Carlo method is predicated on repeated randomized experiments of a statistic of interest (for instance, the future value of an investment portfolio). The idea is to model the random behavior of a complex system (such as a portfolio of investments) by repeatedly sampling different possible random values of the variables that drive the outcomes of the system (such as the future values of each security in the portfolio).

A Simple Simulation Example: Coin Flipping To illustrate some intuition behind how such models work, let's examine a very simple Monte Carlo simulation. As an example, consider the range of possible values for flipping a coin 10 times in a row and counting up the number of heads observed. The statistic of interest (the number of heads observed in 10 coin flips) can take on values ranging from 0 to 10. However, values in the middle of the range (4, 5, and 6) are much more likely to be observed than values at the extremes (0 or 10). To understand what the range of possible values and their associated probabilities looks like, one can simply simulate the flipping of the coin a thousand times or so on a computer and tabulate the results (you could try this by hand, but you would likely expire from boredom).

To create a single observation, a computer program generates random numbers to simulate 10 flips of a coin with a 50/50 probability of heads or tails in each flip. After the 10 simulated coin flips, the number of heads is counted and stored as one possible value of the statistic (let's say the first count comes up with six heads). Next, the process is repeated to calculate a different possible value for the variable (say, four heads). By running this simple experiment a thousand times, one derives a thousand possible values for the experiment—a distribution of outcomes for the number of heads in 10 coin flips.

Figure 1.3 shows the distribution of outcomes from a thousand random trials of this simple experiment.

This chart is called a *histogram,* and shows not only the range of possible values (0 to 10), but also how often each value is observed—the frequency of occurrence. As you can see, the most frequently observed number is five heads (as we would expect for a fair coin flip); however, there are many observations that are above or below five heads. In fact, for this sample of 1,000 trials, only 24 percent of the time did the number

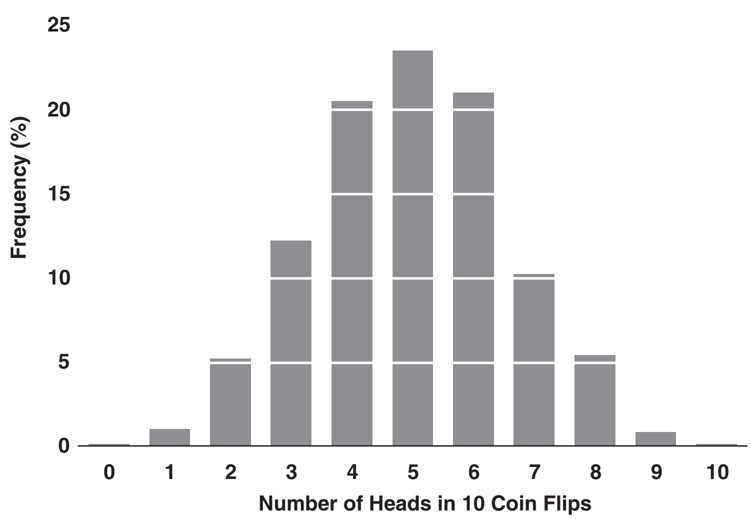

FIGURE 1.3 Plot of the Total Number of Heads Occurring in 10 Coin Flips, Based on 1,000 Trials.

of heads match the expected number of five. The chance of observing a very small or large number of heads is much less common. For instance, in this sample of a thousand trials, there were only eight observations of nine heads (less than 1 percent of all the experiments). In all of the 1,000 simulated trials, there was only one observation of zero heads in ten flips, and only one observation of ten heads in a row. Generally speaking, the further away the value is from the expected value, the lower its probability of being observed. If we were to run this experiment another 1,000 times, we would see a similar, but different set of observations. To get an increasingly accurate estimate of the true probability for each observed number of heads, we would need to run more and more simulations.[5]

Immediately we see that the statistic in this simple experiment (the number of heads in ten coin flips) has *volatility*. It can range over a number of different values, depending on the randomness of the individual coin flips. As human beings, we like to think of outcomes in terms of the typical result. It is important to realize, however, that most of what we experience is not the expected case, but one of many other possibilities. In this experiment, we observe a different number of heads that is different from the expected case (five heads) about 75 percent of the time. This means that you only see the "expected" result about 25 percent of the time. As we will see in later chapters, investment returns have a very similar property—their range of possible outcomes can and often does vary dramatically from the expected value. Making good investment decisions means recognizing that you will more often than not, end up with a different outcome than the expected result.

Simulating Investments Simulating a portfolio of investments is considerably more complex than modeling coin flips (and a lot more interesting, for that matter). A thorough explanation of how modern investment simulation models work would fill a book or two, which thankfully, I won't attempt here. For those that are interested, I provide a brief overview of how such models work, but it is not necessary that you understand all the details in order to benefit from investment simulation.

The hard part in developing a realistic investment simulation is specifying a model for how the returns of investment securities vary over time in response to different realistic economic conditions. This is where financial economics comes into play. Once you have a robust model for the investments, then you can apply the well-known techniques of Monte Carlo simulation to generate results rather easily. But if you have a poor economic model, even a large number of simulations and the fastest computers in the world won't get you good results. Investment simulation is a classic example of garbage in, garbage out. Getting the economic model right is the most difficult part of the process.

While specific simulation models for investments differ in their implementation details, the general process for creating portfolio scenarios is reasonably straightforward. The idea is to generate possible future paths that the economy might take and then determine how a particular portfolio might perform in each set of market circumstances. By generating thousands of such scenarios, one can accurately describe the range of portfolio outcomes that might be experienced and determine the relative probability of observing each outcome.

Generally, investment simulation models are built in layers. At the bottom layer is a model for estimating the joint behavior of key economic variables like inflation and interest rates. The process is complicated by the fact that interactions between economic variables are not purely random, but are partially governed by relationships between them. For example, when inflation unexpectedly increases, this has a predictable impact on current interest rates (they tend to go up). Because of this complexity, simulation models differ in their ability to capture the observed properties of economic and investment variables, and in their consistency with accepted financial economic theory. More sophisticated models use a combination of structural relationships (such as the link between today's inflation and inflation in the previous period) along with random noise to generate possible future values of economic variables. This way the future scenarios have randomness but still abide by theoretical constraints based on the linkage between different types of economic variables.

In the middle layer of most investment simulation models is something called a *factor model*, which generates possible returns for various types of asset classes (for instance, large capitalization stocks, corporate bonds, or international equities) given the values of the key economic variables determined in the first layer. The factor model may be simple, dividing the investment world into as few as three categories: stocks, bonds, and cash, or it may be more complex considering a dozen or more classes of investments. In the case of the Financial Engines simulation model, the investment world is divided into fifteen different asset classes spanning domestic and international equities and bonds.[6]

The top layer of the simulation model captures the linkage between specific securities and the returns in the underlying asset classes. In some models, a specific security may be related to one or more of the underlying asset classes. For instance, a balanced mutual fund might have investments in both bond and equity asset classes. Depending on the mix of assets classes in the specific security, its returns will tend to move with different segments of the market. A given security inherits its *expected return*, the level of return that you can expect, on average, from its underlying investments. For instance, a bond fund will tend to have similar expected returns to bond asset classes, while an equity fund will tend to have similar expected returns to various equity asset classes.

In addition, it is often necessary to adjust the expected returns of a specific security to capture characteristics that make its returns different from those of its underlying investments. For instance, a mutual fund charges fees in the form of an *expense ratio,* which lowers its expected return relative to the assets in its underlying portfolio.

To make the process more concrete, let us examine what it takes to create scenarios for a simple portfolio consisting of a single hypothetical equity mutual fund. At a high level, the simulation process consists of several steps that are repeated many times:

1. Start with current values of key economic variables (interest rates, inflation, etc.).
2. Generate possible values for the economic variables in the next period based on a random draw from an economic model with the current values as inputs.
3. Derive returns for each asset class for that period.
4. Derive returns for the mutual fund based on the underlying asset class returns.
5. Repeat steps 2 through 4 until you have a complete scenario.
6. Repeat steps 1 through 5 to generate multiple scenarios.

For instance, if we start with an initial $10,000, one possible future scenario for an equity mutual fund over a 40-year period might look like Figure 1.4 (assuming no taxes).

This is but just one possible path for how this fund might perform over the next 40 years. Depending on the economic environment and the performance of the securities held by the fund, the outcomes could be very different. This becomes evident as we add additional possible scenarios to the chart, as shown in Figure 1.5.

As you can see, there is already quite a wide range of potential outcomes from the mutual fund portfolio with only five scenarios. With just five scenarios, the $10,000 initial investment outcomes range from as little

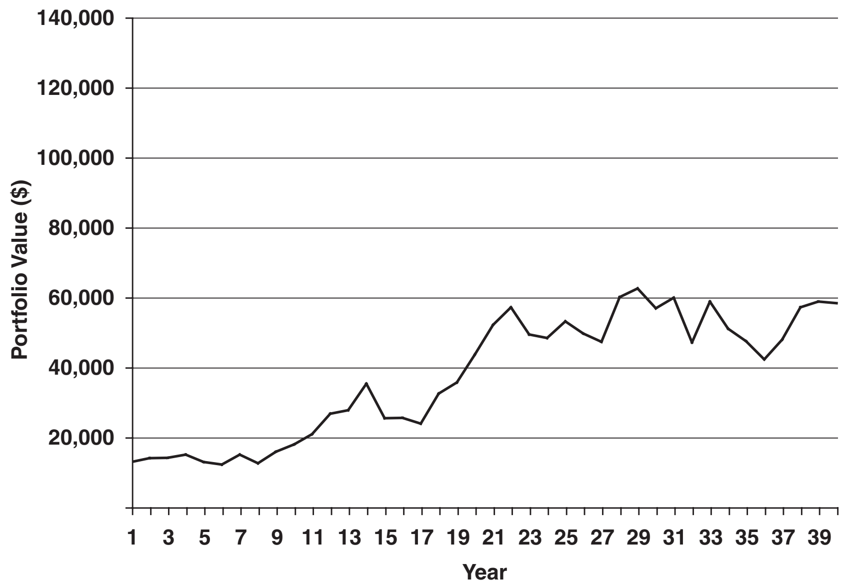

FIGURE 1.4 One Possible 40-year Scenario for an Equity Mutual Fund with a $10,000 Initial Investment.

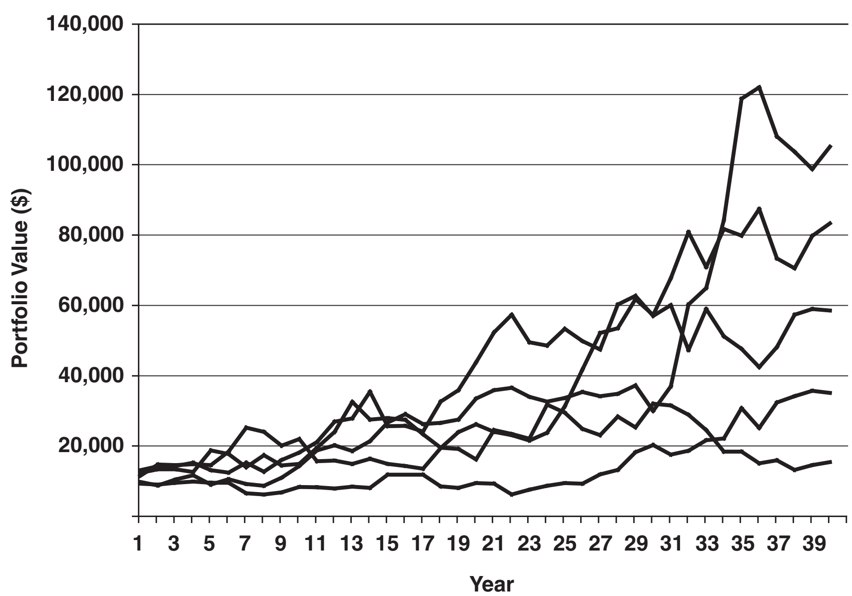

FIGURE 1.5 Five Possible 40-year Scenarios for an Equity Mutual Fund with a $10,000 Initial Investment.

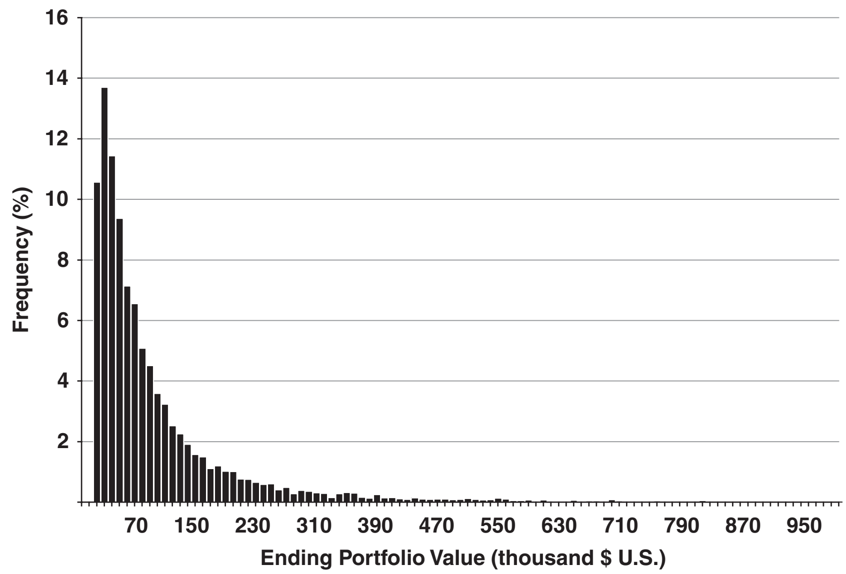

FIGURE 1.6 Histogram of 10,000 Simulations of an Initial $10,000 Investment in a Hypothetical Mutual Fund over 40 Years.

as $15,000 to as much as $105,000 depending on how the market and specific securities held by the fund performed over the 40-year period. If you have ever daydreamed how your life would have been different had you married your high school sweetheart, taken that job in New York years ago, or decided to become an artist in Paris, you have a feel for the simulation concept. With simulation, we let our investments live through many *possible* futures to see how things might play out. By generating thousands of such scenarios, we get a distribution of possible outcomes from which we can calculate many useful statistics. Figure 1.6 shows what the distribution of final wealth outcomes after 40 years looks like for our hypothetical mutual fund after 10,000 simulated scenarios.

The shape of this distribution of portfolio outcomes is very typical for investments with higher levels of volatility (like a stock mutual fund). The minimum outcome is bounded by zero (unless you borrow money to buy the fund, you can't lose more than you started with). About two-thirds of the outcomes lie between zero and $75,000. However, there are a significant number of observations that greatly exceed the expected value. This is called a *skewed distribution,* and is characteristic of portfolio distributions when holding volatile assets like equities over long time periods. Note that there are small probabilities of some very good outcomes. Even though 50 percent of the outcomes are below $46,700 (the *median* or middle value), the average outcome is over $89,000. The reason for the difference is that the very good outcomes drive up the average, yet they do not affect the median case very much since their probabilities are small. The higher the volatility of the asset, the more skewed the distribution will become, and the bigger the difference between the median value and the average value. You can play around with this type of investment simulation and see the results for different kinds of portfolios on the Financial Engines *Personal Online Advisor* service. This service allows you to input the investments in your portfolio and perform simulations of how they might perform over different time horizons. It can be a very enlightening exercise to see the range of possible investment outcomes.

The goal of the preceding analysis is not to create pretty charts, but to calculate statistics that are useful for making decisions. For instance, we can calculate the probability that our investments will exceed a certain target goal. We can also calculate how much we might lose if markets perform poorly over the period. We can calculate the range of possible outcomes, including the downside, median, and upside cases. Of course, knowledge of the distribution of outcomes is very valuable in making informed investment decisions.

In the following chapters, we will use the Financial Engines simulation engine to illustrate and quantify the tradeoffs that exist when developing an investment strategy.

Which Investments Are Best?: Investment Analysis

Another prominent technique in the institutional investor toolkit is the use of analytic methods to determine how specific investments are related to market returns, and how they have performed relative to what we would expect given the type of investment. When selecting investments for a portfolio, it is very useful to know what their expected risk and return characteristics are and how they have performed compared to other similar assets. Financial Engines has adopted these methods used by institutional investors to help individual investors do a better job of investment selection. One important technique for determining what an investment acts like is *investment style analysis*.

Investment style analysis, or *returns-based style analysis,* was developed by William F. Sharpe, a 1990 Nobel laureate and cofounder of Financial

Engines, for use in assessing and measuring the performance of investment managers for pension funds and other institutional money managers. The technique was first published in a paper by Sharpe in 1988.[7] Since then, the method has gained widespread use throughout the investing world due to its unique ability to look through a fund's underlying investments to understand how it behaves.

The key to the technique is to recognize that the performance of a fund is actually driven by a combination of exposures to different parts of the economy. For instance, a large-capitalization growth equity mutual fund may have exposures to growth stocks with large capitalizations. However, most funds have exposures to other parts of the market as well. Perhaps some of the stocks held by the fund may be from small or midsized companies, or from value-oriented stocks. If the fund is actively managed (as opposed to an index fund), it may keep a small percentage of assets in cash in order to better handle daily inflows and outflows. Style analysis of such a fund would yield exposures to large-cap growth equity, and perhaps various other equity exposures, or a small percentage in cash. The important point is that most funds are not easily categorized into a single asset class. In the real world, funds have several exposures and often behave significantly different from their stated investment objective. It's kind of like the difference between the name on the label and the list of ingredients on the back of the can (for instance, check out the ingredients listed on a can of Cheez Whiz).

Style analysis works by examining how a fund's returns move with various asset classes in the market. For example, some funds might move more closely with changes in the value of smaller company stocks, while other funds might be more related to movements in corporate bonds. Sharpe discovered that by setting up the problem a particular way, one could derive accurate estimates of a fund's investment exposures with a relatively small amount of data on historical returns. This opened up a whole new window of analysis of funds and their investment behavior. Instead of relying on a single designated *benchmark* like the S&P 500 Index, one could construct a custom benchmark for each fund that was representative of types of investments selected by their investment strategy. A manager who mostly invested in value stocks but held 20 percent of the portfolio in growth stocks could be measured against the performance of a benchmark based on 20 percent growth equity and 80 percent value equity. With the use of such custom benchmarks, it is possible to more accurately ascertain the extent an investment manager is actually adding value. The use of style analysis techniques among professional investors has greatly expanded in recent years due to its flexibility and utility.

Knowledge of a fund's investment style, as we shall see, is critical in building effective portfolios that maximize the expected return for a given level of risk.

How Do You Construct a Good Portfolio?: Optimization

A third major tool of the institutional investor is *portfolio optimization*. Optimization is a mathematical technique for figuring out the best solution to a particular problem, given a range of possible alternatives. In the context of investing, the purpose of optimization is to select investments that do the best job of meeting a set of objectives. These objectives may be fairly simple (give me the combination of funds with the highest expected return for a desired level of risk), or they may be very complex, with numerous competing goals and constraints.

In the institutional world of pension funds and university endowments, portfolio optimization is used to create alternative investment strategies. Data on each of the investment options, along with information about desired risk levels and other constraints, are fed into an optimization algorithm to yield a recommended portfolio allocation. The most popular forms of portfolio optimization are based on the concept of *mean-variance* portfolio theory. Very simply, this means an optimization where the goal is to maximize the expected return of the portfolio for a given level of variance (volatility). The optimizer takes information on each of the securities available to you and determines the portfolio that offers the highest expected return for the level of risk that you are willing to assume.

Of course, there are many ways to measure risk other than looking at just the volatility of returns. More complex optimization methods are able to accommodate different measures for risk, and the inclusion of various types of constraints. For instance, a given investment may have a minimum investment amount, or you may wish to impose a maximum percentage allocation to a particular security. Moreover, there may be factors other than expected return that you may wish to optimize. For instance, capital gains taxes can impact the desirability of buying or selling certain securities. There may also be preferences for the number of securities in the portfolio, or for concentration in a single security. Each of these constraints or preferences can add significant complexity to the optimization algorithm. Modern portfolio optimization engines are able to handle a wide variety of such real-world complexities. The bottom line is that optimization engines are able to help you build portfolios that achieve desired objectives while abiding by constraints and preferences that you believe are important. They are an invaluable tool for modern portfolio management, both for institutions and for individual investors.

TAKE-AWAYS

- The world has changed—and you had better get used to it. You *will* have personal responsibility for managing retirement assets in the future. There is no going back to the days where your employer or the government took care of your retirement income needs.
- Make a fundamental choice: either educate yourself enough to make informed decisions or admit that you need some help. If you choose to get help, make sure that it comes at a reasonable cost and from a competent and objective expert that you can trust. Be realistic about your ability and interest in managing your investments on your own. But most importantly, don't make the mistake of doing nothing.
- If you do work with an advisor, learn enough to spot obviously bad advice when you see it. Also, be careful of how your advisor gets paid—conflicts of interest can yield advice that is not in your best interest.
- Monte Carlo simulation is a powerful technique for understanding the range of possible investment outcomes you might experience. Remember that the expected outcome is only one possibility. It is more likely that you will experience an outcome that is different from the expected one.
- Leveraging tools from the institutional world provide important insights for making informed decisions. Try out the *Personal Online Advisor* service to see how your own investments might perform at www.financialengines.com/intelligentportfolio.[8]