1

Introduction: The Puzzle of the Missing Billionaires

Any fool can make a fortune; it takes a man of brains to hold onto it.
—Cornelius "Commodore" Vanderbilt

A beautiful statue of Cornelius Vanderbilt, the nineteenth-century rail and shipping tycoon, adorns the outside of Grand Central Station in New York City. It's there because "the Commodore" ordered the station's construction. Although partially obscured today by an eyesore

called the Park Avenue Viaduct, the statue sits right at the heart of midtown Manhattan, the global center of finance, regularly visible to many of today's financial titans.

When Vanderbilt died in 1877, he was the wealthiest man in the world. His eldest son, Billy, received an inheritance of one hundred million dollars—95% of Cornelius' fortune. Unfortunately, it came without even the most basic of instructions on how to invest and spend this wealth over time. Within 70 years of the Commodore's death, the family wealth was largely dissipated. Today, not one Vanderbilt descendant can trace his or her wealth to the vast fortune Cornelius bequeathed.† The Vanderbilt clan grew at a higher rate than the average American family, but even so this outcome was far from guaranteed. If the Vanderbilt heirs had invested their wealth in a boring but diversified portfolio of US companies, spent 2% of their wealth each year, and paid their taxes, *each one living today would still have a fortune of more than five billion dollars*.

What went wrong?

The Vanderbilt experience is noteworthy in scale but not in substance. The dissipation of great wealth over just a few generations is a common enough occurrence to warrant its own maxim: "from shirtsleeves to shirtsleeves in three generations." The truth of this dictum can be seen in how remarkably few of the billionaires in the news these days are the scions of old-money wealth. From these observations, we can tease out a valuable insight: the wealthiest families of the past were not equipped to consistently make sensible investing and spending decisions. If they had been, their billionaire descendants would vastly outnumber today's newly minted variety.

To get a rough count of these "missing billionaires," let's turn back the clock to 1900. At that time, the US census recorded about four thousand American millionaires, with the very richest counting their wealth in the hundreds of millions. If a family with five million dollars back then had invested their wealth in the US stock market and spent from it at a reasonable rate, that single family would have generated about 16

† Hopefully his descendants had a good time while dissipating their wealth. It brings to mind how George Best, the legendary Manchester United winger in the wild 1960s, answered a reporter's question asking how he went through all his money: "I spent a lot of money on booze, birds, and fast cars. The rest I just squandered."

billionaire households today.† If even a quarter of those millionaires in 1900 started with at least five million dollars, their descendants alone should include close to *16,000* old-money billionaires alive today. If we include the private wealth created throughout the twentieth century as well, rather than just a snapshot in 1900, we believe the tally of potential billionaires is vastly greater.

But as of 2022, *Forbes* estimated there were just over 700 billionaires in the United States, and you'll struggle to find a single one who traces his or her wealth back to a millionaire ancestor from 1900. We needn't go so far back in the past to find this pattern. Fewer than 10% of today's US billionaires are descended from members of the first *Forbes* 400 Rich List published in 1982. Even the least wealthy family on that 1982 list, with "just" $100 million, should have spawned four billionaire families today. We recognize that some wealthy families purposefully chose to give away or consume virtually all of their wealth in their lifetimes, but we believe these cases were relatively rare and do not account for the near-complete absence of "old-money" billionaires we see today.

We're not lamenting a scarcity of billionaires in the world today. Our point is that, collectively, *we all face a really big and pervasive problem when it comes to making good financial decisions*. If even the most financially successful members of our society, at least some of whom were smart and capable, and all of whom could afford the "best" financial advice, consistently made atrocious financial decisions, what can be expected from the rest of us? There's a Persian proverb Victor's father was fond of, which seems improbable at first, but the truth of which has become a main motivation for this book: *"It's more difficult to hold on to wealth than it is to make it."* This book sets forth our framework for addressing the challenge faced by all families—not only potential and actual billionaires—to find a path to better financial decisions.

In the chapters that follow, we're going to explore in detail the most common ways in which these missing billionaire families discarded their enormous head start: taking too much or too little risk, spending more than their wealth could support, and not adjusting their spending as their wealth fluctuated. Above all, they did not have a unified decision-making

† We are assuming the families spent and donated 2% of their wealth each year, paid taxes on what they spent, had no other income, and procreated at a typical rate. Full investment in the US stock market as the baseline assumption for these families is meant to be representative of the many attractive ways wealthy families could have invested. We are not suggesting that people could or should have invested that way in the past or should in the future.

framework, which left them susceptible to chasing whatever was hot and dumping it as soon as it was not, anchoring spending decisions on hoped-for portfolio returns, and paying exorbitant fees for poor advice.

Notice that among the primary errors we listed, we did not include choosing bad investments. That's because one thing that did not cause these billionaires to go missing was a poor investment environment. Indeed, it's hard to imagine a better one. The US stock market delivered a 10% pretax annual return over the period, turning one million dollars in 1900 into roughly *one hundred billion* dollars at the end of 2022, a 100,000x return. Instead, perhaps the heart of the problem is one of misplaced attention. In investing, the natural tendency is to focus on the question of *what* to buy or sell. Nearly 100% of the financial press is dedicated to this question, so it's reasonable to suppose that the *"what"* decision is the most important thing we should be thinking about. It isn't.

We'll explain that the most important financial decisions you need to get right are of the *"how much"* variety. How much should you buy of a good investment; how much should you spend today and over time; how much tax should you defer to the future; how much should you spend to insure against low probability, high consequence events. Implicit in these questions is the recognition that risk is present in just about every good thing we come across. So, whenever we're trying to figure out how much of a good but risky thing we should do, we need to weigh the greater expected benefit from doing more versus the cost of taking more risk. We hope to leave you with a practical framework for making these sizing and risk-taking decisions consistently and confidently.

Why do we think sizing is so important? Consider this: if you pick bad investments but do a good job sizing them, you should expect to lose money, but your losses won't be ruinous. You'll be able to regroup and invest another day. On the other hand, if you pick great investments but commit way too much to them, you can easily go broke from normal ups and downs while waiting for things to pan out.

Our own personal experience backs up the proposition that the sizing decision, often an afterthought, is actually the most critical part of investing. We've both experienced first-hand the impact of getting the "how much" decision wrong, losing the majority of our personal wealth in the process. Victor was a founding partner at the hedge fund Long-Term Capital Management (LTCM). In 1998, at age 36, he took a nine-figure hit when LTCM was undone by that decade's second financial crisis. The monetary loss was compounded by the psychological blow of the *business'* failure and its impact on the 153 employees who

worked there, as well as the impact on the reputations of all involved. A decade later, James at 28 lost a smaller sum but still a material fraction of his wealth in part through his investments in the hedge fund where he worked.

In each case, we believed we'd selected investments with an attractive risk/reward profile that were highly likely to pay off in the long run. The trades that took down LTCM in 1998 were money-makers over the ensuing years, and the hedge fund that employed James also bounced back to generate strong returns following its precipitous swoon in 2008. Unfortunately, the short run always comes before the long run, and neither of us got to enjoy the rebound of these investments. The lesson: good investments plus bad sizing can result in cataclysmic losses.

The LTCM story has been told countless times—several books, many articles, and even a Harvard Business School case study. It was colorful, involving Wall Street traders straight off the pages of Michael Lewis' *Liar's Poker* together with a band of highly respected professors, including two Nobel Laureates. To the outside world, it appeared that they had built a money machine (one of the books about LTCM was rather hyperbolically titled *Inventing Money*), so it's not surprising that most of the LTCM partners were heavily invested in the fund they managed. The fund had returned more than 40% per year from inception to shortly before its decline, with an annual volatility of about 12%. And prior to founding LTCM, the older partners could look back on their 20 years of very positive experience doing the same kind of investing at Salomon Brothers.

Victor had about 80% of his family's liquid wealth invested in the LTCM hedge fund. With the benefit of hindsight, this level of investment concentration was a mistake. But the more useful question to ask, which hasn't been addressed anywhere in the LTCM literature, is this: what analysis should Victor have done to determine how much to invest in the fund, and what would that analysis have said? We'll provide a detailed answer, and we'll also show why "as much as you can get" is almost never the correct answer to the question "how much of a good thing is right for you?"

Is this poor sizing judgment confined to the likes of your authors when they were young, highly confident bond arbitrage traders, and to wealthy families living in the first half of the twentieth century? Unfortunately, there's quite a bit of evidence that this failing is more widespread and persistent, suggesting that investors are systematically hurt

by poor money management skills. For example, individual investors in aggregate severely underperform market returns, in both absolute and risk-adjusted terms. Some of this underperformance comes from paying high fees, but much arises from having too much or too little at risk and usually at just the wrong times. A landmark study of individual brokerage accounts by University of California Professors Brad Barber and Terrance Odean, aptly titled, "Trading Can Be Hazardous to Your Wealth," found that individuals who actively traded their portfolios underperformed market returns by 6% per annum. Other researchers have found that the aggregate returns that investors in mutual funds experience are typically several percentage points per year below the returns that a buy and hold investor in those same funds would have earned.[1]

In an attempt to better understand how people deal with the sizing of attractive investment opportunities, Victor and his co-researcher Richard Dewey conducted an experiment in 2013 that was later published in the *Journal of Portfolio Management*. In a controlled setting, they invited 61 financially and quantitatively trained individuals to play a simple coin-flipping game. The participants were informed that the (virtual) coin was programmed to have a 60% chance of landing on heads. They were each given $25 and allowed to bet any way they wanted. They were told that at the end of 30 minutes—time enough for about 300 flips—they'd be paid however much the $25 had turned into, subject to a maximum payout of $250.

As we'll discuss in more detail throughout the book, there are a range of simple, sensible betting strategies that would result in an expected payout in the vicinity of the $250 cap.† So, how did the players do? More than 30% of them lost money, and incredibly, over 25% went bust! Only 20% made it to $250, a far cry from the more than 90% who should have. Our thesis in this book is that the ideas leading to doing well in our coin-flipping game are equally helpful in making good decisions on the important financial matters in our lives.

Investing isn't the only area where making sound financial decisions under uncertainty is important. Everyone needs a coherent saving and spending plan for all stages of life, especially if they expect to enjoy periods of retirement. Even when it comes to the seemingly simple decision of at what age to start taking Social Security benefits, a recent study

†Not a bad return for 30 minutes of quickly pecking at a keyboard, especially for a bunch of university students!

found that Americans are leaving more than a trillion dollars on the table by making suboptimal decisions, primarily by taking the benefits sooner than they should.[2] There's also good evidence that many people spend too much on certain kinds of insurance, from protection on their kitchen appliances to comprehensive car insurance, while spending too little on other types of insurance products, such as annuities. Sadly, personal finance books just don't provide enough good advice either, as documented in a recent easy-to-read survey of the 50 most popular ones by Yale Professor James Choi.[3]

The decision-making problems we focus on have three common features:

1. They require a decision to be made in the face of uncertain outcomes.
2. Some of the outcomes will have a meaningful impact on your happiness or welfare.
3. The impact on your welfare will not exactly mirror the monetary outcomes.

Philosophers, economists, and mathematicians have been thinking about how to systematically evaluate these kinds of decisions for almost 300 years, since Daniel Bernoulli introduced the concept of a "utility function" in order to provisionally solve the St. Petersburg Paradox, which we describe in Chapter 6. It's virtually a law of human nature that we experience a diminishing marginal benefit from further and further increases in spending (or wealth). This is why the utility we derive from spending our wealth doesn't mirror the monetary sums involved. Thinking this way makes it clear that we should make decisions that maximize our utility rather than purely monetary outcomes. These ideas are at the root of economic thought and permeate even through the younger branches of the economics tree, such as behavioral finance.

The utility-based decision framework has at its core three main steps for any given financial decision. First, we need to assess possible monetary outcomes and estimate their associated probabilities. Then we need to map these monetary outcomes into utility outcomes. Finally, we need to search over the range of different possible decisions, to find the one which produces the highest Expected Utility.

The framework of maximizing Expected Utility is the foundation of the field known as "decision-making under uncertainty" as well as most other theories of individual choice and human interaction. Many of the

ideas that form the basis of this book are found in the sub-field known as "lifetime consumption and portfolio choice," an area first formulated and explored in the 1960s by Nobel Prize-winning economists Paul Samuelson and Robert C. Merton. The basic problem addressed is to determine the choices that a person should make with regard to investing, saving, and spending that will maximize Expected Lifetime Utility.

Unfortunately, many of the original papers in this field have been too complex or abstract to gain traction in practice, with the result that very few people are using these tools to make better decisions. In the pages that follow, we hope to present these important ideas in a way that will not only resonate, but also change your thinking. It can take a surprisingly long time for a great idea to make the jump from theory to practice; we believe these are great ideas whose time has come.

The framework of making decisions to maximize Expected Utility has been criticized for not being a realistic representation of how people actually behave. Indeed, researchers have documented myriad cognitive biases that result in inconsistent and irrational decision-making. We agree with this evidence wholeheartedly, and that's exactly why we've decided to write this book—in order to help people make better decisions by understanding what happened to all those missing billionaires and overcoming some of those cognitive biases! No framework would be needed if people were *naturally* excellent at making decisions involving a complex web of probabilities, scenarios, and transformations of monetary sums to personal utility. But we know that we aren't.

The fact that the framework and actual behavior don't always agree is a sign that the framework can add value. People will undoubtedly still choose to make seemingly suboptimal decisions in some ways—that's part of being human—but better to do so knowingly and with an ability to estimate the cost rather than out of ignorance or for lack of a better alternative. In this book, we are not going to describe how people *do* make important financial decisions; our goal is to explain how they *should* be making them. We hope to change behavior by convincing our readers to adopt a framework centered on Expected Utility for making the important financial decisions in their lives.

The Expected Utility you get from a particular course of action in an uncertain world is a model—it is a simplified representation of reality. Like any model, it is a framework for organizing and simplifying a complex situation. Sometimes we get so familiar with a model that we don't even realize that's what it is. For example, in the world of government

bonds, the concept of yield to maturity is a model that really came into its own when computer power became cheap enough to quickly make the relevant calculations for any bond, on any settlement date, at any price.[4] Without the model of yield to maturity, it would be very hard to compare bond A, with a price of 102.15 maturing in 4 years and 11 months paying a coupon of 3%, to bond B, with a price of 125.67 maturing in 5 years and 2 months paying a coupon of 8%.† The "yield model" just converts price and other relevant information into a better metric to compare bonds. Expected Utility, just like yield to maturity, is a model. And it's a good model in that it allows us to compare and evaluate different and complex decisions in a way that is intuitive and has an underlying economic rationale.

One of the first steps on the path to better decision-making will be to understand your own personal utility function. This is not as difficult as you may think. In a 2018 survey we conducted with about 30 of our clients—if you haven't guessed, we run a wealth advisory practice—we found that all of them were able to express a relationship between wealth and utility that helped them weigh different risk/reward trade-offs in a consistent manner. Our goal is to increase your comfort with this framework by working through a wide variety of practical case studies, from deciding how much of your wealth you should invest in your own business to when it makes sense to pay capital gains tax. We'll dive deeply into the primary ways that risk impacts your financial well-being. Our aim is to give you core concepts, without dumbing them down to the point where they become hidden behind vague general guidance. There will be some math involved, but nothing beyond ideas you'd have met in high school—and we'll always try to explain the economic logic behind an idea, not just the math.

We recognize that a big part of your financial decisions still involves figuring out *what* to invest in, so we won't ignore the "what" question altogether. For many investors today, the potential investment universe is daunting: index funds, active funds, individual stocks and bonds, real estate, hedge funds, private equity, venture capital, and the list goes on and on. Plenty of books have been published to help you evaluate each of these types of investments, so we'll only dedicate one chapter to discussing how to assess the expected return and risk of the most common

† For readers who tried to figure it out in their heads, the yield to maturity of bonds A and B are 2.53% and 2.65%, respectively.

investments. Most of this book explains how to translate your evaluation of individual investments into a decision on how much risk to take, on each investment and in total, and how to make the connected decision of how to spend your wealth over time.

Most of the methods we discuss will be usable by anyone with a calculator or a simple spreadsheet. Some may require more sophisticated technology to implement them exactly as we do, but with minor simplifications these too can nearly always be worked out on a spreadsheet. We've built a set of modeling tools to help work through some of the most common and important problems that benefit from them, and they're available on our website, www.elmwealth.com. If you find yourself wanting to use a piece of the framework but can't find a good tool for it, please do feel free to engage with us, and we can very likely help or point you in the right direction.

We've broken the book into four main segments:

- Investment Sizing
- Lifetime Spending and Investing
- Where the Rubber Meets the Road
- Puzzles

We know that for most readers, this book won't make for light and relaxing reading. One suggestion we have to make things a bit easier going, if desired, is to begin reading just a subset of chapters that cover core concepts, skipping over those that primarily discuss illustrations and applications of those ideas. We hope you'll find the skipped chapters calling your name later. We think the following 10 chapters are important to read on your first pass through the book: 2–6, 9, 12–14, and 24. In whatever order you get to them, we hope you'll find the other 14 chapters informative and helpful in bringing the core ideas of the book to life.

There are places where we'll describe key concepts with symbols and formulas as well as words. Have no fear; there are only four formulas in the entire book, which you can find for reference in the section "A Few Rules of Thumb" near the end. One of the four formulas, which we call the Merton share, is simple and important enough that we hope you'll finish the book with it indelibly recorded in your mind! We hope that in these few instances, stating ideas both in natural language and the language of mathematics will make these concepts clearer and more memorable.

We've written this book with our younger selves never far from our minds. We hope that, if we had read it early in our careers, it would have helped us to avoid the investment sizing mistakes we made. Indeed, the mission of Elm Wealth, our investment management and wealth advisory business, is to help our clients benefit from our experience in the broad area of financial decision-making under uncertainty. Returning to our coin-flip experiment, our optimism that this book can do good stems in part from the fact that once we explained a more sensible sizing strategy to the participants, they took that on board and would mostly have reached the maximum payout if given another chance to play. If people are exposed to a good framework, they will often readily understand and embrace it.

Making better decisions is not a zero-sum game. Sound financial decision-making, consistent with your individual preferences, will not only increase your and your family's expected happiness, but will dramatically increase the welfare of our entire society.

An investment in knowledge pays the best interest.

—Benjamin Franklin