Don’t Run Out

A Practical Guide to Tax-Smart Retirement Income, Social Security, and Medicare, revised for 2026

Jeff Batson, RFC®

Copyright © 2026 Jeff Batson

All rights reserved.

No part of this publication may be reproduced, distributed, or transmitted in any form or by any means, including photocopying, recording, or other electronic or mechanical methods, without the prior written permission of the author, except in the case of brief quotations embodied in critical reviews and certain other noncommercial uses permitted by copyright law.

Revised for 2026.

goatplanning.com

Contents

Introduction: Why I Wrote This Book.....5

Cast of Characters.....9

Chapter 1: The Real Risk Isn't the Market.....14

Chapter 2: Sequence of Returns — The Risk No One Warns About.....19

Chapter 3: The Volatility Buffer.....30

Chapter 4: Social Security 101.....41

Chapter 5: Claiming Strategies — When to Start and Why It's (Mostly) Irreversible.....54

Chapter 6: Medicare — Four Parts, Two Paths, and One Hard Deadline.....79

Chapter 7: Part D — Drug Coverage, the New Ceiling, and the Mistakes That Cost the Most.....94

Chapter 8: Beyond the Premium — IRMAA, Long-Term Care, and Medicare's Gaps.....106

Chapter 9: Required Minimum Distributions — The IRS Wants Its Money.....120

Chapter 10: Roth Conversions — Paying Taxes on Your Terms.....132

Chapter 11: The Case for a Bond Alternative.....152

Chapter 12: How a Fixed Indexed Annuity Actually Works.....162

Chapter 13: Indexed Universal Life — Tax-Free Income and Legacy Planning.....178

Chapter 14: Tax Mitigation in Retirement — Building a Multi-Year Plan.....193

Chapter 15: Generational Wealth Transfer.....207

Chapter 16: Putting It All Together.....220
Important Notice.....232
Glossary.....234
Sources and Where to Verify the Numbers.....239
About the Author.....241

Introduction: Why I Wrote This Book

Benjamin Franklin said the only two certainties in life are death and taxes. The good news is you only have to die once.

Taxes, on the other hand, show up every year, and in retirement they show up in places most people never expect: a Medicare premium that suddenly jumps for no obvious reason, an IRA withdrawal that pushes you into a higher bracket, an inheritance that comes with a ten-year clock attached. I wrote this book because I kept watching smart, careful people get surprised by these things, not because they hadn't saved enough, but because no one had ever walked them through how the rules actually work once the paychecks stop and the withdrawals begin.

I'm Goat Planning's founder, based in Staunton, Virginia, and for the last several years I've sat across the table from people in exactly that position. I'm a Registered Financial Consultant®, licensed in life and health insurance, and, through additional state-required coursework and exams, authorized to offer non-variable annuities, long-term care insurance, Medicare Supplements, Medicare Part D drug plans, and Medicare Advantage plans. I'm also pursuing my Social Security certification through the National Social Security Association. Consider that last sentence my one permitted humble brag, and now you

know exactly where I'm coming from before you read another page.

Something most books like this skip: some of the readers of this book will go on to work with me directly, including on products discussed in these chapters, such as annuities and certain types of life insurance. I think you deserve to know that up front, plainly, and only once. What I won't do is use these pages to make a pitch. Instead, I'll explain, as clearly and honestly as I can, what each tool is actually for, who it tends to help, and who it doesn't, the same way I'd want a professional to explain it to my own parents. If something in this book ever sounds like a pitch, I haven't done my job. And if you do decide you'd like personalized help, know that I'm legally required to act appropriately for your specific situation under whatever rules apply to the product or service we're working on, and personally I try to hold myself to that standard whether or not the law requires it. You'll find a way to reach me in the back of the book. Nowhere else.

That's the spirit this whole book is written in: education first. Not because it sounds good on a back cover, but because I've found it's the only approach that actually helps. The retirement income decisions in this book (when to claim Social Security, which Medicare path to choose, how to manage taxes across a multi-year withdrawal plan) are decisions you'll likely

make once and live with for decades. They deserve more than a sales pitch. They deserve an explanation.

So that's why I wrote this book: because the distribution phase of retirement, turning what you've saved into income that lasts, is at least as complicated as the saving phase that came before it, and almost nobody prepares you for it. The accumulation years come with decades of practice: open a 401(k), pick some funds, watch the number grow. Time is on your side: a bad year can be outlasted, a downturn can be waited out, mistakes have years to correct themselves. The distribution years don't offer that luxury. Time is no longer on your side; if anything, it works against you. The decisions compound, the deadlines are often final, and the cost of learning by trial and error can be measured in thousands of real dollars, or in premiums, taxes, and benefits you can never get back.

This book won't try to make you a financial professional. It will try to make you a confident, informed client: someone who can sit across from any advisor, agent, or tax preparer and ask the right questions, recognize a good answer, and make decisions that fit your own life rather than someone else's sales quota.

Throughout the book, you'll meet a handful of entirely fictional people I've invented to make these ideas concrete: Dave and Mary, a couple facing the everyday version of these decisions; Linda, planning on her own;

the Whitfields, with a bit more complexity to manage; and Walter, further down the road than the rest, showing what these decisions look like once they've already been made. None of them are real, and none are based on any specific client, but the situations they face are, in some form, ones I see almost every week. A full introduction to each of them follows this Introduction, so you'll have a page to flip back to if you ever lose track of who's who.

I'll leave you with one more thing before we begin. Wherever you are right now (still working, newly retired, or years into it), the goal of this book is the same for every reader: to help you build a plan that lets you stop worrying about whether your money will last, so you can get back to actually living the retirement you saved for.

Let's get started.

Cast of Characters

Throughout this book, the same six people return. They're entirely fictional, invented to make abstract rules concrete, but the situations they face are drawn from the kinds of conversations financial advisors have every week. This page is here to flip back to whenever it's hard to keep track of who's who.

Dave & Mary Thompson, 64 and 61

Dave and Mary are the book's anchor couple, and they appear more than anyone else. Dave still picks up part-time consulting work; Mary retired from teaching last year and is still getting used to the rhythm of it. They've done most things right (they saved, they paid off the house, they raised two kids who turned out fine), and now they're facing the same question almost every couple their age eventually asks: how to turn what they've built into income that lasts? Dave is also five years into managing an IRA he inherited from his father, who passed away at 79, already several years into his own required withdrawals, which means Dave has been taking annual distributions of his own, with five years left before the account must be fully emptied. Mary also built up a health savings account during her working years, a detail that matters once Medicare enters the picture.

Traditional IRA: $100,000 (Dave) · Roth IRA: $150,000 (Mary) · Mary's HSA: $28,000 · Mary's pension · Home in

Staunton, VA, paid off · Children: Kevin and Allison · Dave's $60,000 inherited IRA: father passed 5 years ago at 79; year 5 of the 10-year distribution window.

Linda Castellano, 58

Linda spent thirty years in healthcare administration and built a solid 401(k) along the way: a mix of traditional contributions, a smaller Roth portion she started later in her career, and years of employer match, which by law always lands in the traditional, not-yet-taxed bucket. Divorced, with one grown daughter, Maya, she's planning retirement the way she's done most things: on her own. Linda's chapters look different from the Thompsons': there's no spouse to coordinate Social Security with, no second income to fall back on, and a sharper edge to the question of what happens if it's just her for the next thirty years. Linda also carries an inherited IRA with a more complicated history. Her maternal grandmother passed away eight years ago, already taking her own required distributions, and Linda's mother inherited the account as a non-spouse heir, which started a 10-year distribution clock with annual withdrawals required along the way. When her mother passed away two years ago, Linda inherited the same account as what's

known as a successor beneficiary, which means she didn't get a fresh 10-year window of her own. She's now racing the final two years of her mother's original clock, a useful and often misunderstood example later in the book.

401(k): $420,000 total: $340,000 traditional/match, $80,000 Roth · Taxable brokerage: $35,000 · Condo, small mortgage remaining · Daughter: Maya · Maternal grandmother passed 8 years ago (already taking RMDs); Linda's mother inherited as a non-spouse original beneficiary and passed 2 years ago; Linda is a successor beneficiary with 2 years left on the original 10-year window.

Robert & Susan Whitfield, 67 and 65

Robert sold the small business he spent two decades building; Susan retired from a corporate VP role not long after. Between the two of them, they've accumulated more than most (a sizable IRA, a taxable brokerage account, a vacation home on the Outer Banks), which means they also have more decisions to make, and more at stake if they get them wrong. The Whitfields show up most often in the later chapters, where Roth conversions, annuities, and legacy planning actually have room to matter.

Traditional IRA: $1.2 million · Taxable brokerage: $400,000 · Susan's former 401(k): $200,000 · Primary home and a paid-off vacation property · Three adult children.

Walter Higgins, 84

Walter is further down the road than the rest of the cast, and that's exactly the point. A retired postal worker and steady saver his whole life, he retired in 2008 at 66, right as the financial crisis hit, and watched the downturn do enough damage to his portfolio that he went back to work for a few more years to recover. Once he was working again, he suspended his Social Security benefit rather than keep collecting it, letting it grow through delayed retirement credits until he restarted it at 69, at a noticeably higher monthly amount than what he'd locked in at 66. He's been taking required minimum distributions for years now, not reading about them, living them. His wife, Carol, passed three years ago; his daughter Patricia has started helping manage the details. Walter is also Susan Whitfield's father, which means his story and the Whitfields' will eventually intersect, a connection that pays off later in the book.

Traditional IRA: $340,000, in RMD distribution · Social Security: claimed at 66 (2008), suspended during return to work, restarted at 69 with delayed retirement credits · Survivor pension from Carol's former employer · On a Medicare Advantage plan since 65; now managing a heart condition · Home paid off · Daughter Patricia holds financial power of attorney · Father of Susan Whitfield.

Chapter 1: The Real Risk Isn’t the Market

Who This Book Is For

Most people approaching retirement already suspect they’re missing something. They’re usually right.

Maybe it’s a Social Security question no one can quite give a straight answer to. Maybe it’s a feeling that the tax bill in retirement should be lower than it is, if only it were clear which levers to pull. Maybe it’s just a number that won’t stop running in the back of the mind at two in the morning: how long will this last? Whatever it is, nobody opens a book like this because they think everything’s fine. They open it because some part of them suspects it might not be, and they’d rather find out now, from a chapter, than later, from a bank statement.

The plain truth is this: most people approaching retirement are missing at least one piece of this puzzle, not because they’re careless but because nobody ever sat them down and explained how Social Security, Medicare, and taxes actually interact once the paychecks stop. And one part might come as a surprise. Some of what this book covers, including how Social Security claiming actually works, isn’t something a financial professional gets paid to explain. It’s covered here because it matters, not because it benefits anyone explaining it.

Dave and Mary Thompson appear throughout this book. Dave is sixty-four and still picks up consulting work; Mary is sixty-one and retired from teaching last year. They've done a lot of things right: saved steadily, paid off the house, raised two kids who turned out fine. Yet they still lie awake some nights wondering what they haven't thought of yet. That unspoken worry, common to almost everyone approaching retirement, is exactly what this book is for.

Four Fears, One Book

People bring specific questions to retirement planning: when to take Social Security, which Medicare plan is right, what to do about an old 401(k). Underneath them are usually the same four fears, in one form or another.

First, the fear of running out of money. Second, the fear of paying more in taxes than necessary. Third, the fear of getting a Social Security or Medicare decision wrong, and not finding out until it's too late to fix it. And fourth, the fear of a market downturn arriving at exactly the wrong moment and taking a bite out of savings there's no time left to rebuild.

This book is built to address all four, directly, one at a time. Not with vague reassurance, but with the actual mechanics: how Social Security claiming works, what Medicare actually costs and when, how required

minimum distributions and Roth conversions interact with a tax bill, and how to structure income so a bad market year doesn't become a bad decade. By the end, all four of these fears should get smaller, not because the risks disappear, but because they become understandable enough to plan around.

Two Very Different Jobs

The part almost nobody explains clearly is this: saving for retirement and living off retirement are two completely different jobs, and most people only ever get trained for the first one.

During the working years, the accumulation phase, time is on the saver's side. Contributions go into a 401(k), funds get picked, and decades do the heavy lifting. A bad year in the market is an inconvenience, not a crisis, because there are years, even decades, to recover from it. Markets have always eventually recovered, and the luxury of waiting them out, while still contributing, has always been there. A downturn during these years even has a silver lining: new contributions buy shares at a discount, which means a falling market is working in favor of anyone still adding money to the pile.

The distribution phase, the years of withdrawing income instead of adding to it, plays by entirely different rules. Time stops being an ally. A market downturn that hits right as withdrawals begin doesn't just dent a balance; it can permanently shrink how

long the money lasts, in a way that the exact same downturn wouldn't have if it had happened five years earlier or five years later. There's a name for this: sequence of returns risk, and its damage can be measured in real numbers. For now, the point is simpler. The strategy that served a saver well for thirty years is not automatically the right strategy for the next thirty years of spending. Retirement isn't a finish line. It's a job change, and nobody hands out the new job description.

Why This Matters More Than Most People Think

The uncomfortable math is this: someone who retires at sixty-five has a real chance of needing this money to last thirty years or more. That's not a pessimistic guess. That's just what longer lifespans look like now. Thirty years is longer than most marriages, longer than most careers, longer than the time most people spend in their childhood home. Planning to make money last that long isn't paranoia. It's just math.

And unlike the accumulation years, where a mistake in year three can be corrected in year fifteen, distribution-phase mistakes tend to be loud and permanent. Claiming Social Security at the wrong time can lock in a smaller check for life. Missing a Medicare enrollment deadline can attach a penalty that follows a person for decades. Taking too much from a traditional IRA in the wrong year can trigger a Medicare premium

surcharge nobody saw coming, on top of the tax bill they did expect. None of these mistakes are dramatic in the moment. They're just expensive, quietly, for a long time.

What This Book Will and Won't Do

This book will not turn its reader into a financial professional, and it won't try to. What it will do is make that reader a far more confident, informed client: someone who can sit across from any advisor, agent, or tax preparer, understand what's actually being recommended and why, and ask the questions that matter. That's a different goal than most books in this space set for themselves, and arguably a better one.

> These ideas play out through Dave and Mary, and through four other people introduced on the Cast of Characters page: Linda, making these decisions on her own; Robert and Susan Whitfield, with more complexity and more at stake; and Walter, further down the road, living out decisions the others have yet to make. None of them are real. The decisions they face are.

One risk catches more people off guard than any other, not because it's complicated, but because few explain it before it matters.

Chapter 2: Sequence of Returns — The Risk No One Warns About

Walter's Story

In 2008, Walter Higgins did everything according to plan. He'd spent a career with the postal service, saved steadily, and at sixty-six, just past his full retirement age of sixty-five and ten months, he claimed Social Security and settled into retirement. Then the bottom fell out. The S&P 500 lost 37 percent that year, the worst calendar-year decline since the Great Depression. Walter's portfolio, like millions of others, took a hit it had no time to recover from before he needed to start drawing on it.

Walter didn't wait to see what would happen. He went back to work immediately. Continuing to withdraw from a portfolio that had just lost more than a third of its value would have meant locking in those losses for good: selling shares at the bottom, permanently shrinking the base that needed to recover, and risking running out of money years earlier than planned. Going back to work wasn't about earning extra income to rebuild his nest egg. It was about removing the need to touch it at all, so it could simply wait out the downturn until the market recovered. While he was working, he also suspended his Social

Security benefit rather than keep collecting it. That was not the point of the plan, but a welcome bonus: it let the benefit grow through delayed retirement credits instead of staying frozen at the amount he'd locked in at sixty-six. By the time he restarted it at sixty-nine, it came with roughly three additional years of higher income behind it, plus a meaningfully larger monthly check for the rest of his life.

The numbers tell part of the story. What they don't capture is the trip Walter had planned for that first year of retirement: flying out to spend real time with his daughter Susan's three kids while they were still young enough to remember it. That trip didn't happen. Neither did the next few. By the time Walter's finances were steady enough to make the visit, his grandchildren were teenagers with their own lives, not the kids he'd pictured teaching to fish. Walter's story has a recovery in it, financially. Not everyone's does. But even his recovery came with a cost that never shows up on a balance sheet, the kind of cost to keep in mind every time the math turns to percentages and portfolios. The real stakes are never really about the numbers. They're about what the numbers were supposed to buy.

That risk has a name: sequence of returns risk. It's one of the least understood ideas in retirement planning,

and one of the most important, because unlike most investment risk, it isn't really about how the market performs. It's about when.

Same Average Return, Two Very Different Outcomes

Consider a hypothetical that makes the idea concrete. Picture two retirees, each starting with exactly $2,000,000, each withdrawing exactly $200,000 a year for fifteen years. Over those fifteen years, both portfolios experienced the exact same set of annual returns, averaging 11.31 percent a year, a perfectly respectable result. The only difference between them is the order those returns arrive in.

In the first scenario, the strong years come early and the weak years come late.

Scenario A — Strong Returns Early, Weak Returns Late

Year	Start of Year	Return	Withdrawal	Net Spendable
1	$2,000,000	5.78%	$200,000	$200,000
2	$1,904,040	18.59%	$200,000	$200,000
3	$2,020,821	31.16%	$200,000	$200,000
4	$2,388,189	6.08%	$200,000	$200,000
5	$2,321,231	27.04%	$200,000	$200,000
6	$2,694,902	15.55%	$200,000	$200,000
7	$2,882,983	(4.64%)	$200,000	$200,000
8	$2,558,492	31.65%	$200,000	$200,000
9	$3,104,955	18.53%	$200,000	$200,000
10	$3,443,243	6.59%	$200,000	$200,000
11	$3,456,973	(7.03%)	$200,000	$200,000
12	$3,028,008	23.72%	$200,000	$200,000

13	$3,498,811	36.90%	$200,000	$200,000
14	$4,516,072	(25.95%)	$200,000	$200,000
15	$3,196,051	(14.42%)	$200,000	$200,000
TOTAL	**$2,564,021**	**11.31% avg**	**$3,000,000**	**$3,000,000**

After fifteen years and $3,000,000 in total withdrawals, this portfolio still has $2,564,021 left in it. Now look at the same fifteen annual returns, in the opposite order: weak years first, strong years last.

Scenario B — Weak Returns Early, Strong Returns Late

Year	Start of Year	Return	Withdrawal	Net Spendable
1	$2,000,000	(14.42%)	$200,000	$200,000
2	$1,540,483	(25.95%)	$200,000	$200,000
3	$992,686	36.90%	$200,000	$200,000
4	$1,085,202	23.72%	$200,000	$200,000
5	$1,095,142	(7.03%)	$200,000	$200,000
6	$832,251	6.59%	$200,000	$200,000
7	$673,935	18.53%	$200,000	$200,000
8	$561,760	31.65%	$200,000	$200,000
9	$476,246	(4.64%)	$200,000	$200,000
10	$263,427	15.55%	$200,000	$200,000
11	$73,290	27.04%	$73,290	$73,290
12	$0	6.08%	$0	$0
13	$0	31.16%	$0	$0
14	$0	18.59%	$0	$0
15	$0	5.78%	$0	$0
TOTAL	**$0**	**11.31% avg**	**$2,073,290**	**$2,073,290**

Same starting balance. Same fifteen annual returns. Same average: 11.31 percent. The only thing that changed was the order. And by year eleven, this portfolio is empty, unable to fully fund even that year’s withdrawal, let alone the four years after it. The retiree in Scenario A finishes with two and a half

million dollars still working for them. The retiree in Scenario B runs out five years before the table even ends.

This is the part to remember at the next backyard barbecue. Ask either of these two retirees, “Is your financial planner any good?” and they could both truthfully answer, “They’re great. I’ve averaged 11.31 percent a year for fifteen years.” One of them is describing the best financial decade and a half of their life. The other is describing how they ran out of money. Average return never lies. It’s just rarely the whole truth, and it’s an easy number for an advisor, or an investor bragging at a barbecue, to hide behind.

That same blind spot shows up in a more sobering way in how investors actually behave, not just how a hypothetical portfolio could behave. Every year, DALBAR’s Quantitative Analysis of Investor Behavior compares the return the market actually delivered to the return the average investor actually captured. Over the twenty years ending in 2024, the S&P 500 returned 10.35 percent annualized. The average equity fund investor, over that same twenty years, earned 9.24 percent, a real and persistent gap, and one DALBAR has documented almost every year since it began tracking this in 1985. It used to be considerably worse: over the thirty years ending in 2015, DALBAR found the S&P 500 had returned 10.35 percent annualized while the average equity mutual fund investor captured just 3.66 percent, a gap of more

than six and a half percentage points a year, year after year, for three decades. The reason isn't fees, and it isn't bad luck. It's behavior: investors tend to sell during downturns, out of fear, and buy back in once the market has already recovered, out of confidence. They miss the worst days and the best days for the same emotional reason, and the math of missing the best days is brutal, since those recovery days are often what makes a portfolio whole again.

There is a deeper reason that 11.31 percent figure is almost beside the point. A CFP® doesn't need any particular skill to get a client roughly the S&P 500's return: buying a low-cost index fund and leaving it alone accomplishes that with almost no effort at all. That part isn't where real value gets added or lost. What actually determines whether a retiree ends up like Scenario A or Scenario B has nothing to do with which fund was chosen and everything to do with what happens to that same return once withdrawals begin, whether sequence of returns risk gets managed, ignored, or simply never explained. A forced sequence of returns problem (a downturn that happens to land in the first few years of withdrawals) and an unforced one (panic-selling during that downturn, locking in losses that a calmer year-eleven version of the same investor never would have) end the same way: a portfolio that can't recover, regardless of how good the index's long-term average turns out to be.

This deserves dwelling on, because the accumulation-phase planning industry, including plenty of well-credentialed advisors, has historically spent far more energy on picking investments than on managing the behavior of the person holding them, or on what happens the moment that person starts withdrawing instead of contributing. A CFP® who built a perfectly reasonable portfolio can still preside over a client who panic-sells in a downturn and buys back in too late, and "I averaged 11.31 percent" doesn't tell anyone whether that happened. Sequence of returns risk and investor behavior are two different mechanisms, but they share the same lesson: the number on the statement rarely tells the whole story, and retirement planning that stops at picking good investments is only half the job.

This is a hypothetical example built to isolate one variable, the order of returns, while holding everything else constant. Real portfolios don't move in such a clean pattern. But the principle it illustrates is real, well documented, and exactly what happened to Walter Higgins in 2008: the market doesn't care whether a downturn arrives the year someone retires or twenty years into their working life. For a saver, it's a rough patch. For a retiree pulling out a fixed dollar amount every year, an early downturn forces the sale of more shares at depressed prices just to generate the same income, shares that are then gone and can't participate when the market eventually recovers.

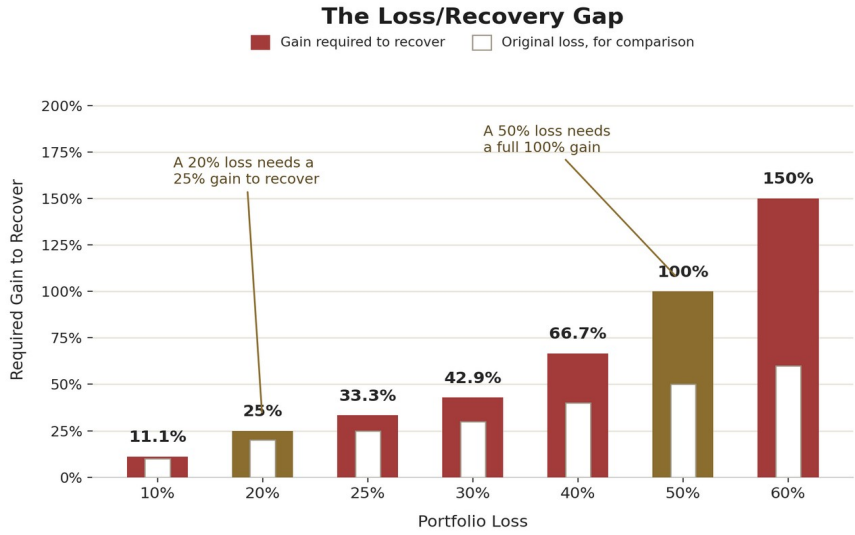
Why the Order Matters More Than the Average

Start with a piece of arithmetic that catches almost everyone off guard the first time they see it: losses and gains are not mirror images of each other. A 20 percent loss doesn't need a 20 percent gain to break even. It needs a 25 percent gain, because the loss happened on a larger number than the recovery has to work with. A 50 percent loss is worse than it sounds for the same reason: it takes a full 100 percent gain, a complete doubling, just to get back to where things started. The deeper the hole, the more disproportionately steep the climb out of it.

This idea has been reduced to two rules, famously, by one of the most successful investors of the twentieth century.

> *"Rule No. 1: Never lose money. Rule No. 2: Never forget rule No. 1."*
>
> **— Warren Buffett**

It sounds almost too simple to be useful, until the arithmetic above makes clear why he says it. Avoiding a loss in the first place is mathematically easier than recovering from one, not just emotionally easier, but literally, provably easier, by the numbers. For someone still decades from retirement, that's a nice principle to keep in mind. For someone newly retired and withdrawing from the very portfolio that just took the hit, it's closer to the whole ballgame.



The mechanism is simple once it's visible. Withdrawing a fixed amount from a shrinking portfolio takes a bigger bite out of it, percentage-wise, than withdrawing the same amount from a growing one. A 20 percent loss in year one of retirement, combined with a withdrawal on top of it, can shrink a portfolio so much that even strong returns in later years are growing from a permanently smaller base. The same 20 percent loss in year fifteen barely matters, because there's no withdrawal compounding the damage and less time left for it to matter anyway.

This is exactly why average return, the number most people fixate on, can be dangerously misleading for anyone in or near retirement. An average masks the order. Two portfolios can carry the identical average and produce wildly different real-world outcomes, one funding a comfortable three decades, the other running dry with years of retirement still ahead.

It’s Not Just 2008

Walter’s story happened to land on the worst sequence-of-returns year in recent memory, but 2008 isn’t the only example. In 2022, the S&P 500 fell roughly 18 percent for the year, the worst calendar year since 2008 itself, driven by rising interest rates and inflation fears rather than a financial crisis. Anyone who retired in early 2022, expecting to draw a steady income from a portfolio built during the previous decade’s strong returns, absorbed that loss in the very first year of withdrawals, the single most dangerous moment for it to happen.

One scenario is worth running in real numbers: if a downturn like 2008 or 2022 happened again this year, what would it do to the retirement plan a person is counting on right now? Not in the abstract, but in real numbers. How much would need to be withdrawn regardless of what the market just did? How many years of spending would come directly out of a portfolio that just lost a fifth, or more, of its value? Most people have never actually run that number. A plan built with care can answer it.

The Volatility Buffer: A First Look

Sequence of returns risk is the backdrop that Social Security, Medicare, and tax decisions all get planned against. A Social Security claiming strategy, a Roth

conversion schedule, an annuity decision: all of it interacts with this same underlying problem, which is that retirement income needs to keep flowing even in the years markets don’t cooperate.

The most direct tool for managing this risk is a volatility buffer, a deliberately built cash or short-term reserve designed for one job: covering living expenses during a downturn so that investments are never forced to sell at a loss simply to generate this month’s income. Walter didn’t have one in 2008. The point is to make sure that’s a choice, not an accident.

Chapter 3: The Volatility Buffer

A wise man once lost a great deal of weight, and a friend asked him how he'd done it. "Don't eat," he replied.

It's a ridiculous answer, and also, technically, correct. A similarly ridiculous, similarly correct answer fits a much more expensive question: what should a retiree do when the market drops sharply, right when withdrawals are supposed to begin? Don't sell. Get the money somewhere else.

That's the entire idea behind a volatility buffer, stated in one breath. The rest is just the part where "somewhere else" gets defined, because saying it and building it are two very different projects.

What Walter Didn't Have

Walter Higgins didn't have a volatility buffer in 2008. What he had instead was a job he could go back to, which worked, but only because he was lucky enough to still be employable at sixty-six. He improvised, under pressure, exactly what a volatility buffer is designed to do on purpose: remove the need to sell investments at a loss while the market recovers. A paycheck became his "somewhere else." Not everyone has that option. A volatility buffer is how a retiree builds the same protection deliberately, years in

advance, instead of hoping a job offer shows up at exactly the right moment.

Two Different Jobs, Two Different Tools

Before going further, it helps to separate two ideas that sound similar but aren't: an emergency fund and a volatility buffer. An emergency fund covers the unexpected (a roof, a transmission, a medical bill) and it belongs in cash or a cash-equivalent account, typically six to twelve months of expenses, because it needs to be boring, liquid, and available on short notice. Nothing here changes that. Every retiree should still keep one.

A volatility buffer solves a different problem, and the math behind it rules out cash as the tool for the job. Think about what pure cash would actually require: a retirement can easily run twenty-five to thirty years, and a market downturn severe enough to matter can show up more than once in that span: 2008, 2022, and whenever the next one turns out to be. Self-funding a cash reserve large enough to ride out every one of those events, for decades, without ever touching the rest of the portfolio, would require setting aside an enormous share of the nest egg, money that would then sit there earning next to nothing for years at a stretch, often for no reason at all, since most years aren't down years. That's not a minor inefficiency. Over a multi-decade retirement, it's a real and growing

drag on the very wealth the buffer is supposed to protect.

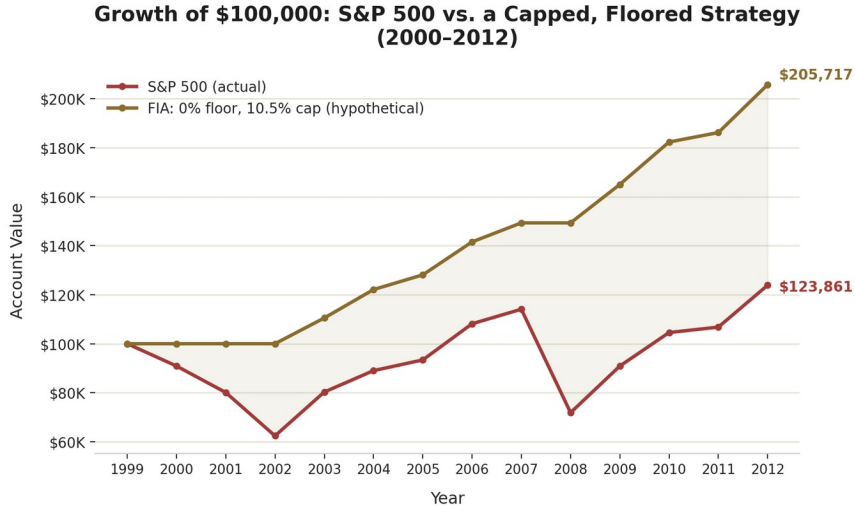
A Different Kind of Buffer

This is where a different tool tends to fit better: a fixed indexed annuity with a 0 percent floor, used specifically as the volatility buffer itself rather than as a cash substitute.

The mechanism is simple. A fixed indexed annuity credits interest based on the performance of a market index, commonly the S&P 500, but with two guardrails built into the contract. On the upside, the gain is limited by a cap, a participation rate, or a spread, so a year when the index returns 20 percent might only credit 8 to 12 percent, depending on the specific contract. On the downside, the floor is 0 percent: if the index loses value that year, the credited interest is zero, not negative. The account value does not decline because of market performance. A year that would have meant a painful, locked-in loss in a stock portfolio simply credits nothing in an FIA: disappointing, perhaps, but not damaging.

Real numbers make the effect concrete. The chart below applies a hypothetical FIA structure, a 0 percent floor and a 10.5 percent cap, to the S&P 500's actual year-by-year returns from 2000 through 2012, a stretch that includes both the dot-com crash and the 2008 financial crisis. $100,000 invested directly in the S&P 500 over those thirteen years grew to $123,861.

The same $100,000, credited under the FIA structure, grew to $205,717.



The gap isn't magic, and it isn't because the FIA outperformed the market in any single good year. It never did, since the cap always limited the upside. The gap exists entirely because of what happened in the bad years. In 2000, 2001, and 2002, the S&P 500 lost a cumulative 37.6 percent. The hypothetical FIA, with its 0 percent floor, lost nothing in any of those three years. It simply sat still while the index fell. By the time the recovery came in 2003, the FIA was starting from its full original base; the S&P 500 was starting from a hole it still had to climb out of. The same pattern repeated in 2008. This is the achieving-versus-protecting idea made visible: the FIA never tried to outperform a good year. It just refused to participate in the bad ones.

$100,000 Invested, 2000–2012: S&P 500 vs. a 0% Floor / 10.5% Cap FIA Structure

Year	S&P 500	FIA Credit	S&P Cumulative	FIA Cumulative
2000	(9.10%)	0.00%	(9.1%)	0.0%
2001	(11.89%)	0.00%	(19.9%)	0.0%
2002	(22.10%)	0.00%	(37.6%)	0.0%
2003	+28.68%	+10.50%	(19.7%)	+10.5%
2004	+10.88%	+10.50%	(11.0%)	+22.1%
2005	+4.91%	+4.91%	(6.6%)	+28.1%
2006	+15.79%	+10.50%	+8.1%	+41.5%
2007	+5.49%	+5.49%	+14.1%	+49.3%
2008	(37.00%)	0.00%	(28.1%)	+49.3%
2009	+26.46%	+10.50%	(9.1%)	+65.0%
2010	+15.06%	+10.50%	+4.6%	+82.3%
2011	+2.11%	+2.11%	+6.8%	+86.2%
2012	+16.00%	+10.50%	+23.9%	+105.7%
TOTAL	**—**	**—**	**$123,861**	**$205,717**

A few caveats belong right next to this table. The S&P 500 figures are real, historical, total returns, including dividends. The FIA figures are hypothetical, built by applying a 0 percent floor and a 10.5 percent cap to those same real index returns, which is a reasonable but simplified way to illustrate the mechanism. An actual contract would also reflect rider costs if any were elected, and the specific cap, participation rate, or spread offered by an actual carrier in an actual year, which moves over time and varies by contract. This isn’t a guarantee that any specific FIA would have produced exactly this result. It’s a demonstration of why the floor-and-cap structure tends to matter most precisely when it’s needed most.

What This Isn’t

Fixed indexed annuities are insurance products, not bank accounts, and they come with real terms that deserve the same scrutiny as any other financial decision. Most carry a surrender period, commonly seven to ten years, during which withdrawing more than a contractually allowed amount triggers a surrender charge. Caps and participation rates aren’t fixed for life; carriers can adjust them at each contract anniversary, within whatever minimum the contract specifies. The guarantee itself rests on the claims-paying ability of the issuing insurance company, not the FDIC. None of this makes the product wrong for the job; it makes it a contract that should be read, understood, and compared carefully before signing, exactly like any other meaningful financial decision. Contracts vary meaningfully by carrier, and a good retirement income advisor earns their keep here: researching which specific FIA, among the many available, actually fits a given client’s goals, time horizon, and risk tolerance, rather than treating every contract as interchangeable.

One more point belongs here, since fees determine how much of a return an investor actually keeps: many fixed indexed annuities, particularly ones built on simple index-crediting strategies without optional riders, carry no ongoing fee at all. That stands in real contrast to trading costs and the expense ratios charged by many mutual funds, both of which are

subtracted from returns year after year, whether the account goes up or down. A 0 percent floor combined with no ongoing fee means a larger share of whatever the contract credits actually stays in the account and keeps compounding: nothing is skimmed off the top before it has the chance to grow.

“Not the FDIC” is worth pausing on, since it’s exactly the kind of phrase that should raise a question rather than get skipped past. Every state runs its own guaranty association, funded by the insurance industry itself, which steps in if a carrier becomes insolvent, typically covering annuity benefits up to $250,000 per owner, per insurer, often with a higher aggregate cap depending on the state. It isn’t the same as federal deposit insurance, and claims can take time to process. But the underlying risk it protects against is genuinely rare: insurer insolvencies have stayed low for decades, including through 2008, when hundreds of banks failed nationwide and the S&P 500 lost 37 percent. The insurance industry’s core claims-paying obligations held up, even as one well-known insurance-adjacent company nearly failed because of a trading unit unrelated to the policies it had written. None of this is a guarantee about any specific company. It’s a reason the structure, used carefully, has earned the trust retirees place in it.

What It Frees Up Elsewhere

There's a second-order effect worth naming, because it changes how the rest of the portfolio gets built. Many pre-retirees and retirees handle market-downturn risk the same way: shift more of the portfolio from stocks into bonds as retirement approaches, on the theory that bonds are the conservative, lower-risk choice. There is a real crack in that theory: 2022 alone should put a dent in the idea that bonds are automatically safe.

A volatility buffer built on a 0 percent floor changes the math. If a meaningful piece of the portfolio is already protected from loss, not just lower-risk, but structurally unable to lose value to a market decline, the overall portfolio's risk has already come down before the stock-versus-bond allocation question even gets asked. That can free the rest of the portfolio to stay more invested in stocks than a retiree might otherwise feel comfortable with, rather than retreating into bonds that carry more risk than their reputation suggests. The buffer isn't replacing growth. It's making room for it.

A Second Option Worth Knowing About

A fixed indexed annuity isn't the only tool that can play this role. The cash value inside an indexed universal life insurance policy works on a similar principle: growth linked to a market index, with a floor that prevents losses in a down year. For some households it can serve as a second volatility buffer or a complement

to the first. IUL is a more complex product with its own set of considerations, including how it's funded and how its costs change over time, and it deserves its own complete treatment rather than a quick summary here. That fuller treatment comes later.

One quick comparison while it's fresh: a simple money market fund can also function as a short-term volatility buffer, and for some retirees that's exactly the right call. The trade-off is tax treatment. Interest earned in a money market fund is taxable income in the year it's earned, every year, whether or not it's touched. The growth inside an FIA or an IUL, by contrast, isn't taxed until it's actually withdrawn. Over a buffer that may sit untouched for years at a time, that difference compounds, quietly, but for real.

How Big Should It Be?

The sizing logic is similar in spirit to the cash-based version of this idea, even though the vehicle is different: the right size depends on reliable income already in place, risk tolerance, and how the rest of the portfolio is allocated. A retiree with a pension and Social Security covering most basic expenses needs less of this protection than someone relying almost entirely on portfolio withdrawals. That reliable income is already doing part of the buffer's job. Someone who would lose sleep and panic-sell at the first bad headline benefits from a larger allocation than someone who can calmly ride out a downturn. There's

no single percentage that fits everyone, which is exactly why this is a conversation to have with a professional who can see the whole picture, not a number to copy out of a book.

A Simple Way to Think About the Pieces

One useful way to organize this is to think in three pieces, each with a different job.

Piece	When It's Used	What It Holds	Job
Emergency Fund	Immediate	Cash, high-yield savings	Cover the unexpected: repairs, medical bills, surprises
Volatility Buffer	Ongoing	Fixed indexed annuity, 0% floor	Absorb down-market years without loss; never sold at a loss
Growth Portfolio	Long-term	Diversified stocks and bonds	Growth engine; replenishes the buffer in good years

The emergency fund stays small, liquid, and boring, exactly as it always has. The volatility buffer is the piece doing the real protective work, absorbing the years the market drops, without ever locking in a loss, while still being positioned to earn something in the years it doesn't. The growth portfolio remains where the bulk of long-term growth happens, since that money has the time to ride out volatility that the buffer and the emergency fund were never meant to take on.

The Cost of Doing This

The trade-off, since no strategy is free, is this: money inside an FIA-based volatility buffer is, by design, giving up some upside in the strong years in exchange for never losing principal in the weak ones. That's not a flaw. It's the price of the floor, paid in the form of a cap, a participation rate, or a spread, for the years a downturn would otherwise have forced a sale at the worst possible moment. Compare that trade-off against Scenario B, a portfolio fully depleted by year eleven, and it looks considerably more reasonable.

> *"Don't sell. Get the money somewhere else."*
>
> — the whole idea, in one sentence

That is the core idea. Everything else in retirement income planning (Social Security timing, Medicare decisions, Roth conversions, annuities) gets easier to think about once this piece is in place, because none of those decisions have to be made under the pressure of a forced sale. The first of the big, hard-deadline decisions is when to claim Social Security, and why that choice is mostly irreversible once it's made.

Chapter 4: Social Security 101

Here is a question worth sitting with first: have you ever gotten a tattoo, or a genuinely bad haircut? Both have something in common with claiming Social Security that almost nobody thinks about in the moment. Once it is done, it is done. A haircut grows back eventually. A Social Security claiming decision, in most circumstances, does not. It sets a number that follows a person for the rest of their life, with nothing but an annual cost-of-living adjustment to soften it.

> Walter Higgins found a narrow exception to that rule. After claiming at sixty-six in 2008, the market crash forced him back to work, and he suspended his benefit rather than keep collecting it. That was a rare second chance to improve the timing, but one that only exists for a specific window, under specific conditions, for people who happen to be between full retirement age and seventy. Most claiming decisions offer no such escape hatch. Get it wrong at sixty-two, and the reduction is permanent. There is no going back.

Here is the frame that makes all of this easier to think about. Strip away the politics and the paperwork, and Social Security is a simple thing wearing a complicated costume. It is a lifetime annuity. A person pays in during their working years, and in return gets a monthly check that keeps coming for as long as they live, rises a little each year to track inflation, and is backed by the federal government rather than a single insurance company. Private insurers sell products that try to copy that. None

of them can match the price, and none can match the inflation adjustment. The strange part is that the most valuable annuity most people will ever own is also the one they tend to switch on by accident. That is what the basics are for: so the decision gets made on purpose.

A quick scope note before starting. Social Security runs several different programs under one name, and the focus here is retirement benefits for aged workers and their families. It is not about Social Security Disability Insurance, which pays workers who become unable to work before retirement age, and it is not about Supplemental Security Income, a separate, needs-based program for people with very limited income and resources that is not actually funded by Social Security taxes at all. Both matter. The system was built to catch people who fell early, not only those who grew old on schedule. Close to one in five Social Security dollars goes to someone collecting for a reason that has nothing to do with age: a worker disabled at fifty, a widow raising two children after her husband died at forty-one, a child who lost a parent. That part of the system rarely makes the brochures. It is still doing quiet work in the background.

One Leg of a Three-Legged Stool

For a long time, retirement in America rested on three legs: Social Security, an employer pension, and whatever a person managed to save and invest on their own. The stool still stands, but the middle leg has gotten shorter for almost everyone. Traditional pensions, the kind that pay a guaranteed monthly amount for life, have mostly vanished from private companies. A person who still has one today probably worked for a

government or belonged to a union. Everyone else got handed a 401(k) and the job of turning it into income themselves.

That shift puts more weight on the two legs that remain, and it makes the timing of the Social Security leg matter more than it used to. For a lot of households, Social Security is now the only piece of the plan that behaves like an actual pension, which is exactly why the decision of when to start deserves careful attention rather than a shrug.

Older Than the Rules It Runs By

> Walter was born in 1942, and his full retirement age was sixty-five years and ten months, which means claiming at sixty-six put him just two months past it. Anyone born in 1960 or later has a full retirement age of sixty-seven, a full fourteen months later than Walter's.

Congress raised the full retirement age gradually, starting with people born in 1938, in response to people living longer than the system's original designers expected. (Also, the money was running out). The increase phased in slowly, a couple of months at a time depending on birth year, until it leveled off at sixty-seven for anyone born in 1960 or later, where it has stayed since. The schedule looks like this:

Birth Year	Full Retirement Age
1943–1954	66
1955	66 and 2 months
1956	66 and 4 months

Birth Year	Full Retirement Age
1957	66 and 6 months
1958	66 and 8 months
1959	66 and 10 months
1960 or later	67

Anyone born from 1955 through 1959 sits in the transition zone and should confirm their exact full retirement age directly rather than assume it. Everyone born in 1960 or later has a flat, simple answer: sixty-seven. That is the number most readers will plan around, even though Walter's story, and plenty of other people's stories, ran on a different one.

What "Full Retirement Age" Actually Means

Full retirement age, usually shortened to FRA, is the age at which Social Security pays a worker's full earned benefit: not a penny reduced, not a bonus added. Every other claiming age is measured against it. Claim before FRA, and the benefit is permanently reduced. Claim after FRA, and it permanently grows, up to age seventy. FRA is the fulcrum the whole decision balances on.

The technical name for the full, unadjusted benefit is the Primary Insurance Amount, or PIA, a term that comes back when spousal and survivor benefits enter the picture, since those are calculated from it too. For now, it is enough to know that the PIA is the benefit Social Security has determined a worker earned, based on their highest thirty-five years of income, paid in full at their personal FRA.

The Three Numbers That Matter Most

Three claiming ages are worth holding in mind: sixty-two, FRA, and seventy. Sixty-two is the earliest age to claim a worker's own retirement benefit, and it comes with the steepest discount. For someone with an FRA of sixty-seven, claiming at sixty-two locks in a permanent 30 percent reduction. Not temporary, not something that corrects itself later. Just a smaller check for the rest of that person's life. Seventy is the other end of the range. Delaying past FRA earns delayed retirement credits worth roughly 8 percent a year, capping out at age seventy with a 24 percent increase over the FRA amount. Credits stop accruing after seventy, so there is no reason to wait longer than that.

To put dollar figures next to the percentages: for a worker with maximum lifetime earnings and an FRA of sixty-seven, the 2026 maximum monthly benefit is $2,969 at sixty-two, $4,152 at FRA, and $5,181 at seventy. Most readers will not have maximized every working year, so their own figures will be smaller. The average retired worker collects closer to $2,071 a month as of the most recent cost-of-living adjustment. The percentages hold regardless of income level. The gap between claiming at sixty-two and claiming at seventy is roughly a 75 percent difference in monthly income, locked in for life, based entirely on one filing decision.

A Bigger Piece Than People Think

It is easy to treat Social Security as the small, dependable background hum of a retirement and to put the planning energy into the portfolio. The numbers argue against that. Measured across all older

households, Social Security provides roughly a third of total income, and for the average retiree it lands somewhere between a third and two-fifths of the money arriving each year. The share climbs sharply further down the income ladder, where it can be nearly everything, and it shrinks for the wealthy. Even affluent retirees tend to lean on it for the one thing the rest of the portfolio cannot provide: a check that shows up every month, adjusts for inflation, and never has a bad year in the market. A plan that treats it as a rounding error is usually a plan nobody has looked at closely.

Getting a Personal Statement

Every number here is a generality. The number that actually matters is the one tied to a reader's own earnings record, and Social Security makes it easy to get. There are two ways to request a Social Security Statement.

The fast way is online: create a free "my Social Security" account at https://ssa.gov/myaccount. It takes a few minutes. Proving identity through Login.gov or ID.me is the only step that takes any effort, and the Statement appears instantly, available to view, save, or print anytime afterward. The slower alternative is Form SSA-7004, downloadable from ssa.gov, filled out by hand and mailed to the Social Security Administration's Wilkes-Barre Direct Operations Center at the address printed on the form. The paper version usually takes four to six weeks to arrive.

The Statement itself is worth a close look. The redesigned version opens with a bar chart showing a personalized benefit estimate at nine claiming ages,

from sixty-two through seventy, a direct visual answer to the question that matters most. Below that sits a full, year-by-year earnings history: every year of income on record, along with an estimate of Social Security and Medicare taxes paid, broken out separately for worker and employer. The Statement also includes a disability benefit estimate, survivor benefit estimates for eligible family members, and a plain-language summary of Medicare eligibility.

One detail deserves close attention. The earnings history on the Statement is the actual record Social Security will use to calculate a benefit someday. An employer that misreported a name, a year of self-employment income that never made it into the system, a gap that should not be there: these errors are common, and they can quietly shrink a future benefit if they are never caught. Reviewing the Statement once a year, and flagging anything wrong using the instructions printed on it, is a small habit with outsized consequences.

The Help Desk That Is Required to Stay Neutral

Here is something that surprises almost everyone who telephones the Social Security Administration (SSA). The SSA is not allowed to advise a person on when to claim. A representative can tell someone exactly what their benefit would be at sixty-two, at full retirement age, and at seventy. They can process the application and answer questions about the rules. What they cannot do, by policy, is say which choice is best. They will not say "wait," and they will not say "claim now." That line is drawn on purpose. It also means the single most

consequential financial decision in front of many retirees comes with a help desk that is required to stay neutral. The mechanics to claim are available from the agency. The judgment is not, and it has to come from somewhere else.

There is one more thing the agency does not do: coordinate the benefit with the rest of a person's money. Social Security has no idea what is in an IRA, when a pension starts, or how a Roth conversion this year will interact with a tax bill. It pays what its formula says to pay, on the schedule a person chooses. Fitting that stream together with everything else, so the pieces do not collide and spike a tax bill or a Medicare premium, is left entirely to the household. That coordination is the whole subject of the chapters that follow.

Working While Collecting: The Earnings Test

Claiming before FRA while still working introduces a wrinkle that catches people off guard. For anyone under FRA for the entire year, Social Security withholds $1 in benefits for every $2 earned above $24,480 in 2026. In the calendar year someone actually reaches FRA, the limit jumps to $65,160, with a gentler $1 withheld for every $3 over that amount, counting only earnings in the months before the birthday. Once FRA arrives, the earnings test disappears completely. There is no limit at all on how much can be earned alongside a full Social Security check.

One detail eases a lot of unnecessary panic: money withheld under the earnings test is not gone. Social Security recalculates the benefit upward once FRA is reached, crediting back the months that were reduced.

The practical effect is closer to an involuntary, temporary pause than a true loss. It is still worth planning around, since it affects cash flow in the meantime, but it is not a reason to assume that early claiming and working simply cannot mix entirely.

How Social Security Gets Taxed

There is a second, separate way income interacts with Social Security, and it has nothing to do with claiming early or still working. It applies at any age, to anyone with enough other income. Up to 85 percent of a Social Security benefit can be counted as taxable income at the federal level, depending on something the IRS calls provisional income: adjusted gross income, plus tax-exempt interest, plus half of the Social Security benefit itself. For a single filer, provisional income above $25,000 makes up to half the benefit taxable; above $34,000, up to 85 percent. For a married couple filing jointly, the thresholds are $32,000 and $44,000. These numbers were written into law decades ago and have never been adjusted for inflation, which is precisely why they catch more retirees every year even when nothing about their circumstances has changed.

The mechanics create a trap that the brackets alone do not reveal. In the zone where provisional income climbs through the low thirties and into the forties, each additional dollar of ordinary income can drag another fifty to eighty-five cents of benefit into the taxable column. The effective tax rate on that next withdrawal is far higher than the stated bracket suggests, which is why a claiming decision and a tax plan cannot be made in separate rooms. A temporary senior deduction added in 2025 softens the bite for some retirees, but it does not

repeal the provisional-income formula, so the planning still matters. This is ordinary income tax, calculated the same way Roth conversions and required minimum distributions affect a tax bill, which is exactly why it returns in the later tax-planning chapters.

A Quick Word on Spousal Benefits

Married couples have an additional layer to think about. A spouse can claim a benefit worth up to 50 percent of the other spouse's PIA, even with little or no earnings history of their own. A few rules shape how that works. The higher-earning spouse generally has to be receiving their own benefit before the other spouse can claim against it. Spousal benefits cap at 50 percent of the worker's FRA amount and never grow with delayed retirement credits, so there is no reward for a spouse waiting past their own FRA to file for one. Claiming a spousal benefit early carries its own, steeper discount: claimed at sixty-two, it shrinks to roughly 32.5 percent of the worker's PIA, a sharper cut than the worker's own early-claiming reduction. Whoever claims a spousal benefit always receives the higher of their own benefit or the spousal amount, never both stacked together.

Divorce and Death Change the Rules

A marriage ending, whether through divorce or death, does not necessarily end a spouse's claim on a former partner's earnings record. The rules differ depending on which one happened, and both come up often enough to deserve a clear explanation.

A divorced spouse can claim on an ex-spouse's record, even if that ex has remarried, as long as the marriage lasted at least ten years, both people are at least sixty-

two, and the person claiming is currently unmarried. The mechanics mirror a current spouse's benefit: up to 50 percent of the ex's PIA, reduced if claimed early. This claim has no effect whatsoever on the ex-spouse's own benefit, or on what a new spouse might receive. Social Security treats it as an entirely separate entitlement, not a slice carved out of someone else's check.

Survivor benefits, triggered by an ex-spouse's or spouse's death, work on a different and considerably more generous scale. Rather than capping at 50 percent, a surviving spouse or surviving divorced spouse (the same ten-year marriage rule applies) can receive up to 100 percent of what the deceased was actually receiving, including any delayed retirement credits the deceased had earned by waiting past their own FRA. Survivor benefits can also start earlier than any other claiming age, as early as sixty, two full years before a worker's own retirement benefit becomes available, with a maximum reduction of 28.5 percent for claiming that early. One rule trips people up more than any other: remarrying before age sixty ends eligibility for survivor benefits on a deceased spouse's record, permanently. Remarrying at sixty or later does not affect it at all. That single birthday is a hard, sometimes overlooked, financial line.

Is Social Security Going Away?

No, but it helps to understand exactly what is true behind the worry, since a vague anxiety is worse than a clear picture of the mechanics. Social Security's trust funds are projected, under current law and current projections, to deplete their reserves sometime around

2034. That date shifts slightly with each year's report, but it has stayed in the same neighborhood for a while.

Depletion does not mean the program runs out of money. Social Security is funded mainly by ongoing payroll taxes, not just the trust fund, and that revenue keeps flowing in regardless of what happens to the reserves. Under current projections, if Congress does nothing at all between now and then, incoming payroll tax revenue alone would still cover roughly 80 percent of scheduled benefits after depletion. That would be an automatic, across-the-board cut, legally mandated under current law, but a long way from the program disappearing.

A candid opinion: an actual benefit cut is unlikely to happen. Congress has faced a Social Security funding crisis before, and the last major fix, in 1983, was enacted only once the system was genuinely close to a shortfall. The same political pattern seems likely to repeat. Raising taxes or trimming benefits is unpopular no matter who proposes it, and no politician wants to be the name attached to reducing or delaying Social Security checks. That is a prediction about political behavior, not a guarantee, and it should not be mistaken for one. The useful takeaway: depletion is a fixed date on the calendar, the mechanics are well understood, and Congress has a long history of acting right around the deadline rather than years before it.

Two Things Worth Knowing (Even If They Are Not the Point)

It is easy to assume American Social Security is generous. Measured against other wealthy countries, it is not. For an average earner it replaces roughly 40

percent of prior wages, which lands below the average among developed nations. By international standards, the U.S. benefit is modest.

It is also, by design, progressive. The formula deliberately replaces a larger share of income for people who earned less over their lifetimes. A low-wage worker gets back about half of prior earnings; a high earner gets back closer to a third. One consequence is that the program disproportionally favors people who earned less, and by a disproportionally higher percentage of earned wages. Neither point changes a single reader's claiming math, but both are useful to know about the thing everyone argues over and few people describe accurately.

What Is Left for Later, On Purpose

The information so far in this chapter is the foundation, not the strategy. It has not touched the harder questions: when it actually makes sense to claim early versus wait, how survivor benefits change the calculation, what a couple should do when their ages and earnings records do not match neatly. Those are the work of claiming strategy, and they only make sense once the mechanics here are solid. For now the goal is simpler: know what FRA means, know that the gap between claiming ages is serious money, and know to check the specific number that applies personally rather than assume it matches a parent's or grandparent's experience. The rules have changed, and Walter's story is proof of it.

Chapter 5: Claiming Strategies — When to Start and Why It's (Mostly) Irreversible

Picture two people with identical careers: the same job, the same paychecks, the same Social Security record down to the dollar. They make exactly one different choice. One claims at sixty-two, the other waits until seventy. For the rest of their lives the second person's monthly check is about 77 percent larger. Same work, same record; nothing separates them but the calendar. Over a long retirement, that single choice can be worth six figures, which is why it deserves to be made on purpose rather than by reflex. There is no one right age that fits everyone. It turns on health, on whether a person is married, on what other money is available, on how long the family tends to live, and on how the tax code treats the result. But there is a right way to think it through.

Most people never get that far. Roughly a third claim at the first opportunity, at sixty-two, and many do it for reasons that have little to do with the math: a fear that the program will not last, a vague sense that a bird in the hand beats two in the bush, or simply because sixty-two is the first door that opens. The instinct is understandable. It is also, for many people, expensive. One widely cited analysis estimated that the average retiree leaves roughly $182,000 on the table by claiming at a financially suboptimal time. The number is large because the decision is permanent.

Nine Thousand Ways to Get This Wrong

The size of the problem is easy to underestimate, so here is the arithmetic. Each spouse in a married couple can start a benefit in any month between sixty-two and seventy, which works out to ninety-seven possible starting months apiece. Ninety-seven times ninety-seven is 9,409 combinations, and that is before layering in spousal and survivor rules. The Social Security Administration has published the math itself. The point is not that a couple should evaluate all nine thousand options. The point is that "when should we claim?" is not a yes-or-no question, and treating it like one is how households leave serious money behind without ever knowing it was on the table.

Smaller Checks Sooner, Bigger Checks Later

The trade is simple to state. Claim early and the checks start sooner but each one is smaller, for life. Wait and each check is larger, for life, but there are fewer of them. The machinery is straightforward: claiming at sixty-two with a full retirement age of sixty-seven cuts the benefit by about 30 percent, while every year of waiting past full retirement age adds 8 percent. Stretched from full retirement age to seventy, that works out to a permanent 24 percent increase above the full benefit. End to end, the benefit at seventy lands at 124 percent of the full amount and the benefit at sixty-two at 70 percent of it, a swing of more than three-quarters between them.

For someone with the maximum earnings history in 2026, that is the difference between about $2,969 a month at sixty-two and roughly $5,181 at seventy, for the identical work record, with the only variable being

the year the person decided to start. Most people are nowhere near the maximum; the average retired worker collects closer to $2,071 a month. But the percentages apply at every income level, which means the size of the decision scales with the size of the benefit.

Age benefits begin	% of full benefit	2026 max monthly
62	70%	$2,969
Full retirement age (67)	100%	$4,152
70	124%	$5,181

Monthly benefit by the age benefits begin, assuming a full retirement age of sixty-seven. The percentages apply to any benefit amount.

And it is, in almost every case, a one-time decision. There is no annual re-vote. Apart from two narrow exceptions covered at the end of this chapter, the number a person locks in at the moment of claiming is the number that follows them, and their surviving spouse, for the rest of their lives, lifted only by the annual cost-of-living adjustment.

What the Calculators Can't See

The traditional tool for comparing now against later is the break-even analysis. It finds the age at which the larger delayed checks, added up, finally overtake the smaller checks that started earlier. Before that age, claiming early comes out ahead; after it, waiting wins.

Take Linda Castellano, a fictional single saver. She is fifty-eight, single, and a healthcare administrator. Suppose her benefit at full

retirement age works out to $2,400 a month. Claiming at sixty-two would reduce it to about $1,680; waiting to seventy would raise it to roughly $2,976. By starting at sixty-two she banks eight years of checks before the patient version even begins. But once the larger check is flowing, it gains about $1,296 every month on the smaller one, and somewhere around age eighty the totals cross. Live past eighty, and waiting was the better financial call. Die before it, and claiming early was.

A break-even view of Linda's numbers. The crossover ignores future cost-of-living increases and taxes, which can move the crossover age by years.

That clean little crossover is where most people stop. It is also where the analysis misleads them, because a break-even calculation only works if a person knows the one thing nobody knows: the date they will die. Every break-even age ever produced is the answer to a question that cannot be answered in advance.

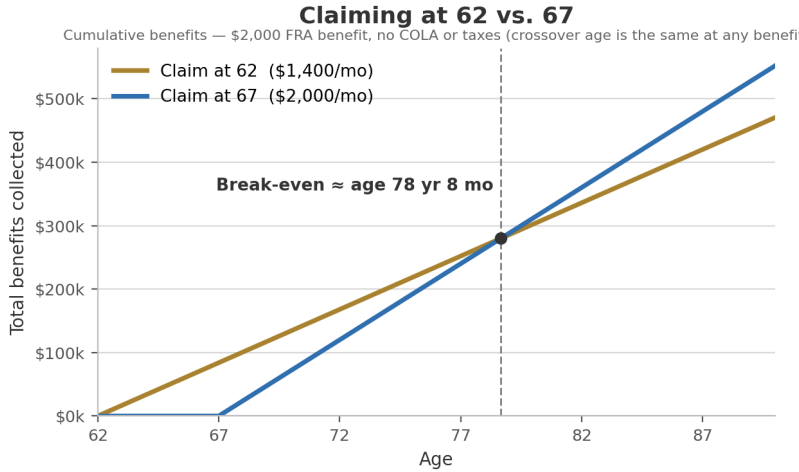
There is more the calculators leave out. Most of the free online tools compute the crossover on the benefits alone. They do not know a person's tax situation, and the share of a Social Security check that is taxable can move the answer by years. They do not account for the way a larger benefit can push income into a higher Medicare premium bracket, the IRMAA surcharge, or for what happens to the rest of the portfolio while a person spends down other accounts to keep waiting. A break-even age that ignores taxes, Medicare, and the drawdown of everything else is a tidy number resting on a partial picture.

Better tools exist for anyone who wants to run the numbers. A free online break-even calculator (the one at safemoney.com is a clean example) will show the crossover age in a few clicks, and more complete planning software such as Boldin (boldin.com) can model the claiming decision alongside taxes, withdrawals, and the rest of the retirement picture. But the most reliable option is usually a person rather than a program: a retirement-income advisor who holds a Social Security education certification can account for the variables a calculator cannot see and fit the timing to one household's whole situation. None of this makes break-even analysis useless. It makes it one input among several, a way to see the shape of the trade-off rather than a verdict. The more revealing question is usually not when the lines cross, but what the decision is actually protecting against.

The Three Classic Break-Evens

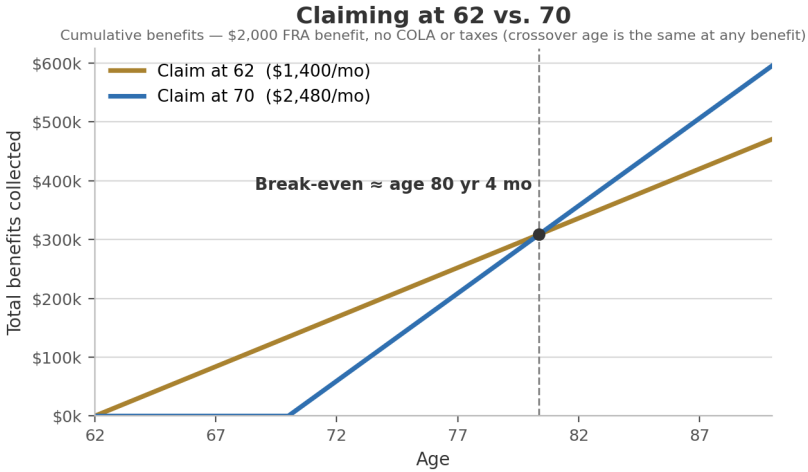
To make that trade-off concrete, it helps to see it drawn. Almost every do-it-yourself claiming decision runs through the same exercise, a break-even between two ages, and three comparisons cover the choices most people face: sixty-two against sixty-seven, sixty-two against seventy, and sixty-seven against seventy. The examples use two round benefits at full retirement age: $2,000 a month, close to the average, and $4,000 a month, near the high end and typical of a career professional. One fact comes first: the break-even age is the same at $2,000 and at $4,000, because the percentages scale evenly. A larger benefit changes the dollars at stake, not the age where waiting pulls ahead. The graphs use the $2,000 benefit and run the classic

way, cumulative dollars collected with no cost-of-living adjustment and no taxes.



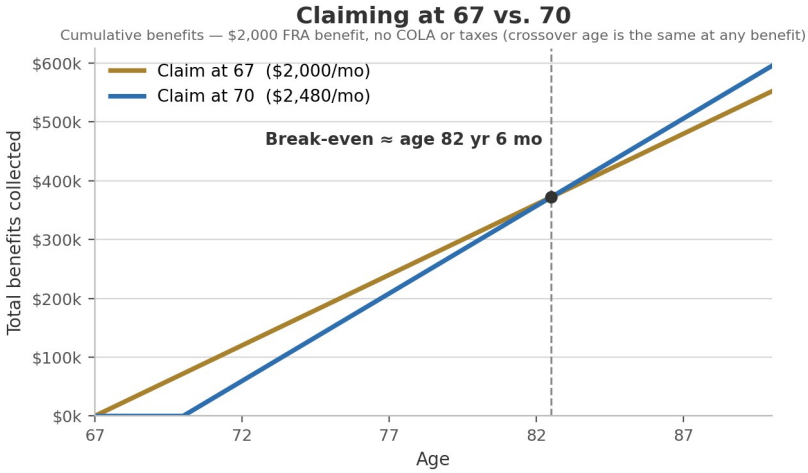
Claiming at 62 versus 67. Crossover near age 78 years, 8 months.

Start with sixty-two against sixty-seven. On the $2,000 benefit, claiming at sixty-two pays about $1,400 a month; waiting to sixty-seven pays $2,000. The early claimer collects for five extra years and builds a lead of about $84,000 before the patient version starts. From sixty-seven on, the larger check narrows that gap by $600 a month, and the two totals draw even at about age seventy-eight years and eight months. On a $4,000 benefit the lead and the gap both double, and the crossover holds at the same age. Live past it, and waiting won; die before it, and starting early won.



Claiming at 62 versus 70. Crossover near age 80 years, 4 months.

Sixty-two against seventy is the widest common choice. On $2,000, seventy pays about $2,480 a month, a swing of more than three-quarters above the sixty-two figure. Claiming eight years early builds a head start of about $134,000, but the monthly gap after seventy is large too, about $1,080, and the lines meet at close to age eighty years and four months. On a $4,000 benefit the head start is roughly $269,000 and the crossover is unchanged.



Claiming at 67 versus 70. Crossover near age 82 years, 6 months.

Sixty-seven against seventy is the question for someone who has already ruled out claiming early: stop at full retirement age, or push the last three years. On $2,000, the head start from claiming at sixty-seven is about $72,000, and the $480 monthly gap after seventy takes until age eighty-two years and six months to erase, the latest crossover of the three. On a $4,000 benefit the dollars double and the age stays put.

A pattern stands out: the longer the wait, the later the break-even, from about seventy-nine for the shortest delay to past eighty-two for the last three years. Waiting always costs money up front and returns it only to a person who lives past the crossover. Since about one in three of today's sixty-five-year-olds reaches ninety and one in seven reaches ninety-five, a large share of retirees live well past every one of these ages.

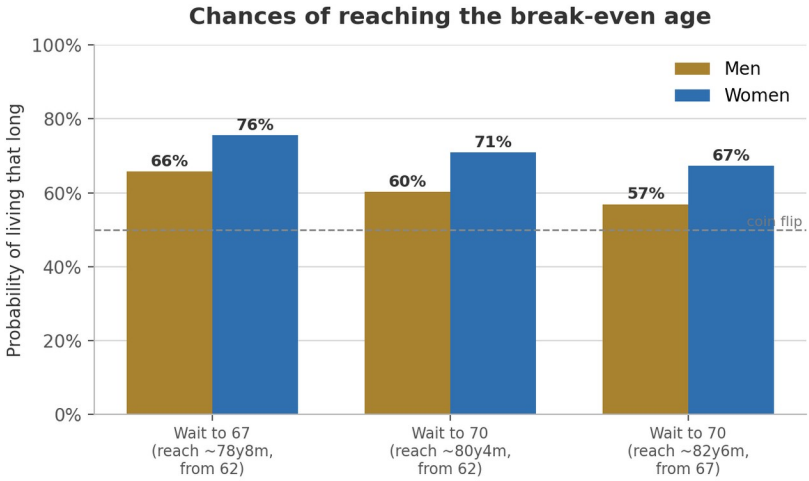
These are the classic lines, and the classic version is also incomplete, exactly as the calculators warned. A cost-of-living adjustment raises the larger delayed check by more dollars each year and pulls every crossover a year or two earlier. Taxes and the Medicare surcharge push the other way and depend on the household. Treat the three ages as the shape of the trade, not the verdict.

The arithmetic here is the simple part. The claiming decision is harder, because the age chosen also moves ordinary income taxes, provisional income, and the IRMAA surcharge, and for couples it turns on two full retirement ages, an age gap between spouses, differing life expectancies, and sometimes a divorce. Break-even is a basic and important input, not the answer, and the

claiming decision is itself one moving part of a whole retirement income and withdrawal strategy. The sections that follow take up the pieces the arithmetic leaves out.

Will You Reach the Crossover?

A break-even age only matters if a person reaches it, so the honest next question is what the odds are. By the Social Security Administration's 2023 period life table, most people do. A sixty-two-year-old man has about a 66 percent chance of living to the seventy-eight-years-eight-months break-even for waiting to sixty-seven; a woman's chance is about 76 percent. Extend the comparison to seventy, with its crossover at about eighty and a third, and the odds ease to roughly 60 percent for men and 71 percent for women. Even the latest crossover, age eighty-two years and six months for the sixty-seven-versus-seventy choice, beats a coin flip for men at about 57 percent and reaches two in three women.



Odds of living to each break-even age, SSA 2023 period life table.

For a person in average health, the base-case odds tilt toward waiting on all three comparisons. The reason to claim early is seldom that the average person will not reach the break-even. It is that a particular person has a particular reason: poor health, a family history of short lives, a need for the income now, or a spouse whose survivor benefit is better protected another way.

The mirror image of the survival question is the size of the prize. Suppose a person knew they would reach ninety. Waiting would carry no gamble at all, and the reward is easier to see in dollars. On a $2,000 benefit, in nominal dollars, waiting to sixty-seven instead of sixty-two hands a ninety-year-old about $81,600; waiting to seventy instead of sixty-two is worth about $124,800; and waiting the last three years adds about $43,200. On a $4,000 benefit each figure doubles, to roughly $163,200, $249,600, and $86,400. Cost-of-living adjustments would push all of them higher.

Decision	Break-even	Odds (M/W)	+$ by 90, $2k	+$ by 90, $4k
Wait to 67 (vs 62)	~78y 8m	66/76%	+$81,600	+$163,200
Wait to 70 (vs 62)	~80y 4m	60/71%	+$124,800	+$249,600
Wait to 70 (vs 67)	~82y 6m	57/67%	+$43,200	+$86,400

SSA 2023 period life table. "+$ by 90" = extra lifetime benefit by age 90 on a $2,000 vs $4,000 FRA benefit, nominal dollars.

Nobody gets that certainty. But the two numbers together, the odds of reaching the crossover and the dollars waiting past it, frame the decision better than a break-even age alone. A majority reach the line, and the reward for those who do runs into tens of thousands of dollars.

Social Security is full of deliberate design choices that favor some people over others, and the numbers in this chapter put a few of them in plain view. The benefit formula is steeply progressive by design. For a worker retiring at sixty-seven, it replaces about three-quarters of a very-low earner's career-average wages, roughly half of a low earner's, and only about a third of a high earner's, sliding to about 27 percent for someone who earned the taxable maximum every year[1]. And because wages above that maximum, $184,500 in 2026[2], are neither taxed nor counted, the replacement rate keeps falling for anyone who earns more: a person making a million dollars a year has Social Security replacing only a sliver of actual income. The survival odds show a quieter tilt in the same spirit. Women live longer, so a larger share of them clear every break-even and collect for more years, yet the full retirement age, set in part around life expectancy, makes no adjustment for the life expectancy gap between men and women. Imagine the political debate if females benefits were cut because they are women. The spousal and survivor rules are another piece of social engineering, letting a person draw a benefit built entirely on someone else's earnings. A divorced spouse who never worked, provided the marriage lasted ten years, or a surviving spouse, can end up with a larger monthly check than someone who worked continuously for forty years at modest wages. Whatever one makes of these choices, they are choices,

[1] Replacement rates: Social Security Administration, Actuarial Note No. 2026.9, “Replacement Rates for Hypothetical Retired Workers” (worker retiring at age 67): very-low earner 75.5%, low 55.0%, medium 41.0%, high 33.7%, steady maximum 26.9%.

[2] Social Security Administration, “Contribution and Benefit Base,” 2026: $184,500 (ssa.gov/oact/cola/cbb.html).

and they deserve a clear look before a person builds a plan around the program.

The size of these tilts is easy to underestimate, so here are two concrete cases. Take life expectancy first. By the Social Security Administration's 2023 period life table, the average sixty-five-year-old woman will live about 20.7 more years and the average man about 18.1 a gap of roughly two and a half years. With an identical earnings record and the same claiming age, that difference alone comes to about thirty extra monthly checks: roughly $61,000 more in lifetime benefits on a $2,000 benefit, or about $122,000 on a $4,000 one, paid to the woman for nothing more than living longer. The spousal tilt runs larger still. A person who worked full-time at the federal minimum wage for forty years retires on roughly $950 a month. A spouse who never held a job, but who was married for at least ten years to someone who earned the taxable maximum and delayed to seventy, can draw a survivor benefit of $5,181 a month, the 2026 maximum[4], more than five times as much, on a record they never paid a dollar into. Even the ordinary spousal benefit, half of the high earner's full amount, comes to about $2,076 a month, still more than double the lifelong minimum-wage worker's own check.

[3] Social Security Administration, Period Life Table, 2023 (used in the 2026 Trustees Report). Life expectancy at age 65: 20.66 years (female), 18.12 years (male).

[4] Social Security Administration, “Maximum Social Security Retirement Benefit,” 2026: $5,181/month at age 70, $4,152 at full retirement age, $2,969 at age 62.

Longevity Insurance, Not a Bet

Here is the reframe that changes how the whole decision looks. Waiting to claim is not a bet that a person will live a long time. It is insurance against the risk that they do.

The danger in a long retirement is not dying early; a person who dies early, by definition, does not run out of money. The danger is living to ninety-three on a portfolio built for eighty-three. The years when that risk bites hardest, the late eighties and nineties, when other accounts may be thin and going back to work is no longer an option, are exactly the years when a larger, inflation-adjusted, government-backed check matters most. Delaying Social Security buys more of that check. The odds of needing it are higher than most people assume: about one in three of today's sixty-five-year-olds will reach ninety, and one in seven will reach ninety-five.

This is where the claiming decision connects back to the volatility buffer. A bad sequence of market returns early in retirement can hollow out a portfolio a person is drawing on to live, and a buffer of dependable income that does not move with the market is the protection against it. A larger Social Security benefit is precisely that kind of buffer: guaranteed, inflation-adjusted, and payable for life. Every dollar of it is a dollar the portfolio does not have to produce in a down year. Seen this way, the 8 percent a year that waiting adds is not a return to measure against the stock market. It is the price of buying more of a guaranteed lifetime annuity, often more cheaply than an equivalent income could be purchased anywhere else. For a healthy person who can

cover the gap years another way, it is one of the most efficient purchases in all of retirement planning.

Social Security Against the Six Risks

A retirement plan is not about beating the market. It is about surviving six specific risks, and Social Security touches every one of them. It handles some better than any other asset a person owns, and quietly worsens one or two. Most people have never lined it up against the list.

Longevity is the risk of outliving the money, and it is the one Social Security answers best. A larger, inflation-adjusted check that cannot stop is the cleanest longevity insurance available, and, as above, the odds of needing it are rising as people live longer.

Liquidity is where the picture is split. The check itself is about as liquid as money gets: it lands in the account every month, spendable immediately, with no asset to sell, no market to time, and no fund to wait on. The asset behind the check is not liquid at all, because a person cannot borrow against future benefits or cash them out in a lump sum. Social Security delivers perfect income liquidity and zero principal liquidity, which is a fair trade for a retiree who needs steady spending money more than a pile they might raid.

Inflation is where Social Security beats almost everything else in the portfolio. It carries a built-in cost-of-living adjustment, and most assets do not. A fixed pension check buys a little less every year until, over a long retirement, it can lose a third of its purchasing power. A bond pays the same coupon whether bread costs two dollars or five. Social Security raises the

payment to track consumer prices, automatically, for life. For a benefit that might be collected for thirty years, that feature is worth more than it looks.

Market volatility is the risk Social Security simply is not subject to. A market crash the year a person retires can gut a portfolio they are drawing on; it does nothing to the Social Security check. That comfort comes with one honest caveat. Market-proof is not the same as risk-free. The trust fund reserves are projected to run short around 2034, and if Congress does nothing, incoming payroll taxes would still cover roughly 80 percent of scheduled benefits. That is a legally mandated gap under current law, not the program vanishing, and Congress has closed such gaps before at the last minute. But "probably" is a political forecast, and a plan should know the difference between market-proof and risk-free.

Mortality is the risk Social Security handles worst, and the one couples overlook most. When one spouse dies, the household does not keep both checks. It keeps the larger one, and the smaller one stops, often at the same moment the survivor moves from married-filing-jointly to single, where the brackets are tighter and more of the remaining income is exposed. Expenses rarely fall by half when a spouse dies. Income can. That single shift is why having the higher earner wait, locking in the biggest possible survivor check, is the most durable decision in the whole exercise.

Taxes come last, and Social Security is not tax-free for most people. Half of the benefit is folded into the provisional-income formula that decides how much of the check itself is taxed, and the same income can lift

Medicare premiums through the IRMAA surcharge. That is one more reason a claiming decision and a tax plan cannot be made in separate rooms, a thread the later chapters pull all the way through.

Line the six up, and a question falls out of them, the one that actually matters. Are these risks, and the specific way Social Security answers each, written into the plan? Or is the biggest inflation-indexed annuity a person will ever own just running in the background while the plan focuses on the parts that are more fun to talk about?

Why People Claim Early, and When They Are Right

The standard advice is to wait, and for a lot of people that advice is correct. But "always wait" is lazy, and it ignores the perfectly good reasons a person might start early. There are at least five, and they pull in different directions for different households.

The first is simply wanting to retire. If the whole point of the money is to stop working and start living, and the numbers hold together, then holding out until seventy for a larger check can mean spending the healthiest years watching the clock. A slightly smaller benefit a person actually gets to enjoy can beat a larger one claimed at an age when the knees no longer cooperate.

The second is needing the money. Someone laid off at sixty-three with a thin savings account, or a person whose body can no longer do the work, may not have the luxury of bridging eight years on other income. For them, claiming early is not a mistake. It is the floor doing exactly what a floor is for, and the unhurried "just

wait" advice is a luxury of people who have other money to live on in the meantime.

The third is health. Break-even math only works if a person knows how long they will live, and nobody does, but a serious diagnosis or a family history of short lives tilts the odds. If the years ahead are genuinely limited, collecting a benefit for more of them is a reasonable call. The opposite is also true: someone whose parents are still arguing at ninety should think hard about waiting.

The fourth is the earnings test, and it cuts in favor of the lower earner. Claiming before full retirement age while still working withholds a dollar of benefits for every two dollars earned above $24,480 in 2026. A person working part-time and staying under that cap gives nothing back, and the money withheld from higher earners is not even lost, since the agency credits it back later. The test punishes the still-working high earner and leaves the modest earner alone.

> That is the situation Dave Thompson is in. At sixty-four he is still doing consulting work, and if he claimed now, much of the benefit would be withheld anyway.

The fifth is being the lower earner in a marriage. The most valuable move in couples' planning is usually to have the higher earner wait, because their benefit becomes the survivor's benefit. The lower earner's delay does not protect the survivor in the same way, so the lower earner is often free to claim early, bringing income into the household while the higher earner's

benefit grows in the background. If a spouse out-earned you, starting your own check sooner may be the sensible half of a two-part plan. The underlying principle is simple: protect the income a retiree needs, so they can free up the income they want.

When Two Records Have to Talk to Each Other

For married couples, the claiming decision stops being one decision and becomes two that must be made together, because the timing of one spouse's claim reaches into the other's lifetime income. This is the part of Social Security most often handled by accident, and the accidents are costly.

The hinge is the survivor benefit. When one spouse dies, the survivor does not keep collecting both checks: they keep the larger of the two, and the smaller one stops. A household living on two Social Security checks becomes a household living on one, the bigger one, at exactly the moment its expenses rarely fall by half. Which check survives, and how large it is, was largely decided years earlier, by when the higher earner claimed.

That single fact drives the most valuable move in spousal planning: for most couples, the higher earner should delay, often all the way to seventy. Doing so does two jobs at once. It raises that spouse's own benefit for life, and it raises the survivor benefit the other spouse will fall back on, a benefit that lasts as long as either of them is alive. Delaying the higher earner's claim is the one decision that keeps paying after one of them is gone.

Dave and Mary Thompson are a textbook case. Dave is the higher earner; suppose his benefit at full retirement age is about $2,900 a month. By waiting until seventy he lifts it to roughly $3,596, and that larger figure becomes the floor under Mary's income for the rest of her life if she outlives him, which, as the younger spouse, she statistically will. Because Dave is still consulting and would run into the earnings test anyway, waiting costs him very little in the meantime.

Mary, the lower earner, does not need to wait, and waiting would not help the survivor calculation. She can claim her own benefit earlier to bring income into the household while Dave's grows. One subtlety trips up even careful planners: delaying lifts the survivor benefit, but it does not lift the spousal benefit. A spousal benefit is capped at 50 percent of the higher earner's full-retirement-age amount, no matter how long that spouse delays. So once Dave files, Mary may add a small spousal top-up, but the number that truly changes her life arrives later, as a survivor, and that one reflects every dollar of Dave's delay.

A mechanical note that catches people: under current rules, someone who files is generally treated as filing for every benefit they qualify for at once. The old maneuvers, claiming a spousal benefit while letting one's own benefit grow, or the "file and suspend" that once let a spouse collect early, were closed off by Congress for nearly everyone born after the mid-1950s. For most couples today the planning is simpler, and it

comes down to sequencing: who claims early for cash flow, and who delays to protect the survivor.

> The higher a couple's net worth, the more freedom they have to play it this way. Robert and Susan Whitfield do not need Social Security to pay their bills at all; for them the decision is almost purely about survivor protection and taxes. Robert, the higher earner, can delay to seventy and bridge the gap from their own accounts, locking in the largest possible survivor benefit for Susan while the delay also trims his later required withdrawals.

The Question Every Teacher Asks

There is one group for whom this used to carry an asterisk, and Mary belongs to it on paper: public employees with a government pension. For decades, two rules reduced Social Security for people who had also earned a pension from work that did not pay into the system. Those two rules confused and angered millions of teachers, firefighters, and police officers. As of 2025, both are gone. That is worth explaining, because a reader who spent a career in public service may still be planning around rules that no longer exist.

The first rule was the Windfall Elimination Provision, or WEP. It reduced a public employee's own Social Security benefit, the one earned from other, covered jobs, if they also drew a pension from non-covered work. The second was the Government Pension Offset, or GPO, and it was harsher: it cut any spousal or survivor benefit by two-thirds of the government pension, which for many people erased that benefit entirely. About

three out of four people hit by the GPO saw it reduced all the way to zero.

The Social Security Fairness Act, signed in January 2025, repealed both. The repeal reaches back to January 2024, and the Social Security Administration moved fast: by the middle of 2025 it had sent more than three million people over $17 billion in higher benefits and back pay. Anyone whose spousal or survivor benefit had once been zeroed out by the GPO, and who therefore never bothered to apply, should apply now, because the offset that blocked them is gone.

Whether any of this touches a particular person depends entirely on one question: did the public job pay into Social Security?

> Here Mary's case turns out to be reassuringly ordinary. Virginia is one of the states whose teachers do pay into Social Security, so a Virginia teacher earns a state pension and a Social Security record side by side, and WEP and GPO never applied to Mary in the first place. Her own benefit, and any survivor benefit from Dave's record, are both hers in full.

But that is a feature of Virginia, not of teaching. In roughly fifteen states, large groups of teachers and other public workers belong to systems that never paid into Social Security, and there the old rules used to fall hard. The takeaway is the same for everyone: before building a plan around a public pension, confirm whether that job was covered by Social Security, and know that, covered or not, the WEP and GPO reductions

that used to complicate the answer are no longer on the table.

Flying Solo

Single people face a cleaner version of the decision, though not necessarily an easier one. With no spouse to protect, the survivor benefit drops out of the math entirely, which removes the single strongest argument for delaying. What is left is the person's own longevity and their own resources.

> That can cut both ways. Back to Linda: because she has no second income and no spouse to fall back on, the case for buying herself a larger lifetime benefit is, if anything, stronger than it is for a married person. She is her own safety net, and a bigger guaranteed check is the most reliable net there is. She is also still working at fifty-eight, so claiming before full retirement age would run her straight into the earnings test. For Linda, leaning toward a later claim does double duty: it sidesteps the earnings test now and buys more longevity insurance for later. Because she was married more than ten years before her divorce, she could in principle claim on her ex-husband's record, but only if that benefit beat her own, and hers is higher, so it changes nothing.

The Two Escape Hatches

The decision is mostly irreversible, and that word carries weight, because two narrow do-overs exist. Neither is a casual undo; both come with conditions.

The first is withdrawal. Within twelve months of first claiming, a person can formally withdraw the application, repay every dollar of benefits received, and reset as if they had never claimed, which allows a larger benefit later. It can be used only once in a lifetime, and the repayment comes all at once, so it is a remedy for a recent decision that quickly proved wrong, not a long-term strategy.

The second is suspension, and it is the one Walter Higgins used. A person who has reached full retirement age but is not yet seventy can ask Social Security to suspend the checks; the benefit then grows by those same delayed retirement credits until it restarts, automatically at seventy or sooner if the person asks. Walter claimed at sixty-six in 2008, was pulled back to work when the crash gutted his savings, and suspended rather than keep drawing a benefit he did not yet need, turning a too-early claim into a second chance. But suspension exists only in that narrow window between full retirement age and seventy. Outside it, there is no reset.

The summary is this: it is far better to get the first claim right than to count on undoing it. The escape hatches exist, but they are small, and most claiming mistakes qualify for neither.

Before Filing

So what does a deliberate decision look like? It begins with information most people have never looked at. Every worker can pull a personalized Social Security

statement, online in a few minutes at ssa.gov or by mail, and it lays out the estimated benefit at sixty-two, at full retirement age, and at seventy: the three numbers the whole decision turns on. The statement is free and takes only minutes to retrieve, and no plan can be built on a benefit nobody has bothered to look up.

With those numbers in hand, the questions fall into a manageable order. How long do the family's people tend to live, and how is this person's health? Is there other money to live on while a benefit grows, or is the income needed now? Is the person still working, and would the earnings test eat into an early claim? And, for anyone married, what happens to the survivor, since that is the decision that outlives both spouses? A married couple that gets nothing else right can still do enormous good simply by having the higher earner wait.

None of these questions has a universal answer, which is exactly the point. The worst Social Security decision is the one made on autopilot, by defaulting to sixty-two because the door happened to be open, or by waiting to seventy because someone said waiting is always smart. Neither is a strategy. Both are reflexes. Made deliberately, with the statement on the table and the survivor in mind, Social Security can be the steadiest income of a person's life. Made by accident, it is a haircut that will not grow back.

And there is one more decision waiting right behind this one, on a tighter clock. Around the same age these questions come up, a person turns sixty-five and meets Medicare, with its enrollment deadlines, its penalties for signing up late, and a set of premiums that climb

with income in ways a claiming decision can quietly trigger. Medicare comes next.

Chapter 6: Medicare — Four Parts, Two Paths, and One Hard Deadline

Most people picture Medicare as the easy part of retirement. A person turns sixty-five, a red-white-and-blue card arrives in the mail, and the long worry about health coverage is finally over. That picture is comforting, and it is wrong in three expensive ways. Medicare is not automatic for everyone. It is not free. And several of its deadlines, once missed, attach a penalty that is not a one-time fee but a surcharge added to every premium for the rest of a person's life. The gap between what people assume about Medicare and what is actually true is where some of the costliest retirement mistakes take root.

Medicare runs on a tighter clock, and this is it. Around the same time the Social Security questions come up, a person turns sixty-five and meets a program with its own enrollment windows, its own penalties, and its own fork in the road that, once taken, can be hard to walk back. The good news is that the whole system rests on a small number of decisions, and none of them is difficult once the pieces are laid out. Here are the parts, what they cost, when to sign up, and the one big choice that shapes everything after it.

Four Parts, Not One

Medicare is not a single plan. It is four parts, labeled with the first four letters of the alphabet, and the labels are worth memorizing because the rest of the

system is built on them. Part A is hospital insurance: inpatient stays, skilled nursing care after a hospitalization, hospice, some home health. Part B is medical insurance: doctor visits, outpatient care, lab work, preventive services, and durable medical equipment. Together, Parts A and B are what most people mean by “Original Medicare,” and they are administered directly by the federal government.

Part C and Part D come from private insurance companies that Medicare approves and oversees. Part C, better known as Medicare Advantage, is an all-in-one alternative that bundles Parts A and B, and usually drug coverage, into a single plan from a private insurer. Part D is stand-alone prescription drug coverage, the piece Original Medicare does not include on its own. Part D and the cost of drugs get their own full treatment; the focus here is the larger architecture.

Part	What it covers	Who runs it	What it costs in 2026
A	Hospital insurance: inpatient stays, skilled nursing after a hospital stay, hospice, some home health	Federal government	Premium-free for most (40+ work quarters); $1,736 hospital deductible per benefit period
B	Medical insurance: doctor visits, outpatient care, lab work, preventive care, equipment	Federal government	$202.90/month standard premium; $283 annual deductible; then 20% of most costs, with no cap

Part	What it covers	Who runs it	What it costs in 2026
C	Medicare Advantage: an all-in-one private plan bundling A, B, and usually D, often with extras	Private insurers (Medicare-approved)	Varies; many plans add $0 beyond the Part B premium; networks and copays apply
D	Prescription drug coverage	Private insurers (Medicare-approved)	Varies widely by plan; many under $50/month, some $0

Two facts on that table do most of the damage when people miss them. Part A is usually premium-free, because most workers paid for it through decades of payroll taxes, but Part B is not. The standard Part B premium in 2026 is $202.90 a month, and it climbs higher for high earners. And Original Medicare, for all its strengths, covers only about eighty percent of the cost of covered services. It pays its share, the patient pays the other twenty percent, and here is the part that surprises people: there is no annual limit on that twenty percent.

The Eighty-Percent Problem

That missing cap is the single most important thing to understand about Original Medicare. Most insurance from a job has a maximum a person can be made to pay in a bad year, after which the plan covers everything. Original Medicare has no such ceiling. The twenty-percent coinsurance applies to every covered service with no stopping point, which means a serious

illness (a long hospitalization, a course of expensive infusions, a major surgery with complications) can run the patient's share into the tens of thousands of dollars. Eighty-percent coverage sounds generous until the bill is large enough that the remaining twenty percent is a catastrophe.

This is why almost no one should stop at Parts A and B alone. The remaining twenty percent is the gap, and filling it is the real Medicare decision. There are two ways to do it, and they lead down two very different roads.

Two Paths

The first path keeps Original Medicare and adds two things to it: a Medicare Supplement policy, almost always called Medigap, to cover the gaps, and a stand-alone Part D plan for drugs. Medigap is private insurance whose only job is to pay the deductibles, coinsurance, and copayments that Original Medicare leaves behind. Buy a comprehensive Medigap policy and the twenty-percent problem largely disappears: the patient pays a predictable monthly premium and very little else, and can see any doctor or hospital in the country that accepts Medicare, with no networks and no referrals.

The second path replaces Original Medicare with a Medicare Advantage plan. Instead of the government paying providers directly, a private insurer takes over, bundling hospital, medical, and usually drug coverage

into one plan, frequently with a low or even zero monthly premium beyond Part B, and often with extras Original Medicare never offers, such as dental, vision, hearing, and fitness memberships. In exchange, the plan runs on a network of doctors and hospitals, may require referrals to see specialists, and can require prior approval before it covers certain care. For anyone who carried employer health coverage during their working years, the arrangement will feel familiar: a Medicare Advantage plan works much like that group plan (a lower premium paid month to month, care delivered through a network, a copay at each visit, and a cap on the year's costs), and the insurer behind it is, in many cases, the very same company. Importantly, every Medicare Advantage plan does carry an annual out-of-pocket maximum, the ceiling Original Medicare lacks. In 2026 that ceiling can be no higher than $9,250 for in-network care, and the typical plan sets it closer to $5,400.

Within Medicare Advantage there are two common shapes. A Health Maintenance Organization, an HMO, keeps costs down by confining care to its network and usually requiring a referral to see a specialist. A Preferred Provider Organization, a PPO, costs a little more but allows care outside the network at a higher share of the cost. The trade in both cases is the one that runs through the entire decision: lower, more predictable cost in exchange for less freedom, or more

freedom in exchange for higher or less predictable cost.

	Original Medicare + Medigap	Medicare Advantage
Monthly cost	Higher: Part B premium plus a Medigap premium plus a Part D premium	Lower: often just the Part B premium; many plans add $0
Doctors and hospitals	Any provider in the U.S. that accepts Medicare; no networks	A network of providers; out-of-network costs more or isn't covered
Referrals and approvals	None	Often required, especially in HMOs
Out-of-pocket cap	Medigap keeps costs small and predictable; little gets through	A hard annual cap ($9,250 max in-network in 2026; often lower)
Extra benefits	Not included; bought separately	Dental, vision, hearing often bundled in
Prescription drugs	A separate Part D plan	Usually built in
Switching later	Can be denied for health reasons after the initial window	Can change every year during open enrollment
Best for	Predictability, travel, and freedom of choice	Lower premiums, extra benefits, comfort with a network

Neither path is the right answer for everyone, which is exactly why the choice exists. The Whitfields, a fictional couple, have ample

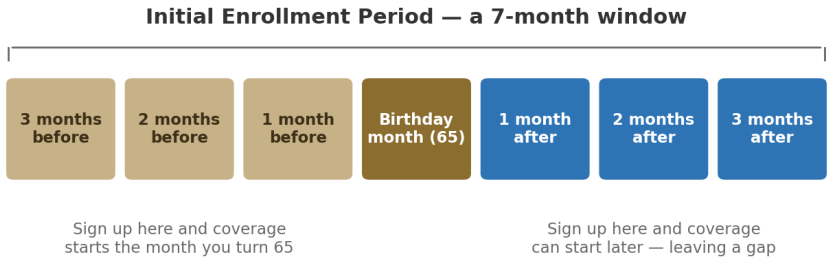
resources and travel often; they took the first path, Original Medicare with a comprehensive Medigap policy, because they value seeing any specialist anywhere without a referral and want their health costs to be a fixed, knowable line in the budget. For them the higher premium buys precisely what they want: certainty and freedom. Someone in good health on a tighter budget, comfortable using a local network of doctors, might reasonably reach the opposite conclusion, taking a low-premium Advantage plan and accepting its rules in exchange for the extras and the lower monthly cost.

One row of that table deserves its own warning, because it turns a reversible-looking choice into a nearly permanent one. A person can switch from one Medicare Advantage plan to another every year. But switching from Medicare Advantage back to Original Medicare with a Medigap policy is a different matter: outside a one-time window early on, a Medigap insurer in most states can look at a person’s health and charge more, or refuse to cover them at all. Walter Higgins learned this the hard way. He chose an Advantage plan at sixty-five for its low premium, stayed healthy for years, and later wanted to move to Medigap so he could see a specialist out of state without a fight, only to find that, with a heart condition now on his record, no Medigap insurer

in his state would take him at a price he could afford. The choice he made at sixty-five had become the last one he could freely make. A companion trap runs in the other direction. Moving from Original Medicare and a Medigap policy into an Advantage plan ends a stand-alone drug plan automatically, because a person cannot hold both at once, but it leaves the Medigap policy untouched. That policy has to be cancelled by the person, directly with the insurer; nothing cancels it automatically. And because a Medigap policy is barred from paying anything toward an Advantage plan, a supplement left in place keeps drawing a premium for coverage that can no longer pay a cent, a recurring waste until someone notices it.

The Deadline That Doesn’t Wait

All of this assumes a person enrolls on time, and on-time is a narrower target than most people realize. The main window is the Initial Enrollment Period: a seven-month span that opens three months before the month a person turns sixty-five, includes the birthday month, and closes three months after. Sign up in the first three months and coverage begins the month of the sixty-fifth birthday; wait until the birthday month or later, and coverage can be pushed back by a month or more, opening a gap exactly when it is least wanted.



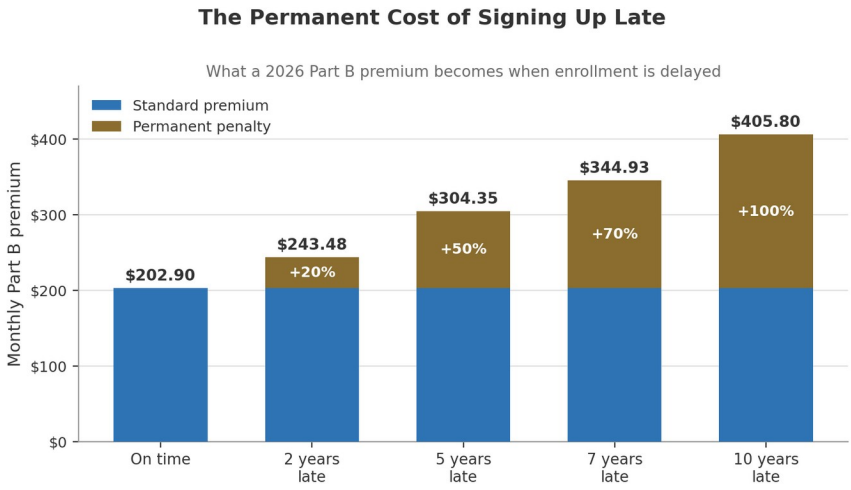
The Initial Enrollment Period. Signing up in the first three months keeps coverage from starting late.

Whether a person has to do anything at all depends on Social Security. Someone already collecting Social Security at sixty-five is enrolled in Parts A and B automatically. The card simply arrives. But someone who has delayed Social Security, as the higher earner in many couples should, is signed up by no one. Robert Whitfield, who is delaying his own benefit to seventy, will get no automatic enrollment and no reminder; he has to act on his own, three months before sixty-five, or risk a penalty for a deadline nobody told him was running. The very strategy that maximizes his Social Security check removes the safety net that would otherwise have enrolled him in Medicare.

The one situation that genuinely allows a person to delay Part B without penalty is active coverage through a current employer, their own or a spouse’s, at a company with twenty or more

employees. That exception is real, and it is also narrower than people think. It does not cover retiree health benefits. It does not cover COBRA. And it does not cover self-employment with an individual policy. Dave Thompson, still consulting at sixty-four, assumes that because he is working and insured he can put Medicare off. But his consulting income is self-employment, and the individual plan he buys is not active employer coverage, so when he turns sixty-five, the clock applies to him in full. If he waits, the penalty waits with him.

And the penalties are why the deadline matters. Miss the Part B enrollment window without a qualifying exception, and the premium rises by ten percent for every full twelve-month period a person could have enrolled but didn’t, not once, but added to every monthly premium for as long as they have Part B. The surcharge is recalculated off the premium each year, so it grows over time rather than fading. Two years late adds twenty percent; five years late, fifty; the cost of a forgotten deadline compounds steadily for the rest of a person’s life.



The Part B late-enrollment penalty is permanent and grows with each year of delay.

Part D carries its own late penalty, a smaller monthly surcharge, and it, too, is permanent. A person who misses the initial window entirely still has a yearly fallback, the General Enrollment Period that runs from January through March, but coverage bought then can start late and any penalties may already have attached. The lesson under all of it is the same: with Medicare, the cheapest move is almost always to sign up on time.

Where Medicare Meets Social Security

Medicare and Social Security are separate programs, but they are wired together in ways that catch people off guard. The most direct connection is the Part B premium: for anyone already collecting Social Security, it is deducted automatically from the monthly benefit, so the check that arrives is already net of it. For

someone who has delayed Social Security, like Robert, there is no check to deduct from yet, so Medicare sends a bill directly, usually every three months, and it has to be paid like any other. Delaying Social Security does not delay the Medicare premium.

A second surprise catches married couples: Medicare does not work like a family health plan. There is no “spouse and dependents” coverage. Each person enrolls individually, on their own sixty-fifth birthday, and pays their own premium. Mary Thompson is younger than Dave, so when Dave turns sixty-five and moves onto Medicare, Mary is not covered by it. She needs her own coverage, whether from a marketplace plan or elsewhere, to bridge the years until she reaches sixty-five herself. Couples who assume one spouse’s Medicare will cover the other can find a dangerous gap in the middle.

Mary’s case adds one more piece worth knowing, because it rewards a little foresight. During her working years she built up a health savings account, the rare account that goes in untaxed, grows untaxed, and comes out untaxed when it is spent on medical costs. Medicare changes how that account works in two ways. First, once a person enrolls in any part of Medicare, new contributions to a health savings account must

stop; it can no longer be added to, though whatever is already in it stays fully available. Second, and more usefully, that balance does not go to waste: it can be spent tax-free on most Medicare costs, including the Part B, Part D, and Medicare Advantage premiums that otherwise come straight out of pocket or out of a Social Security check. The one notable exception is a Medigap premium, which a health savings account cannot cover tax-free. For someone like Mary, years of steady saving become a ready source of tax-free dollars for the very Medicare premiums in question.

Finally, for higher-income households, both the Part B and Part D premiums climb with income through a surcharge called IRMAA, the income-related monthly adjustment amount, calculated from a tax return two years earlier. For the Whitfields, that means the Social Security and retirement-account decisions they make in their early sixties can quietly raise their Medicare premiums later. IRMAA is its own subject, and it deserves close attention; for now it is enough to know that with Medicare, as with so much of retirement, this year’s income echoes forward into future costs.

Before Sixty-Five

A deliberate approach to Medicare, like a deliberate approach to Social Security, starts before the deadline rather than at it. Three months before the sixty-fifth birthday is the moment to act, and a short set of questions organizes the whole decision. Is the person already collecting Social Security, or will they need to enroll on their own? Is there genuine active-employer coverage that allows a penalty-free delay, or does the clock apply now? Which path (Original Medicare with Medigap, or Medicare Advantage) fits this person's health, budget, doctors, and tolerance for networks? And because the first choice can be the hardest to undo, does it leave room for the life they expect to live, including the travel and the specialists they may want a decade from now? None of these questions has to be faced alone; a licensed Medicare broker can walk a person through every one of them at no cost.

Answer those, enroll on time, and the hardest parts of Medicare are behind a person. What remains is the piece set aside: the drug coverage under Part D, where the plans are many, the costs are uneven, and a recent and sweeping change has finally put a ceiling on what anyone can be made to pay in a year. Part D comes next.

Linda Castellano, at fifty-eight, is still several years from any of this. But the same clock bearing down on the Thompsons and the Whitfields is coming for her too, and the decisions will be

easier for having watched them play out. The worst Medicare mistakes are made by people who never realized a decision was being made at all. The reader, at least, will not be among them.

Chapter 7: Part D — Drug Coverage, the New Ceiling, and the Mistakes That Cost the Most

Of Medicare's four parts, Part D is the one people are most tempted to treat as optional. It does not cover hospital stays or doctor visits; it covers prescriptions, and a person taking nothing at sixty-five can be forgiven for thinking it is a bill worth skipping. That instinct is a trap. Part D is the piece Original Medicare leaves out entirely, the late penalty for going without it is permanent, and the rules governing it just changed more than they have at any point since the benefit began in 2006. A confusing gap that punished people in the middle of the year is gone, and in its place stands something Part D never had before: a hard ceiling on what anyone can be made to pay.

One piece of Medicare was set aside for closer study. This is that piece. The drug benefit has its own way of being obtained, its own costs, its own penalty, and a recent overhaul that every person approaching sixty-five should understand before choosing a plan.

Two Doors to the Same Coverage

There are two ways to get Part D, and they follow directly from the two paths through Medicare. The first is a stand-alone drug plan attached to Original Medicare, the route taken by anyone who keeps Original Medicare and a Medigap policy (the terms Medigap and Medicare Supplement are interchangeable), including the Whitfields, a fictional couple. The second is to let a Medicare

> Advantage plan carry it: most Advantage plans, the network-based option, fold drug coverage in, which is why Walter Higgins, on an Advantage plan since sixty-five, already has it built into the plan he holds. Not all of them do, though: some Advantage plans are sold without drug coverage, designed for people who already hold creditable coverage from another source, such as the Department of Veterans Affairs, and these are frequently marketed to veterans. As a rule, a person uses one door or the other, and the rule has teeth: enrolling in a stand-alone drug plan while on an Advantage plan automatically drops the person out of that Advantage plan and back to Original Medicare. That holds even for the drug-free, veteran-oriented plans just described: adding a separate Part D plan to one of them can cost a person the Advantage plan itself (a couple of uncommon plan types aside).

Whichever door a person walks through, the plan behind it is run by a private insurer that Medicare approves and oversees. Each plan sets its own monthly premium, keeps its own list of covered drugs, the formulary, sorts those drugs into pricing tiers, and contracts with its own network of pharmacies. The federal rules underneath are identical from plan to plan; everything stacked on top of those rules is not. Two plans can cover the same medication and charge very different amounts for it.

The Donut Hole Is Gone

For most of Part D's history, its costs ran through four confusing stages, and the second of them was notorious. After a person's drug spending crossed a threshold, coverage suddenly thinned and out-of-pocket costs jumped, a stretch known as the coverage gap, or more commonly the "donut hole," where people who had been paying small copays abruptly found themselves paying far more until they spent their way out the other side. It surprised people every year, always somewhere in the middle of the year, and always when they could least afford it.

That structure no longer exists. The Inflation Reduction Act rebuilt the Part D benefit from the ground up; as of 2025 the donut hole was eliminated, and the benefit now moves through three straightforward phases instead of four. The change matters in a practical way: the old thresholds, the numbers from the donut-hole era, describe a system that is gone, and any plan summary or article still citing them is out of date.

Phase	What the enrollee pays	2026 figures
Deductible	Full negotiated price for covered drugs until the plan's deductible is met	Up to $615; many plans set it lower or at $0
Initial coverage	25% of the cost of covered drugs (copay or coinsurance)	Continues until out-of-pocket reaches the annual cap
Catastrophic	$0 for covered Part D drugs for the rest of the calendar year	Begins once out-of-pocket reaches $2,100

The Part D benefit in 2026: three phases, replacing the old four-stage design and its coverage gap.

The third phase is the heart of the reform. Once a person's out-of-pocket spending on covered drugs reaches the annual cap, $2,100 in 2026, they pay nothing more for covered Part D drugs for the rest of the calendar year. Part D had never had such a ceiling. Before it existed, a person on an expensive specialty medication could spend many thousands of dollars a year with no end in sight; now the spending stops at a known number. That number rises a little each year (it was $2,000 in 2025), so it is worth confirming the current figure rather than assuming a round one. Only a person's own out-of-pocket cost for covered drugs counts toward it; monthly premiums do not.

What a Plan Actually Costs

Premiums vary more than almost anything else in Medicare. Some Part D plans carry no premium at all; others run past a hundred dollars a month, with many sitting well below fifty. One number gets quoted so often that it is widely mistaken for the typical premium, $38.99 in 2026, but it is not what any particular person pays. It is the national base beneficiary premium, a benchmark the government uses to calculate subsidies and the late-enrollment penalty. The premium that actually lands on a person's statement is set by the plan they choose, not by that benchmark.

For higher-income households there is a surcharge on top. Just as income raises the Part B premium, it raises the Part D premium too, through the same income-related adjustment, IRMAA, calculated from a tax

return two years earlier. For the Whitfields, whose income sits well above the threshold, the drug-plan premium will not be the modest figure on the plan's brochure; it will be that figure plus a surcharge tied to income they earned two years before. IRMAA reaches across both Part B and Part D, and it is involved enough to warrant a chapter of its own, which is where the next one goes.

Two Wins Worth Knowing

Two of the law's changes are worth singling out because they help the people who need help most. The first is insulin. A month's supply of each covered insulin product is capped: the enrollee pays no more than the lowest of thirty-five dollars, a quarter of the plan's negotiated price, or a quarter of Medicare's Maximum Fair Price for that insulin. The deductible does not apply to insulin, so the cap holds from the first fill of the year rather than only after a separate hurdle is cleared.

The second is vaccines. Every adult vaccine on the federal recommended schedule, including the shingles vaccine, which used to cost some people around two hundred dollars, is covered with no cost-sharing at all under Part D: no copay, no deductible, even at an out-of-network pharmacy. For anyone weighing whether a recommended shot is worth the price, under Part D there is no longer a price to weigh.

Item	2026
Annual out-of-pocket cap	$2,100, then $0 for covered Part D drugs the rest of the year
Maximum plan	$615; plans may set it lower or at

Item	2026
deductible	$0
National base beneficiary premium	$38.99, a benchmark for penalties and subsidies, not a typical premium
Insulin, per month's supply	No more than the lowest of $35, 25% of the negotiated price, or 25% of Medicare's Maximum Fair Price; no deductible applies
Recommended adult vaccines	$0 cost-sharing, including the shingles vaccine
Late-enrollment penalty	1% of $38.99 per full uncovered month, added to the premium permanently; triggered by a 63-day gap without creditable coverage

Key 2026 Part D figures.

The Penalty for Waiting

Part D's late-enrollment penalty is separate from Part B's, and like Part B's, it is permanent. It is triggered by going sixty-three days or more without "creditable" drug coverage after the initial enrollment window closes. Creditable coverage means a drug plan expected to pay, on average, at least as much as standard Part D, the kind of solid coverage some people still have through an employer or a union. Coverage that falls short of that bar, or no coverage at all, starts the clock.

The penalty itself is small per month and relentless over time. It adds one percent of that national base premium (the $38.99 figure) for every full month a person went without creditable coverage, rounded to the nearest dime and added to the premium for as long as they hold Part D. Wait fourteen months and the

surcharge is fourteen percent, attached to the premium permanently and recalculated upward as the base premium climbs. It is the same design as the Part B penalty: a forgotten deadline that quietly bills a person for the rest of their life.

This is the trap the healthy walk into. A person who takes no medication looks at the premium for a plan they don't think they need and skips it, then develops a condition at seventy-two, signs up at last, and finds the penalty has been compounding the entire time. Dave Thompson is closer to this edge than he realizes. Still consulting and insured through an individual policy he buys himself, he assumes that being covered is enough; but an individual plan is not creditable employer coverage, so the Part D clock applies to him exactly as the Part B clock did. Mary takes almost nothing, yet she enrolls in a low-premium drug plan anyway, partly to keep the penalty from ever attaching, and partly because the health savings account she built in her working years can pay that premium with tax-free dollars. A few dollars a month now is cheaper than a surcharge forever.

Spreading the Cost

A new option, available since 2025, changes when a person pays rather than how much. The Medicare Prescription Payment Plan lets an enrollee spread their out-of-pocket drug costs into level monthly installments across the calendar year, billed by the plan, instead of paying the full amount at the

pharmacy counter. There is no interest and no fee to use it.

The critical thing to understand is what it does not do. It does not lower the total cost of anything: it is budget-smoothing, not a discount, reshaping the timing of the bill while leaving its size unchanged. Its value shows up for someone facing a large charge early in the year. Walter, eighty-four and managing a heart condition, could fill an expensive medication in January and run through much of the annual cap in a single trip; the payment plan turns that one steep bill into twelve smaller, predictable ones. A person already enrolled is carried into the next year automatically unless they opt out, but switching to a different plan means opting in again with the new one, a small step easy to forget in the shuffle of changing plans.

Why Prices Are Moving

Behind these changes runs a larger one. For the first time, Medicare can negotiate the price of certain high-cost drugs directly with the manufacturers, and the first ten negotiated prices took effect on January 1, 2026. The drugs chosen are among the most widely used in Part D, and several treat heart disease and diabetes, the very conditions common at Walter's age, so a person taking one of them may see a lower price at the counter this year than last. How much any single drug saves depends on the drug and the plan and is best checked against a current source rather than assumed, but the direction is the point: the ground under Part D pricing is shifting, and for some of the medications that matter most, it is shifting downward.

Choosing — and Re-Choosing — a Plan

The cheapest plan on paper is not always the cheapest plan in practice. What determines a person's real cost is not the premium alone but the total, premium plus deductible plus what their specific drugs cost on that plan's formulary and tiers, filled at the pharmacy they actually use. A plan with no premium can end up costing more than a forty-dollar plan if it places a needed medication on a high tier or leaves it off the formulary entirely. The premium is the part people compare; the drug costs are the part that decides it.

Sorting that out is less daunting than it sounds, because the government provides the tool. Medicare's Plan Finder, at medicare.gov, lets a person enter their exact medications and preferred pharmacy and see the projected annual cost of every plan available to them, ranked. Entering the actual drug list is the single most useful step in the whole decision; it replaces guesswork about premiums with a real comparison of what the year will cost. And the same discipline applies to doctors, not only drugs. Because an Advantage plan delivers care through a network, anyone for whom a particular physician or hospital matters should confirm those providers are in the plan, and check it every year, since a plan's network, like its formulary, can change and drop a familiar doctor without warning. For a person with no such attachment the point weighs less, and under Original Medicare it does not arise at all, since there is no network to leave.

The do-it-yourself tool is not the only route, and for many people it is not the best one. The same decision

(Original Medicare with a supplement, or Medicare Advantage, and which drug coverage sits underneath) can be handled with a licensed Medicare insurance broker, and doing so costs the beneficiary nothing directly: federal rules bar an agent from charging a fee to help someone enroll, and the premium is identical whether a plan is arranged through a broker or straight from the insurer. One practical thing is worth knowing at the outset, because not everyone who offers Medicare plans offers the same thing. An independent broker is not an employee of any insurer; appointed with many carriers, they can compare across most of the market. A captive agent (or a company's salaried or subcontracted sales representative) works for a single insurer and can present only that company's plans, and those representatives are usually who answer the inbound numbers behind the heavy national Advantage advertising the big carriers run on television and in the mail, the kind of campaign an independent broker rarely matches. A person who calls a number from an ad is therefore reaching one company's sales desk by design, not a comparison across the market, so it is fair to ask anyone offering a plan which carriers they actually represent. Within that scope, a broker turns the whole Medicare decision, not just the drug plan, into a guided one at no cost.

What a beneficiary gets in return is expertise that is unusually well-regulated. To offer these plans at all, an agent must be licensed by the state and then, every single year, pass the Medicare and fraud-waste-and-abuse certification that federal rules require before anyone may market coverage (a fifty-question exam

scored at ninety percent to pass, rewritten each year as the rules change), and, on top of that, complete a separate training and test for each individual carrier whose plans they offer, also every year. Their conduct runs under CMS marketing rules strict enough that breaking them can cost an agent both the state license and the carrier contracts that make the work possible, and those rules are designed to keep a recommendation focused on the coverage that fits a beneficiary's health, budget, and timeline. Nothing comparable is required to offer ordinary life or health insurance; this annual re-credentialing is unique to Medicare, and an agent who renews the full set season after season is, by the very structure of the requirement, someone who has chosen to stay current in a field that demands it.

And it is not a one-time exercise. Plans change every year: premiums rise or fall, formularies are rewritten, drugs move tiers, pharmacies drop in or out of the network, and insurers add and discontinue whole plans. The plan that fit best this year may be a poor fit next year, often with no change at all in the person's own prescriptions. Each fall the plan mails an Annual Notice of Change describing what will be different, and the open-enrollment window from October 15 to December 7 is the time to act, the natural moment to rerun the comparison, alone through the Plan Finder or alongside a broker who does it every year for a living. One distinction matters here: a Part D or Medicare Advantage plan can be changed each year with no health questions asked, but a Medigap policy is different, since outside its early window a switch can require medical underwriting, so the freedom to shop

every year is real for drug and Advantage coverage and more limited for a supplement. The most expensive Medicare mistake is rarely choosing the wrong plan once; it is choosing a plan and then never looking again, letting it renew on autopilot while a better-fitting, cheaper option sits one comparison away.

Before the Next Decision

Part D is private coverage, so its surface details differ from plan to plan and from one person to the next. But the floor beneath all of it is sturdier than it has ever been: a hard annual ceiling on out-of-pocket cost, insulin held to a fixed monthly cap, recommended vaccines at no charge. The decisions that remain are within any person’s reach: enroll on time so no penalty ever attaches, match the plan to the medications actually taken rather than to the lowest advertised premium, and revisit that match every single year.

One cost was named and then deferred: the income surcharge that raises both the Part B and the Part D premium for higher earners. It has now surfaced twice without being fully explained, and it has earned a chapter of its own. For the Whitfields, and for anyone whose income approaches the thresholds where it begins, IRMAA is the next piece of the picture, a stealth tax on Medicare that turns the financial decisions of one year into higher premiums two years later. IRMAA comes next.

Chapter 8: Beyond the Premium — IRMAA, Long-Term Care, and Medicare's Gaps

Medicare's enrollment decisions, once made, feel like the end of the matter: the parts chosen, the premium set, a path taken. For a higher-income retiree, and for anyone who lives long enough to need help with the ordinary business of daily life, it is not the end. Two kinds of cost wait beyond the premium: one that pushes it higher, and one that Medicare never agreed to pay at all. Both can drain a retirement income, and both reward the person who sees them coming.

The first cost is a surcharge with an unlovely name: the income-related monthly adjustment amount, known by its initials as IRMAA. The second group of costs are gaps, the everyday care Medicare leaves out, and the long-term care it leaves out almost entirely. None of these is exotic, and none is rare. What ties them together is that each is a cost the headline premium hides, and each connects directly to the income and tax decisions that shape the rest of the plan.

IRMAA: The Surcharge That Income Triggers

For most people, the Part B premium is a single, standard number, $202.90 a month in 2026. For higher earners it is not. Above a certain income, Medicare adds a surcharge to both the Part B premium and the Part D premium, and the surcharge climbs in steps as income rises. It is not a tax in the ordinary sense, and the Social Security Administration, not the IRS, administers it, but it behaves like one: the more retirement income a person draws from all taxable

sources, the more they pay to the government for the same coverage. Most beneficiaries never encounter it; only about seven percent pay it in a given year. For those who do, it returns every year their income stays high.

The mechanism that makes IRMAA so easy to stumble into is its delay. The surcharge in any year is set by the income reported two years earlier: the 2026 premium is determined by the 2024 tax return. The figure Medicare uses is modified adjusted gross income, which for this purpose is simply adjusted gross income plus any tax-exempt interest, such as that from municipal bonds. That same lag explains why IRMAA so often lands hardest at the very start of retirement: the income that sets a newly retired person's first Medicare premium is the income from two years before, frequently a last full year of working salary, far above the modest retirement income that has since replaced it. A single strong year (a large withdrawal, a property sale, a Roth conversion) can raise a premium two years later, long after the money has been spent and the year forgotten.

Worse, the surcharge arrives as a cliff rather than a slope. There is no gradual phase-in: a single dollar of income over a threshold moves a person into the next tier and raises the premium by the full step. The brackets for 2026, set by 2024 income, run as follows.

Income (single filer)	Income (married, joint)	Monthly Part B	Part D surcharge
Up to $109,000	Up to $218,000	$202.90	None
$109,001–$137,000	$218,001–$274,000	$284.10	+$14.50
$137,001–$171,000	$274,001–$342,000	$405.80	+$37.50
$171,001–$205,000	$342,001–$410,000	$527.50	+$60.40
$205,001–$500,000	$410,001–$750,000	$649.20	+$83.30
Above $500,000	Above $750,000	$689.90	+$91.00

2026 monthly Medicare premiums by income, set by modified adjusted gross income from the 2024 tax return. The Part D figure is a surcharge added to whatever the chosen drug plan charges; both spouses on Medicare pay their own. Brackets are inflation-adjusted yearly.

The amounts are not trivial, and they compound across a household. The lowest surcharge tier adds roughly $1,150 a year per person; because IRMAA is charged to each enrolled individual, a married couple who both crossed the first threshold each pay it, for something near $2,300 between them. At the top of the scale the surcharge approaches $7,000 a year for one person. A couple does not split a bracket. They each land in it.

This is where Medicare and tax planning become one subject. The Whitfields, a fictional couple, hold a large traditional retirement account; the required withdrawals from it lift their income high enough that IRMAA is simply part of their

Medicare cost, every year. Roth conversions, required distributions, and capital gains each raise modified adjusted gross income, and each can trip the surcharge two years on. Inheritance can do the same, though not in the way people expect: receiving cash or property is not itself taxable income, but inheriting a traditional retirement account is, because its withdrawals are taxed, and a non-spouse heir must empty such an account within ten years. Linda Castellano, draining an inherited account against exactly that deadline, will face a stretch of unusually high-income years; and because she files as a single taxpayer, the thresholds that catch her sit at half a married couple's.

That same halving falls, without warning, on a surviving spouse. When one partner dies, the survivor often inherits the household's retirement accounts, so the income does not fall by half, but the next year's return is filed as a single person, against thresholds set at half the married level. The result is a quiet cruelty: the very income that kept a couple comfortably below the surcharge can push the widow or widower over it, even as the household's total income drops. A spouse's death is, at least, one of the life-changing events that allows the surcharge to be appealed.

There is a relief valve, but a narrow one. A person whose income has genuinely fallen (through retirement, a spouse's death, divorce, or the loss of a pension) can ask the Social Security Administration to

use a more recent year instead, by filing the life-changing-event form; this is precisely the remedy for the newly retired person whose first-year premium was set by a final year of salary. A voluntary choice does not qualify: selling stock or converting to a Roth on purpose is not a life-changing event, which is why those moves are planned around the brackets rather than appealed afterward. And IRMAA is not a life sentence; it is recalculated every year, so when income falls below a threshold, the surcharge falls away two years later, on the same delay that created it.

The practical lesson is that IRMAA turns the timing of income into a Medicare cost. It is the first place the income plan and the Medicare plan visibly collide, where the same dollar that shapes a tax bill also shapes a premium.

What Medicare Simply Leaves Out

The second kind of surprise is not a surcharge but an absence. Medicare was built in 1965 around hospital and physician care, and whole categories of ordinary health spending were left outside its walls and have stayed there: routine dental care, routine vision care, and hearing. The common assumption that a Medicare Supplement closes these gaps is wrong. Medigap pays the deductibles and coinsurance of Original Medicare, and nothing more. The dental, vision, and hearing gaps stay open under a supplement just as they do without one.

The specifics are stark. Routine dental (cleanings, fillings, extractions, dentures) is not covered; a single crown or implant is paid entirely out of pocket. Vision

is much the same: Medicare covers cataract surgery and the treatment of medical eye conditions, and even one pair of glasses after cataract surgery, but the ordinary eye exam and the eyeglasses most people need are not covered. Hearing follows the same pattern: a doctor-ordered diagnostic test is covered, but hearing aids and the exams to fit them are not, and a pair can run several thousand dollars. A single year that brings a crown, new glasses, and a set of hearing aids can cost several thousand dollars that neither Original Medicare nor a Medigap policy will touch. These everyday gaps, together with the larger ones, line up differently across the three ways of taking Medicare, and the whole picture is easier to weigh side by side.

Care or benefit	Original Medicare	Medigap	Medicare Advantage
Routine dental	Not covered	Not covered	**Often a limited allowance**
Routine vision (exams, glasses)	Not covered	Not covered	**Often a limited allowance**
Hearing exams and aids	Not covered	Not covered	**Often a limited allowance**
Routine foot care (podiatry)	Not covered (routine)	Not covered	**Sometimes included**
Annual physical exam	Not covered	Not covered	**Sometimes included**
Long-term / custodial care	Not covered	Not covered	Not covered
Outpatient prescription drugs	Not covered	Not covered	**Usually included**
Care outside the United States	Not covered	**Emergency (some plans)**	**Emergencies only**
Annual out-of-pocket limit	No annual limit	**Limits exposure**	**Annual cap**

What each path leaves uncovered. Medigap pays Original Medicare's deductibles and coinsurance, and most Medigap plans add limited foreign-travel emergency coverage, but none of the other benefits shown. Medicare Advantage often bundles limited dental, vision, hearing, and similar extras and usually includes drug coverage, yet covers no more long-term care than Original Medicare, which pays only for short-term skilled care after a qualifying hospital stay, never custodial care. "Routine" here means everyday maintenance; medically necessary care (a diabetic foot exam, a bunion, eyeglasses after cataract surgery) can

still be covered. Original Medicare also covers a free yearly wellness visit, a conversation-based prevention-planning appointment, not a hands-on head-to-toe exam, which many people mistake for the annual physical it does not cover. And on the final row the logic flips: Original Medicare alone sets no annual out-of-pocket maximum, while Medigap (by paying the cost-sharing) and Medicare Advantage (through a fixed yearly cap) both limit how far that exposure can run.

There are partial patches. Many Medicare Advantage plans fold in limited dental, vision, and hearing allowances, one of the genuine attractions of that path, and stand-alone dental and vision policies can be bought alongside Original Medicare. But allowances are capped and policies carry their own limits; these arrangements narrow the gap rather than close it. The category of care Medicare excludes most completely, and most expensively, is something else entirely.

The Largest Gap: Long-Term Care

The costliest thing Medicare does not cover is long-term care, and the misunderstanding on this point is nearly universal. Medicare pays for skilled care, and only briefly: up to one hundred days in a skilled nursing facility per benefit period, available only after a qualifying inpatient hospital stay of at least three days. The first twenty days are covered in full; days twenty-one through one hundred carry a daily charge of $217 in 2026; after that, the benefit ends. Medicare pays for recovery, not for residence.

The line that catches people is the one between skilled care and custodial care. Skilled care is rehabilitation, physical therapy after a stroke, wound care, the work of trained clinicians. Custodial care is help with the ordinary acts of living: bathing, dressing, eating, moving from a bed to a chair. That second kind is what most long-term care actually is, and Medicare pays

nothing toward it, whether it is delivered at home, in an assisted-living community, or in a nursing home. Even the three-day rule hides a trap: time spent in a hospital bed under “observation” rather than as a formal inpatient admission does not count, so a person can spend days in the hospital and still fail to qualify for the skilled benefit that follows. A Medigap policy changes none of this.

This is not a remote risk. About seventy percent of those who reach sixty-five will need some form of long-term care; on average a man will need it for roughly two and a quarter years and a woman for nearly four, and about one in five will need it for more than five years.

> Helen, Mary Thompson’s mother, is eighty-six and still lives independently, and stands, for the Thompsons, as a reminder that independence at her age is a stage of life rather than a certainty of one.

The care is expensive, too, and rising faster than ordinary inflation. The national median figures for 2026 tell the story; a projection forward twenty years, roughly the span of a long retirement, tells a sharper one.

Type of care	2026 monthly (median)	Projected 2046
In-home care (home health aide)	$6,673	~$14,600
Assisted living	$6,200	~$13,600
Nursing home, private room	$10,800	~$23,700

National median monthly costs, CareScout Cost of Care Survey, 2025; 2046 figures apply 4% annual compound growth, with the in-home figure based on 44 hours of care a week.

Someone has to pay these bills, and the options are fewer than people hope. Medicare pays almost nothing. Families fill much of the gap themselves (roughly half of all long-term care is provided unpaid by relatives), but that labor is not free either: it interrupts the caregiver's own income and savings, and it rarely removes the need for professional help entirely. The program that does pay for custodial care, Medicaid, requires a person to spend down to near-poverty first, with a multi-year look-back at money given away. Stripped to essentials, the choices are three: self-fund the care, insure against it, or impoverish into the program of last resort.

Long-term-care insurance is the tool built for the problem, and it works on its own logic. Benefits typically begin when a person can no longer perform two of six "activities of daily living" (bathing, dressing, eating, and the like) and usually only after an elimination period, a waiting stretch of weeks the family covers before coverage starts. A traditional policy reimburses care costs, often with inflation

protection and options for in-home care and for spouses. It does not pay for home modifications such as a stairlift or a bathroom remodel, and neither does Medicare. The product is specialized enough that in Virginia an agent must hold a license endorsement beyond the basic health license merely to discuss it, a fair signal of how much detail sits beneath the surface.

One feature of these policies deserves a plain warning, because it surprises many who hold them. The premium is set when the policy is issued and is designed to stay level. But only the coverage is locked in, not its price. The insurer cannot single out one policyholder, and cannot cancel anyone who keeps paying; what it can do is petition the state regulator (in Virginia, the State Corporation Commission’s Bureau of Insurance) to raise rates for an entire class of similar policies at once. And it has happened across the industry, sometimes steeply. The first generations of policies, sold in the 1990s and 2000s, were badly underpriced, because insurers misjudged how long people would keep them, how long they would live, and how little interest their reserves would earn. Increases of fifty percent or more have landed on older policies, and the federal employees’ program raised some premiums by as much as eighty-six percent in a single year. When an increase is approved, the policyholder is left to choose among paying it, trimming benefits to hold the premium down, or accepting a smaller paid-up policy.

Even when a policy performs exactly as designed, the path to a paid claim is rarely smooth, and one example shows why. A woman with a long-term-care policy falls

and breaks a hip. The hospital admits her; Original Medicare covers the inpatient stay, but only after she clears the Part A hospital deductible, $1,736 in 2026, which a Medigap policy would absorb and a Medicare Advantage plan would replace with copays of its own. The ambulance that brought her in may or may not be covered, depending on medical necessity and her plan. From the hospital she is discharged to a skilled nursing facility for rehabilitation, which Medicare covers in full for twenty days, then only in part, and then, once the skilled need ends, not at all. She comes home no longer needing skilled rehabilitation but unable to perform two of the six activities of daily living, and she files a claim on her long-term-care policy. The assessment can take up to a month, and while it runs she pays for home care herself. Once the claim is approved, a common elimination period, often sixty or ninety days, must still pass before benefits begin, and through all of it the bills are hers. And when the policy does pay, it usually pays by reimbursement: she covers the cost first and waits weeks to be repaid. That is the experience of someone who planned and bought coverage. The person without it simply pays for all of it.

A traditional policy is not the only way to fund the risk, and the alternatives each carry their own logic. Hybrid policies pair life insurance with a long-term-care benefit, paying a death benefit if the care is never needed and sometimes returning the premium, with pricing locked rather than open to class increases. Certain annuities carry a long-term-care benefit and can be obtained without medical underwriting, though the benefits may be taxable. A health savings account

built up during working years can pay for qualified care, and even some premiums, with untaxed dollars. A person with enough assets may instead earmark savings and invest them for the purpose, keeping the money liquid enough to reach when care is needed; in Virginia and many other states, the cash value of a life insurance policy or annuity carries statutory protection from creditors, which is one reason some people keep such earmarked dollars inside one of those contracts rather than a plain account, though the protection has conditions, varies by state, and is properly a question for an attorney. And there is family, which, as the numbers make plain, is rarely the free option it appears to be. The annuity and life-insurance tools named here each get a fuller treatment of their own; the point for now is simply that more than one route exists.

The stakes are not abstract. At the median price of a private nursing-home room, two spouses in care at once can run through more than a quarter of a million dollars a year, and it is not unusual for a couple who both need care to watch a seven-figure nest egg drain away into it. That is why, whichever route a person takes, the advantage belongs to those who plan early. Long-term-care coverage is cheapest and most attainable while a person is still healthy enough to qualify and young enough for a manageable premium; the planning is hardest to do well at the moment it is finally needed. Like the Medicare path itself, it is a decision best made deliberately, years ahead of any crisis.

The Costs Behind the Cost

IRMAA and the coverage gaps look like separate problems, but they rhyme. Both are costs the advertised premium conceals, and both are shaped less by Medicare's rules than by decisions made elsewhere in a retirement plan. IRMAA is governed by the shape of a person's income; long-term care, by how early and how deliberately its cost is funded. Neither is visible at the moment a person signs up for Medicare, and both can arrive years later with real force.

This is why the rest of the planning is not only about taxes. Every dollar of income that a required distribution forces or a Roth conversion creates does more than change a tax bill. It can move a Medicare premium two years on. And the income that a long stretch of care can consume is the very income those chapters are built to protect. Medicare's hidden costs are, in the end, retirement-income problems wearing a medical disguise.

> The Whitfields will plan their withdrawals with one eye on the IRMAA brackets. A single filer like Linda will watch thresholds that sit at half a couple's, as, one day, will a surviving spouse. And Helen's long, independent old age reminds the Thompsons that the largest bill of all may be not for coverage but for care. The worst of these costs fall hardest on the people who never knew to look for them; the reader who plans for the costs beyond the premium is the one least likely to be ambushed by them.

Chapter 9: Required Minimum Distributions — The IRS Wants Its Money

A traditional IRA, and the components of a 401(k) that have not already been taxed, are built on a deal with the government. In a 401(k), those untaxed components are the contributions taken from a paycheck before any tax is withheld, plus any employer match, which is never taxed going in. A traditional IRA works on the same principle: a contribution can be deducted from taxable income for the year it is made, which directly lowers that year's tax bill. In both cases the money goes in untaxed and grows untaxed for decades. The catch is in one word. The tax was deferred, not forgiven. The government always meant to collect it, just later. Required minimum distributions are how it does so. Starting at a set age, a portion of every tax-deferred account has to be withdrawn each year and reported as income, whether the retiree needs the money or not.

Two ideas already raised come together here. The income that pushes a retiree into a higher tax bracket, or past an IRMAA threshold, or that makes more of a Social Security check taxable, is often the required distribution itself. The rule is strict, but strict also means predictable. The starting age is fixed, the formula is published, and the math can be run years ahead. That makes a required distribution one of the few retirement costs a person can see coming well in advance, and one of the most expensive to ignore.

When Distributions Begin

Today, required distributions begin at age 73, but that age has moved over the years. It was once 70½, then 72, and the Setting Every Community Up for Retirement Enhancement (SECURE) 2.0 Act of 2022, signed into law that December, raised it to 73. For anyone born in 1960 or later, it rises again to 75, starting in 2033. The rule applies to traditional IRAs and to most workplace retirement plans: SEP and SIMPLE IRAs, 401(k)s, 403(b)s, and governmental 457(b)s. Inside a 401(k), it covers the entire balance, including the employer match. Roth IRAs are the exception. Because the money in them was already taxed, they require no distributions at all during the owner's lifetime, which is worth holding onto, because it becomes a tax-mitigation strategy in its own right.

What matters for this rule is the type of account, not what is invested inside it. A certificate of deposit, a mutual fund, a stock, or an annuity held inside a traditional IRA is subject to the required distribution, because the IRA itself is. The same holdings in an ordinary account are not. So the rule reaches traditional IRAs and employer plans, including whatever is invested within them, but it does not reach a regular taxable brokerage account, a bank savings account or CD held outside a retirement account, the cash value of a life insurance policy, or a Roth IRA during the owner's life. None of those carry a required distribution.

Annuities deserve a closer look because they fall on both sides of that line. An annuity held inside an IRA or an employer plan, called a qualified annuity, is part of

the required distribution system like everything else in the account; once it is turned into a stream of lifetime payments, those payments generally satisfy the requirement for that contract. An annuity bought with after-tax money outside a retirement account, called a non-qualified annuity, is different. It carries no required distribution during the owner's life. But its growth was deferred, not erased: when money is withdrawn, the gains come out first and are taxed as ordinary income, and only after the gains are used up does the original after-tax principal come back tax-free.

The first deadline comes with a trap. The very first distribution is not due until April 1 of the year after the one in which a person turns 73. Every distribution after that is due by December 31. Waiting until that April to take the first one does not buy any time, though, because the second year's distribution is still due that same December. Two full years of withdrawals then land in a single tax year, stacked on top of each other, which can push income into a higher bracket. It can also raise a Medicare premium through IRMAA two years later, because IRMAA is set by the tax return from two years earlier, not by current circumstances. The simplest way to avoid all of this is to take the first distribution in the year of turning 73 and stay on the December schedule from then on.

There is one exception to the starting age. Someone still working at 73, for the employer that sponsors their 401(k), and who owns no more than five percent of the company, can put off distributions from that one plan until they actually retire. The exception is narrow.

It applies only to the current employer's plan, not to an IRA, and not to a 401(k) left behind at an old job. Those accounts start at 73 no matter what, working or not.

How the Amount Is Calculated

The amount of each year's distribution comes from a simple formula. Take the account balance as of December 31 of the previous year and divide it by a number from an IRS table called the Uniform Lifetime Table. That number is the distribution period for a person's age. The result is the required distribution for the year. The current year's market does not factor in; only the prior December 31 balance does. And the whole calculation is redone every year, with a new balance and a new divisor.

The divisor shrinks a little each year as a person ages, so the required withdrawal becomes a larger percentage of the account over time. At 73, the distribution is about 3.8 percent of the balance. By 80 it is near 5 percent, by 85 past 6 percent, and by 90 above 8 percent. A rising percentage does not mean the balance itself necessarily falls. In a strong year, the account can earn more than the withdrawal takes out and still finish higher than it started. But as the required percentage keeps climbing with age, it tends to outpace investment growth, and eventually the balance does begin to come down. A few ages show how the percentage rises:

Age	Distribution period (divisor)	RMD as % of balance
73	26.5	3.8%
75	24.6	4.1%
80	20.2	5.0%
85	16.0	6.3%
90	12.2	8.2%
95	8.9	11.2%

Distribution period and resulting withdrawal rate at representative ages. Source: IRS Uniform Lifetime Table, Publication 590-B.

Walter Higgins, a fictional character, shows how this works in practice. Walter is 84, a retired postal worker, and he has been taking required distributions for years. His traditional IRA held about $340,000 at the end of last year. The distribution period for age 84 is 16.8, so his required withdrawal is roughly $20,240, which is $340,000 divided by 16.8. It is due whether he needs the money or not, and it is taxed as ordinary income at whatever rate the rest of his income sets.

The rules for combining accounts trip up a lot of people, and they differ by account type. With traditional IRAs, the separate required amounts can be added together and the whole total taken from any one of them, in any combination. Workplace plans are stricter. Each 401(k) has to be calculated and satisfied on its own, and an inherited IRA can never be

combined with an IRA the person already owns. The safe approach is to figure each account separately, then check what the law actually allows to be combined before moving any money.

The Penalty for Missing One

Missing a required distribution is one of the more expensive mistakes in retirement. The IRS charges an excise tax on whatever should have been withdrawn but was not. SECURE 2.0 lowered that penalty from a steep 50 percent to 25 percent, and to 10 percent if the shortfall is fixed within two years and reported on Form 5329. Even at the lower rate, it is a heavy price for a missed deadline, and it comes on top of the income tax that was already owed on the withdrawal.

Accounts That Are Exempt

Two kinds of account sit outside these rules entirely. The first is the Roth IRA. Its contributions were taxed before they went in, so those dollars come back out tax-free. The growth on top of them was never taxed, and as long as the account has been open for at least five years and the owner is at least 59½, that growth is never taxed either. Tax-free earnings, with no distribution ever required during the owner’s life, are what make the Roth so valuable. SECURE 2.0 extended the same treatment to Roth accounts inside workplace plans, ending required distributions from Roth 401(k)s and 403(b)s in 2024. The second exempt account is the health savings account, which also carries no lifetime distribution requirement.

Mary Thompson's Roth IRA and her HSA are good examples: nothing forces a dollar out of either one, which is part of what makes them so useful in a tax plan, and part of why deliberately shifting money toward the Roth side matters so much.

When an Account Is Inherited

The rules change completely when an account is inherited, and they got much tougher recently. For deaths after 2019, the Setting Every Community Up for Retirement Enhancement Act of 2019, the original SECURE Act, ended the old approach that let most non-spouse heirs stretch withdrawals across their own lifetimes. The replacement is the ten-year rule: a non-spouse beneficiary has to empty an inherited IRA completely by December 31 of the tenth year after the original owner's death.

That change rewrote what a traditional IRA is good for. For decades it doubled as a wealth-transfer tool: an heir could spread withdrawals across their own life expectancy, a feature known as the stretch IRA, letting the account keep growing tax-deferred for years. The ten-year rule ended that, and some financial commentators now argue the traditional IRA has lost much of its value as an inheritance, because the tax can no longer be stretched thin across a lifetime. For the owner, it sharpens a useful question. If the tax on this money will be paid either way, who should pay it: the owner, during retirement, possibly in a low-income year and at a low rate, or the heirs, who may inherit during their own working years and pay at a higher

one? That question is the heart of the Roth conversion decision.

A second part of the rule decides whether the heir can wait the full ten years or has to take something every year along the way, and it caught a lot of people by surprise. It turns on the age of the person who died. If the original owner was 73 or older and had already begun taking their own required distributions, the heir must keep taking a distribution every year, in years one through nine, and still empty the account by year ten. If the owner died younger, before required distributions began, only the ten-year deadline applies, with nothing required in between. The rule was unsettled long enough that the IRS waived the penalty for missed annual distributions from 2021 through 2024. Full enforcement started in 2025.

Dave Thompson’s inherited account shows the harder version of this. He inherited a traditional IRA from his father, who died at 79, several years into his own required distributions. Because his father was past 73 and already taking them, Dave’s account carries both obligations: a distribution every year and a hard deadline to empty it. He is in year five of ten. Each withdrawal stacks on top of the income he already has, and the closer year ten gets, the larger the remaining withdrawals have to be to clear the account in time. Heirs who treat the early years as optional often pay for it with a punishing final stretch.

Linda Castellano's account adds another layer, and a useful term. A successor beneficiary is the person who inherits an IRA that was already inherited once before: the beneficiary of a beneficiary. Linda's mother had inherited a traditional IRA from her own mother, Linda's grandmother, who had died some years earlier, already taking her required distributions. As a non-spouse heir, Linda's mother was bound by the ten-year rule and was taking annual distributions toward that deadline. When Linda's mother died two years ago, Linda became the successor beneficiary. A successor does not receive a fresh ten years; she inherits only the time left on the original clock. Linda has two years to empty the account, which compresses years of deferred tax into a short, high-income window, all of it taxed at single-filer rates.

A surviving spouse stands apart from every other heir. The ten-year rule does not apply to a spouse. A surviving spouse can roll an inherited account into their own IRA and treat it as if it had always been theirs, taking distributions on the normal schedule rather than any ten-year countdown. That difference is why beneficiary designations, and the order in which accounts pass from person to person, deserve as much attention as the investments inside them.

Ways to Lower the Tax

A required distribution cannot be avoided, but the tax that comes with it can be managed, and the most useful work happens before age 73 ever arrives. The

earlier it starts, the better. Choices made during the working years, such as funding a Roth alongside a traditional account, shape how large the required distributions will eventually be. So do the years between leaving work and the first required distribution, which are often a person's lowest-income years: the paycheck has stopped, Social Security may not have started, and the required distributions have not begun. Every dollar pulled from a traditional account in that low-income window, or converted to a Roth, is a dollar that will not be forced out later at a higher rate. Acting well before the deadline, rather than scrambling at 73, is what makes the difference.

For people who give to charity, there is a way to satisfy a required distribution without counting it as income at all. It is called a qualified charitable distribution, or QCD. Someone 70½ or older can send money straight from an IRA to a qualifying charity, and that amount counts toward the year's required distribution while staying entirely out of taxable income. The limit in 2026 is $111,000 per person. The eligibility age of 70½ actually comes before the required-distribution age of 73, so this option opens up before the requirement does. The rules are specific. The money has to go directly from the IRA custodian to the charity; a check made out to the account holder first does not count. It has to go to a 501(c)(3) public charity, not a donor-advised fund or a private foundation. And it only works from an IRA, not a 401(k). For someone who already donates, it is the most tax-efficient way to do it, because the gift never enters adjusted gross income, which also helps hold

down both Social Security taxation and IRMAA two years down the road.

A few smaller moves help around the edges. Taking the first distribution on time, in the year of turning 73, avoids the two-in-one-year pileup. Someone still working can use the current-employer exception. Spending from taxable accounts and traditional retirement accounts first, while leaving Roth and HSA balances alone to grow, keeps the base that future distributions are calculated on smaller. A qualified longevity annuity contract can move a limited amount out of that base entirely, pushing those dollars and the tax on them out to the mid-eighties, though the payments are taxable once they begin. And because tax withheld from a distribution counts as if it were paid evenly all year, a retiree can have enough withheld from a December distribution to cover the whole year’s tax bill, which avoids quarterly estimated payments and the penalties for underpaying them. None of these makes the distribution go away. Each one changes what it costs.

The Costs a Distribution Triggers

A required distribution is never just a withdrawal. It is income, and it sets off other charges. It raises adjusted gross income, which can make more of a Social Security benefit taxable, push up the tax bracket, and, two years later, trigger an IRMAA surcharge on Medicare.

The Whitfields hold about $1.2 million in a traditional IRA. Their first required distribution

> will be more than $45,000 even at today's balance, and it will climb every year after, on top of everything else they draw. For them, the IRMAA surcharge is not a risk to avoid; it is a fixed part of their Medicare cost.

Every one of those downstream charges traces back to a single set of numbers: the December 31 balances the distributions are calculated from. Shrink the traditional balances before 73, and every distribution after that shrinks too, along with the higher brackets, the Social Security taxation, and the IRMAA surcharges that follow from them. That is why the few years before distributions begin matter so much, and why the main tool for using them is the Roth conversion.

The required distribution rewards one kind of person above all: the one who treats it as a scheduled event instead of a yearly surprise. The age is known. The table is published. The math can be run before the first dollar is ever due. The government will collect the tax it deferred decades ago, and that part is not up for debate. What is still open, and what the next few chapters are about, is the timing and the rate, both of them still in the retiree's hands in the years before the clock starts.

Chapter 10: Roth Conversions — Paying Taxes on Your Terms

There is a lever hidden in the required distribution rules. A required distribution cannot be avoided, but its size is set by one number, the balance in the traditional account on December 31, and that number can be made smaller before age 73 ever arrives. The main tool for making it smaller is the Roth conversion. It is the closest thing retirement planning offers to choosing one's own tax rate.

A conversion does not make the tax on a traditional account disappear. Nothing does. What it does is move the tax forward in time, out of a year the government picks and into a year the retiree picks. For someone with a low-income stretch between leaving work and the start of required distributions, that single choice can be worth tens of thousands of dollars over a retirement. It can also be done badly, in ways that raise a Medicare premium or erase a health-insurance subsidy. Three things matter: the case for converting, the methods, and the traps.

What a Conversion Actually Is

A Roth conversion moves money out of a traditional, pre-tax retirement account and into a Roth account, and the amount moved is reported as ordinary income in the year of the move. The money was never taxed going in, so it is taxed coming across. A $40,000 conversion adds $40,000 to taxable income for the year, stacked on top of whatever else the household earns, and taxed at the rates that income reaches. There is no special conversion rate; the converted

dollars pile onto the year's ordinary income and are taxed like any other.

What changes is everything that comes after. Once the money is in a Roth, it is finished being taxed. It grows tax-free, and qualified withdrawals come out tax-free. It carries no required distribution during the owner's lifetime. And when it passes to an heir, the heir still has to empty it within ten years, but takes those withdrawals tax-free. The whole transaction is one trade: pay a known tax rate now to escape an unknown one later.

An analogy makes the trade concrete. Imagine a farmer who owes tax on his crop, and a county tax collector who, being neighborly about it, lets him choose when to pay. The rate is the same either way. The farmer can pay now, on the value of his seed, or wait and pay at harvest, on the value of the grown corn. The seed is cheap and the corn is dear, so paying on the seed looks like the obvious bargain. A traditional IRA poses the same choice. The tax was deferred, not forgiven, and it will be owed someday. Converting to a Roth pays the tax now, on the seed. Leaving the money to grow and drawing it out later pays the tax on the harvest.

The analogy hides one catch, worth stating because the seed is not always the bargain it looks like. If the rate is truly the same at planting and at harvest, and the tax is paid out of the crop itself, the two choices come out even: a smaller amount taxed and then grown equals a larger amount grown and then taxed. The seed becomes the real bargain under two conditions. The first is that the rate is higher at

harvest than at planting, which is usually the way the bet leans. The second is that the farmer pays the seed tax from his own pocket rather than out of the crop, so the whole crop grows untaxed. That second condition translates directly: paying the conversion tax from a taxable account rather than from the IRA lets the entire converted balance compound tax-free, and it is what lets a conversion pay off even when the future rate is no higher. Both conditions matter throughout.

Why Pay the Tax Early

Paying tax sooner than the law requires sounds like a mistake. Several things make it the opposite. The first concerns required distributions. Every dollar moved out of a traditional account before 73 is a dollar no longer in the December 31 balance, so it no longer drives a required distribution, or the higher bracket, the Social Security taxation, and the IRMAA surcharge a distribution can set off. Shrinking the traditional balance now shrinks all of that later.

The rest of the case follows from the same idea. Money in a Roth grows and comes out tax-free, so years of future growth escape tax entirely. A conversion also locks in today's tax rate against the chance of a higher one later, as the retiree's own income climbs once Social Security and required distributions are both running. And it improves what the owner leaves to heirs. The ten-year rule stripped the traditional IRA of much of its value as an inheritance, because the tax can no longer be spread across an heir's lifetime. A Roth gets around that: the heir still empties it within ten years, but owes no tax doing so. The question of who pays the tax on this money, the owner or the

heirs, is the one a conversion settles, in the owner's favor.

The Bet on Future Tax Rates

There is one more reason, and it sits underneath the others: the future rate on a traditional account is not known today. Deferring tax is a bet that the rate skipped now will be no higher than the rate paid in retirement, and by the standard of the last century that is not a safe bet. Today's top rates are low by historical standards, and the country's finances point one way. The federal debt has roughly quadrupled since 2008 and now runs well above the size of the whole economy; the interest on it has grown into one of the largest lines in the budget; and spending has not fallen. Of the levers left to close that gap, higher tax revenue is among the most likely. None of this is certain, and the 2025 law made the current brackets permanent, but permanent in tax law lasts only until the next Congress decides otherwise. A conversion trades that uncertainty for a rate the retiree can see.

This is the thinking behind what advisors call tax diversification: holding retirement money in more than one tax treatment rather than all in one. Savings tend to fall into three buckets. There is a taxable bucket, such as a brokerage account, taxed along the way. There is a tax-deferred bucket, the traditional 401(k) and IRA, taxed on the way out. And there is a tax-free bucket, taxed on the way in and never again. Most people reach retirement with nearly everything in the tax-deferred bucket and little or nothing tax-free, which leaves a single future rate to govern almost every withdrawal. Spreading across all three means no

one rate can reach the whole pile. The Roth, built through conversions, is the main way to fill the tax-free bucket, though not the only one. The cash value of certain permanent life insurance, such as indexed universal life, can be reached tax-free as well, within limits, and fixed indexed annuities offer tax-deferred growth with a measure of protection against market losses. Each carries its own costs and tradeoffs. The broader point holds regardless: the more retirement income that can be drawn from outside the tax-deferred bucket, the less any future rate increase can reach it.

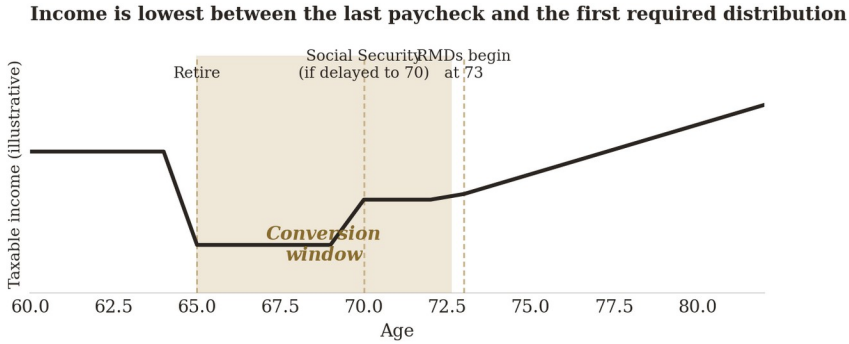
Why the Early Years Are the Cheap Ones

Conversions have a season because not every year of retirement is taxed the same. The cheapest years are usually the ones right after the last paycheck and before required distributions begin. Work has stopped, Social Security may be delayed, and the required distributions have not started, so the lowest tax brackets sit open and unused. Three forces then make the later years more expensive, which is exactly why moving income into the early ones pays off.

The first force is the required distribution itself, which stacks new income on top of everything else once it begins at 73. The second is the way Social Security is taxed. More of a Social Security check becomes taxable as other income rises, so across a wide band of income each extra dollar pulled from a traditional account also drags part of a Social Security benefit into taxable income. The effect, sometimes called the tax torpedo, can push the real tax on that withdrawal far above the bracket it appears to sit in, turning what

looks like a 22 percent rate into something closer to 41 percent. Converting before Social Security begins steps around it. The third force arrives with the death of a spouse. The survivor files as a single taxpayer the very next year, with narrower brackets and a smaller standard deduction, often with close to the same income. A couple who converts while both are alive pays at the wider joint brackets instead of leaving the survivor to draw the same accounts down alone at single rates.

That last point lands directly on these families. Mary Thompson is younger than Dave and, as the statistics go, likely to outlive him; the years they convert together at joint rates are years Mary will not spend drawing a large traditional balance down by herself as a single filer. Robert and Susan Whitfield, a fictional couple, face the same arithmetic at a larger scale. The window itself is easiest to see drawn out.



How taxable income tends to move through retirement. The shaded stretch, after the last paycheck and before required distributions, is usually the lowest-taxed of a retiree's life, and the natural time to convert. Illustrative.

How Much to Convert

The strategy that uses the gap years, the low-income stretch after the last paycheck and before required distributions begin, is called filling a bracket, and it works because of how the tax actually applies. Federal income tax is marginal: each bracket's rate falls only on the dollars inside that bracket, not on the whole income. A married couple in 2026 pays 10 percent on the first $24,800 of taxable income, 12 percent on the dollars up to $100,800, 22 percent above that, and so on, to 37 percent at the top. Filling a bracket means converting just enough to reach the top of the bracket the household already sits in, without pushing the next dollar into a higher one.

The table below shows where the 2026 brackets end. The figures are taxable income, meaning income after the standard deduction, which in 2026 is $32,200 for a married couple and $16,100 for a single filer, has already been subtracted.

Top rate on income up to	Married, filing jointly	Single filer
10%	$24,800	$12,400
12%	$100,800	$50,400
22%	$211,400	$105,700
24%	$403,550	$201,775

Top of each lower 2026 tax bracket, in taxable income after the standard deduction. Above the 24 percent bracket lie the 32, 35, and 37 percent brackets; the 37 percent rate begins above $768,700 for a married couple and $640,600 for a single filer. These thresholds are inflation-adjusted yearly.

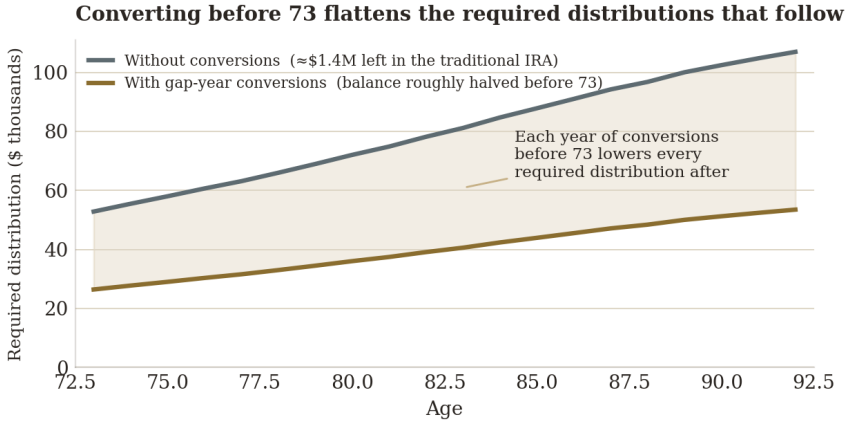
More is not automatically better. The jump from the 12 percent bracket to the 22 percent bracket is steep, so converting one dollar past the top of the 12 percent band nearly doubles the tax on that dollar. For most people this means the whole traditional account should not be converted at once. The useful question is not how much can be converted but how much can be converted at a rate worth paying. A common discipline is to pick a ceiling and stop there: the top of a bracket, the dollar just below an IRMAA threshold, or the rate the household expects to pay later in retirement anyway. Converting past that point trades a low rate now for an even lower one that may never arrive.

Filling the Bracket in Practice

Dave and Mary Thompson, a fictional couple, show the everyday version. Dave is 64 and winding down his consulting work; Mary is 61. Their pre-tax balance is modest: Dave’s traditional IRA holds about $100,000, and Mary’s $150,000 is already in a Roth, so there is nothing to convert there. In a year after Dave’s consulting income has stopped, with both of them delaying Social Security, their taxable income might sit low in the 12 percent bracket. Converting enough to bring that income up toward $100,800, the top of the bracket, would move a large piece of the $100,000 across at 12 percent. Spread over two or three such years, the whole traditional IRA could become a Roth, taxed at 12 percent the entire way, and never appear in a required distribution again. The one piece Dave cannot

touch is the $60,000 IRA he inherited from his father; inherited accounts cannot be converted.

The Whitfields work at a different scale. Between Robert’s $1.2 million traditional IRA and the $200,000 in Susan’s old 401(k), they hold roughly $1.4 million that has never been taxed. Left alone, it produces required distributions in the tens of thousands of dollars every year starting at 73, piled on top of Social Security and everything else. Because Robert plans to delay Social Security to 70, their income stays low from his retirement at 67 until 73, giving six clear years. In those years they can convert up to the top of the 24 percent bracket, $403,550 of taxable income for a couple, and pay the tax from their $400,000 brokerage account rather than from the IRA. Over six years that moves a large share of the $1.4 million into Roth accounts at rates that stop at 24 percent, below the rates their stacked retirement income would otherwise reach. What does not get converted still faces required distributions, but those distributions, and the tax their three children would owe on whatever they inherit, are smaller for every dollar moved.



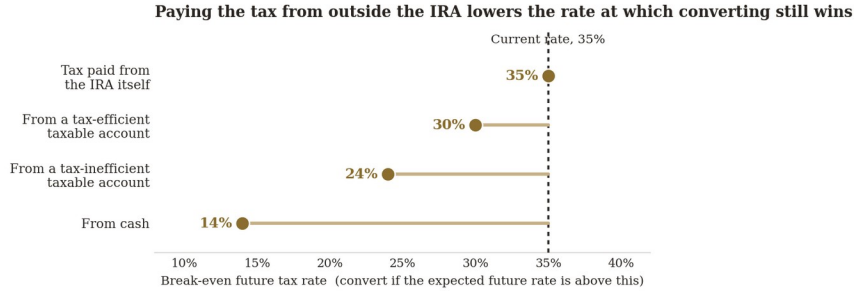
Illustrative required distributions on a $1.4 million traditional balance, with and without conversions during the gap years. Converting before 73 lowers every distribution that follows. Figures are illustrative and assume steady growth.

A Conversion Can Pay Off Even at a Lower Future Rate

The usual rule of thumb is to convert when the future tax rate is expected to be higher than today's, and to hold off when it is expected to be lower. That rule is a good start, but it misses something. When the tax on a conversion is paid from a taxable account rather than from the IRA, the conversion effectively moves money into the tax shelter. The dollars used to pay the tax would otherwise have stayed in the taxable account, generating taxable income year after year; spent on the conversion instead, they let the full converted balance grow tax-free inside the Roth. That hidden shift means a conversion can come out ahead even when the future rate turns out somewhat lower than today's.

Analysts at Vanguard describe this with a single figure they call the break-even tax rate: the future rate at which a person would be indifferent between

converting and not. If the expected future rate is above it, converting wins; below it, keeping the traditional account wins. Paying the conversion tax from outside the IRA pulls that break-even rate below the current rate, and the longer the money will grow and the more heavily the taxable account was being taxed, the lower it goes. Having after-tax money already in the IRA lowers it further. The practical lesson is this: where the tax is paid changes the whole calculation.



The break-even future tax rate falls when the conversion tax is paid from outside the IRA. Above the marked rate, converting wins; below it, keeping the traditional account wins. Illustrative, assuming a 35 percent current rate and a 20-year horizon. Adapted from Vanguard research on break-even tax rates.

Paying the Tax

The benefit of a conversion is largest when the tax on it is paid from money outside the retirement accounts. Paying it out of the converted amount itself shrinks the sum that actually reaches the Roth, and for anyone under 59½, the portion held back to cover the tax counts as a distribution and draws the 10 percent early-withdrawal penalty.

The Whitfields’ brokerage account is what makes their large conversions worth doing. Linda Castellano, a fictional 58-year-old, sits at the other

end of this. With only a small taxable cushion, she can cover the tax on just a modest conversion each year, which caps how much she can sensibly move.

Capital Gains, the Other Rate

Paying from a brokerage account is usually the right move, but it is not free, because not all income is taxed on the same schedule. A Roth conversion is ordinary income, taxed at the marginal rates in the bracket table, and so are short-term capital gains, the profit on anything held a year or less. Long-term gains, on anything held more than a year, run on a separate and gentler track: zero, 15, or 20 percent, with an added 3.8 percent surtax once income climbs high enough. Most retirees who sell a long-held investment pay the 15 percent rate.

This matters because raising cash in a brokerage account to pay the conversion tax usually means selling something that has gained value, and that sale carries a tax of its own. For a long-held position that tax is typically 15 percent, often below the marginal rate the conversion is paying, which is why a taxable account is still usually the best place to draw from. But the two stack. In a low-income year a married couple can hold roughly $98,900 of taxable income, or a single filer about $49,450, before a long-term gain is taxed at all; a large conversion, by filling that room with ordinary income, can push the same gains up out of the zero bracket and into the 15 percent one. The conversion and the gains sold to fund it are best sized together.

The Rules That Trip People Up

Two features of conversions surprise people, and the first is a welcome one: there is no income limit. The income ceilings that block high earners from contributing to a Roth directly do not apply to converting into one, so any amount can be converted in any year, at any income. For anyone shut out of direct Roth contributions, that is a powerful advantage, the open door that makes the backdoor strategies possible. The second surprise is less welcome. A conversion cannot be undone. Until 2018 a person could reverse one, through a move called recharacterization, but the 2017 tax law ended that for conversions. Once the money is converted and the year closes, the tax is owed, no matter what the account does afterward. That makes the size of a conversion something to get right before acting.

After required distributions begin, the order of operations matters. A required distribution itself cannot be converted. Someone 73 or older has to take the full required amount first, as taxable income, and can convert only what is above it. This is part of why conversions belong to the years before 73.

> Walter Higgins, a fictional 84-year-old, shows the limit. He is years into his required distributions, so any conversion would stack on top of a distribution he already has to take and pay tax on, and at his age few years remain for tax-free growth to outrun the upfront cost. For Walter, the window has mostly closed.

A provision called the pro-rata rule decides how much of a conversion is taxable when a traditional IRA holds both pre-tax and after-tax money. A person cannot single out the after-tax dollars to convert. The IRS treats all of someone's traditional, SEP, and SIMPLE IRAs as one combined pool and taxes each conversion in proportion to how much of that pool was never taxed. Workplace 401(k)s and inherited IRAs stay outside the calculation, and spouses are measured separately. For the many people whose traditional balance is entirely pre-tax, the rule changes nothing. For anyone who made non-deductible contributions over the years, it has to be worked out before converting.

Roth accounts run two separate five-year clocks, and they answer different questions. The first applies to each conversion on its own: converted money taken back out within five years, by someone under 59½, is hit with the same 10 percent penalty the conversion was meant to sidestep. The second governs tax-free growth: earnings come out tax-free only once the owner is 59½ and has held a Roth for at least five years. For a retiree already past 59½ who opened a Roth years ago, both clocks are satisfied and neither matters. For someone retiring early, both are worth tracking, and together they are a good reason to open a Roth sooner rather than later, even with a small amount, just to start the clock.

Reaching the Money Early: The Conversion Ladder

That first five-year clock is also the engine of a strategy for people who retire well before 59½, called

a Roth conversion ladder. The idea is to convert a set amount every year, and because each conversion becomes reachable without penalty once its own five years have passed, the conversions begin maturing in sequence. Someone who starts converting at 50 can begin withdrawing the earliest conversions, penalty-free, at 55, then the next at 56, and so on. Each year's conversion forms another rung, and the ladder turns a locked retirement account into a steady stream of spendable money years before the normal age. The converted principal comes out tax-free, because the tax was already paid at conversion.

The catch sits at the front. Nothing is reachable until the first conversion has aged five years, so an early retiree needs five years of spending already saved somewhere accessible, usually a taxable account, to live on while the first rungs mature. For most of these families the ladder is not the main event, since they are at or past 59½. But it matters for anyone leaving work early, and it is one more reason to start the five-year clock as soon as possible.

Two Side Doors for High Earners

Conversions also solve a problem at the other end of a career, for people still working and earning too much to fund a Roth directly. The income limits that block a Roth contribution do not block a conversion, which opens a side door. The first version is the backdoor Roth: a person contributes to a traditional IRA without taking the deduction, then converts that contribution to a Roth. The pro-rata rule is the thing to watch, because if the person holds other pre-tax IRA money,

the conversion is taxed proportionally rather than coming across free.

The second version is larger and runs through a workplace plan. Some 401(k) plans allow after-tax contributions on top of the regular limit, filling the space up to the total plan ceiling, which is $72,000 in 2026 across all sources. Those after-tax dollars can then be converted to Roth, either inside the plan or by rolling them to a Roth IRA, a move often called the mega-backdoor Roth. It only works if the plan specifically allows both the after-tax contributions and the conversion, so the plan document is the place to check first.

	Backdoor Roth	Mega-backdoor Roth
Who it fits	Earners above the Roth income limit	Employees whose 401(k) allows after-tax contributions
How it works	A non-deductible traditional IRA contribution, then a conversion	After-tax 401(k) contributions, then a conversion to Roth
2026 ceiling	The $7,500 IRA contribution limit	Up to the $72,000 total plan limit, minus deferrals and match
Main catch	The pro-rata rule, if other pre-tax IRA money exists	The plan must allow both the contributions and the conversion

Two ways high earners locked out of direct Roth contributions can still reach a Roth. Both rely on the fact that conversions have no income limit. 2026 figures.

The Costs a Conversion Can Trigger

A conversion is income, and like a required distribution, it can trigger charges beyond the tax on it. The main one is IRMAA, the Medicare surcharge. Because IRMAA runs on a two-year delay, a conversion done at 63 or later raises Medicare premiums two years on, once the person is enrolled. The IRMAA

brackets are cliffs: a single dollar over a threshold raises the surcharge for the whole year, sometimes by more than a thousand dollars per person, so conversions near those edges should stop just under them.

The Whitfields are an unusual case here. Their required distributions will put them above the IRMAA thresholds every year regardless, so the surcharge is already a fixed cost and does not argue against converting.

For anyone using a marketplace health plan before Medicare, a conversion runs into a different threshold. If Linda retires before 65, she will buy coverage through the Affordable Care Act marketplace until Medicare starts, and the premium subsidies there shrink as income rises. In 2026 they cut off entirely above 400 percent of the federal poverty level, roughly $62,600 for a single person. A conversion that lifts her income over that line could cost her the subsidy on a full year of premiums, often several thousand dollars, on top of the tax. These subsidy rules have shifted from year to year and should be checked against current law before acting. For Linda, the gap years are still the time to convert, but every conversion has to be weighed against what it does to her health-insurance cost until she turns 65.

Who Should Think Twice

A conversion is not the right move for everyone. Someone who will clearly be in a lower bracket later should not prepay a higher rate now. Someone who can pay the conversion tax only from the IRA itself gives back much of the benefit. Someone who will need the converted money within five years, while still under 59½, runs into the penalty the conversion was meant to avoid. And someone already well into required distributions, like Walter, has too little time left for the math to work. There is no universal right amount. It depends on the length of the gap years, the income already filling the brackets, the filing status, and the IRMAA and subsidy thresholds nearby. The figure that fits the Whitfields would be far too much for Dave and Mary, and either could be wrong for Linda in a year she is protecting a subsidy.

A final point applies to almost anyone weighing a conversion. The decision pulls on more threads than it first appears: the tax brackets, the IRMAA thresholds, the subsidy cliffs, the two five-year clocks, the capital-gains cost of paying the tax, and a forecast of income and tax law years into the future. Each one moves the others, and the answer is usually a multi-year schedule rather than a single number. Tools built for retirement income planning can model that whole sequence and show the lifetime effect of converting a given amount in a given year, which is hard to work out by hand. For many people a conversion is a sensible reason to bring in a professional who specializes in retirement income,

not to hand off the decision, but to see the moving parts modeled together before committing to something that cannot be reversed.

The required distribution is set by a balance and a table. The conversion is the main tool for making that balance smaller before the table ever applies. Neither one makes the tax disappear. What the conversion adds is control over the timing and the rate. The government will collect on every pre-tax dollar in the end, from the owner or from an heir. A conversion lets the owner choose to pay now, at a rate they can see, in a year they pick, instead of later, at whatever rate a future of stacked income and changed law produces. Used in the low-income years before 73, it turns a tax bill that would otherwise arrive on the government's schedule into one settled on the retiree's own.

Chapter 11: The Case for a Bond Alternative

By the time most people reach retirement, they have heard the same advice for years: as the finish line nears, move money out of stocks and into bonds. Stocks are for growing wealth; bonds are for protecting it. The standard glide path shifts a larger and larger share into bonds as a person ages, on the theory that bonds are the safe, conservative anchor of a portfolio. That theory deserves a serious test, because one recent year put a real crack in it.

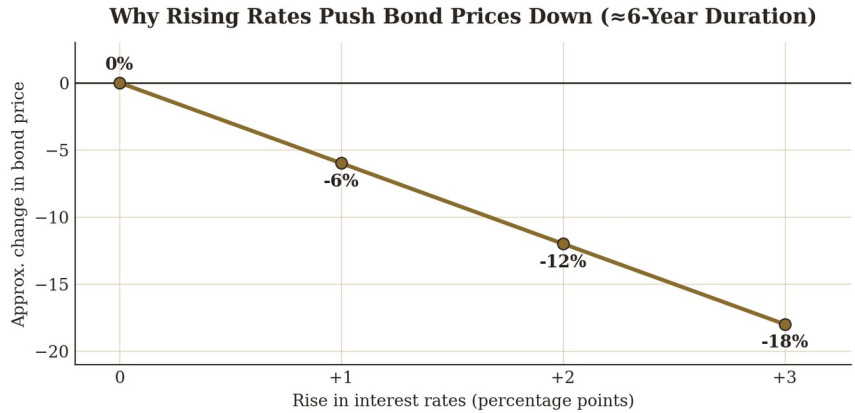
Start with what a bond actually is. A bond is a loan. The buyer hands money to a government or a company, which agrees to pay interest along the way, called the coupon, and to return the original amount, the principal, on a set date called maturity. That structure earned bonds their reputation. The income is steadier than a stock’s, the issuer promises the principal back at the end, and high-quality bonds rarely default. For decades the reputation held up well enough that few people questioned it. Then came 2022.

When the Safe Money Lost

In 2022 the broad U.S. bond market had its worst year on record. The Bloomberg U.S. Aggregate Bond Index, the standard measure of the investment-grade bond market, fell about 13 percent. To see how far outside the normal that was, the index had never lost more than 2.9 percent in a calendar year since it began in 1976. Longer-dated bonds fell much further: intermediate Treasuries lost more than 10 percent,

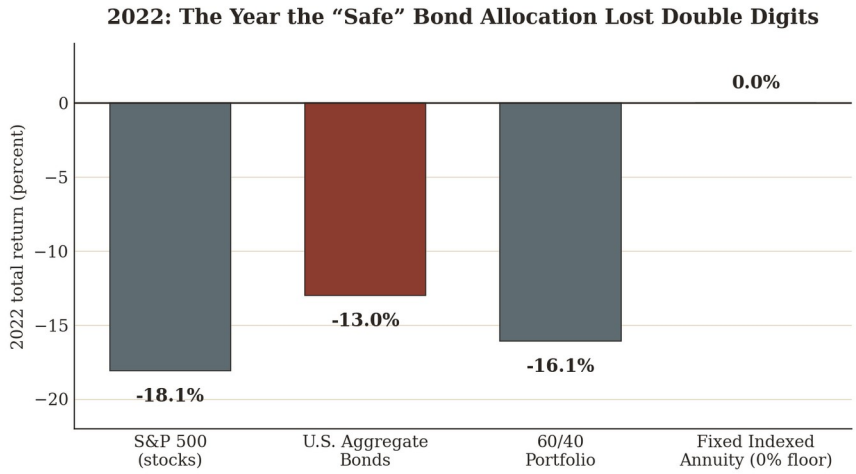
their worst showing in roughly a century, and the longest Treasury bonds lost close to 40 percent. The cause was simple. To fight inflation, the Federal Reserve raised interest rates at the fastest pace in 40 years, lifting its benchmark from near zero to about 4.5 percent in a single year.

Rising rates are the one force that reliably drives bond prices down, and the reason is mechanical. When new bonds come to market paying 5 percent, no one will pay full price for an older bond paying 2 percent. The price of the older bond has to fall until its effective yield matches what a buyer could get elsewhere. The longer a bond has left until maturity, the more its price has to move, a sensitivity measured by a number called duration. A bond or bond fund with a duration of about 6 loses roughly 6 percent of its value for every 1 percentage point that rates rise. In 2022 rates rose far more than a single point, and the price losses followed.



A bond fund's price moves opposite to interest rates, and the longer its duration, the larger the move. At a duration of about six years, each one-point rise in rates costs roughly six percent of value. In 2022 rates rose far more than one point. Illustrative.

The deeper problem in 2022 was not that bonds fell. It was that they fell at the same moment stocks did. The classic retirement portfolio, 60 percent stocks and 40 percent bonds, rests on the idea that the two move differently: when stocks drop, bonds usually hold steady or rise, softening the blow. That cushion is the whole reason the safe sleeve exists. In 2022 it was not there. The S&P 500 lost about 18 percent, bonds lost about 13 percent, and the 60/40 portfolio fell more than 16 percent, its third-worst year ever. A retiree drawing income that year was selling into losses on both sides of the portfolio at once, the exact sequence-of-returns danger that haunts a portfolio in early retirement.



Total returns for 2022. Stocks and bonds fell together, so the standard 60/40 portfolio lost more than 16 percent, its third-worst year on record. A fixed indexed annuity’s 0 percent floor would have credited zero rather than a loss. Index returns are historical; the annuity figure illustrates the floor.

One fair objection deserves an answer. Hold a single bond to maturity, the argument goes, and the issuer returns the full principal, so the dip in price along the

way never becomes a real loss. That is true for one bond held to the end, default aside. But most retirees do not own individual bonds; they own bond funds, and a fund never matures. It holds a rolling pool of bonds, so its share price moves with the market every day, and a sale during a downturn locks in the loss. Even holding a single bond to maturity carries a quieter cost: the money stays parked at the old, lower rate while better rates pass it by. Holding to maturity softens the problem. It does not erase it.

Where Bonds Still Earn Their Place

None of this makes bonds a bad investment, and a fair assessment has to say so. Two points sit firmly in their favor. First, bonds pay far more now than they did before the 2022 selloff. By the middle of 2026 the ten-year Treasury yielded about 4.5 percent, against roughly 1.5 percent a few years earlier, so the income a buyer locks in today is real and meaningful. Second, bonds are liquid and predictable in a way little else is: a Treasury can be sold on any business day, and a bond held to maturity returns a known amount on a known date, which is exactly what a retiree timing a future expense wants. The lesson of 2022 is narrower than the claim that bonds are bad. It is that the word bond and the phrase safe from loss are not the same thing, especially for the bond-fund version most people hold. Higher rates today have not retired that risk either; in 2026 the Federal Reserve was signaling that further increases were still possible.

A Different Place for the Safe Money

This is where the fixed indexed annuity returns to the story, in a more specific role than before. In its first role, it served as a volatility buffer, a way to avoid selling investments at a loss in a down market. Here the question is narrower: could it hold part of the money that would otherwise sit in bonds? A quick reminder of how it works. A fixed indexed annuity credits interest tied to a market index, with two guardrails written into the contract: a cap that limits the gain in a strong year, and a floor of 0 percent that prevents any loss in a weak one. The account value does not fall because the index dropped or because rates rose. Its growth is tax-deferred until the money comes out, and the contract is backed by an insurance company rather than a bank.

That single trait is what speaks to the bond problem. The damage in 2022 came from prices being marked down, bond values rewritten lower as rates climbed. A fixed indexed annuity's account value is not marked to a market. It does not drop when rates rise, and it does not drop when the index falls; in a year like 2022 it would have credited zero rather than losing 13 percent. For the slice of a portfolio whose entire job is to hold its value through a bad year, refusing to lose is the point, and that is the one thing a bond fund cannot promise.

Comparing the Two

A cleaner way to compare the two is to judge each on the five things that matter for any place a person puts money: safety, liquidity, rate of return, tax treatment, and cost. No single vehicle wins on all five, so the real

exercise is to see which trade-offs fit which job. On safety, a bond fund can lose value when rates rise, as 2022 proved, while the annuity's floor prevents a market loss, though its protection rests on the issuing insurer's ability to pay claims, not on federal deposit insurance. On liquidity, bonds win, though by less than it first looks: a bond can be sold any day, while an annuity is less liquid during its surrender period, commonly seven to ten years. Even then, most contracts let the owner withdraw a set share each year, often around 10 percent, with no surrender charge, so only larger withdrawals are penalized. On return, a bond pays a stated yield a buyer can count on, lately around 3 to 5 percent, while an indexed annuity's credit is variable and capped, zero in a poor year and up to its cap in a strong one; industry figures commonly put its long-run average somewhere around 5 to 7 percent, higher than a bond's yield, though that average is earned across good and flat years and is never promised in advance. On taxes, bond interest in a taxable account is taxed every year, while the annuity grows tax-deferred. On cost, many bond funds carry an annual expense ratio, while many plain annuities carry no ongoing fee, with optional riders the main exception.

	Bonds and bond funds	Fixed indexed annuity
Safety	Can lose value when rates rise, as 2022 showed	A 0% floor blocks market loss; backed by the insurer, not the FDIC
Liquidity	Sellable on any business day	About 10% a year withdrawable free; larger withdrawals charged during a 7-to-10-year surrender period
Return	A stated yield, lately about 3% to 5%	Index-linked and capped; 0% in a down year, often near 5% to 7% on average
Taxes	Interest taxed yearly in a taxable account	Grows tax-deferred until withdrawal
Cost	Many funds charge an annual expense ratio	Many carry no ongoing fee; riders cost extra

Bonds and a fixed indexed annuity, judged on the five traits that matter for any holding. Neither wins on all five. Specific terms vary by contract and carrier, and the return figures are typical ranges, not guarantees.

Read straight across, the table says something other than that one is better. The annuity wins on protection from market loss, on tax deferral, and often on cost. Bonds win decisively on liquidity and on paying a known, spendable stream of income. That is why a fixed indexed annuity is not a wholesale replacement for bonds. It is a candidate for the part of the safe

money that does not need to be liquid and whose main job is simply to avoid loss, which frees bonds and cash to do the jobs they do best.

Some historical evidence sits behind the comparison, though it has to be read with care. The economist Roger Ibbotson, known for the long-run studies of stock and bond returns, published a 2018 paper simulating how a fixed indexed annuity would have fared against bonds across the 90 years from 1927 to 2016. Net of fees, his simulated annuity returned about 5.8 percent a year, against about 5.3 percent for long-term government bonds, and it never posted a losing year. Two cautions belong with that figure. The study modeled an idealized, uncapped annuity that did not exist for most of the period, so it is a back-test rather than a real track record, and the research was funded by the annuity industry. Real contracts cap the upside, which pulls the number down, and in low-rate stretches such as the years after 2008, when caps were thin, actual indexed-annuity returns often ran in the low-to-mid single digits. The fair reading is not that an annuity reliably beats bonds, but that over a long horizon it has held up well against them while sidestepping their losing years, and that its edge tends to appear exactly when bonds do worst.

The penalty-free withdrawal allowance does more than soften the lockup. The roughly 10 percent a contract typically lets out without charge is close to what a holder might draw if the annuity were later turned into lifetime income, and it is enough for the same money to serve as the volatility buffer. In a down year, the owner can pull income from the annuity without selling

anything at a loss, because its value never fell. The protection from loss and the spendable cash come from the same slice of the portfolio.

Putting It in the Plan

Robert and Susan Whitfield, a fictional couple, show how this fits a real portfolio. At 67 and 65, they hold roughly $1.8 million across an IRA, an old 401(k), and a brokerage account. Conventional advice would have them move a growing share, perhaps 40 percent, into bonds as they age. The year 2022 is a preview of what that supposedly safe 40 percent can do in the wrong conditions. An alternative is to hold part of that conservative sleeve in a fixed indexed annuity, where a repeat of 2022 credits zero instead of losing double digits, while keeping enough in bonds and cash for liquidity and predictable income. The protected slice has a second effect worth naming: because part of the portfolio can no longer lose value to a market decline, the Whitfields can comfortably keep a bit more of the rest in stocks than they otherwise might, which is where their long-term growth comes from. Either way, shifting part of the safe money into the annuity lowers the risk of the entire portfolio, not only of the slice that was moved.

How much belongs in such a slice is not a fixed number, and Dave and Mary Thompson, a fictional couple, show why. Mary’s pension and the couple’s Social Security cover most of their basic expenses, so a smaller share of their savings is

never exposed to a forced sale at the wrong time. That reliable income is already doing part of the safe sleeve's job. Someone living almost entirely on portfolio withdrawals, with little dependable income behind them, has far more reason to protect the conservative slice. The right size depends on how much steady income is already in place, the time horizon, how much cash the household may need on short notice, and plain temperament. It is a question better worked out with someone who can see the whole picture than copied from a book.

The argument here is not that bonds are bad or that an annuity is a cure. It is that treating the word bond as a synonym for safe deserves a harder look, especially after 2022 and especially for the fund version most retirees own. For the portion of a portfolio whose only assignment is to hold its value through a bad year, a 0 percent floor does that job in a way a bond fund cannot match. Bonds keep their place for liquidity and income; a fixed indexed annuity can take over the narrower task of not losing money. How that annuity actually works, what its caps and riders and surrender terms really mean, and how to weigh one contract against another, is the next question.

Chapter 12: How a Fixed Indexed Annuity Actually Works

A fixed indexed annuity has appeared twice already, and both times in a single role. In its first role, it was a volatility buffer, a place to hold money that would not fall in a bad market. In a second role, it became a candidate for the part of a portfolio that would otherwise sit in bonds. Both descriptions rested on the same short summary: the account earns interest tied to a market index, a cap limits the gain in a strong year, and a floor of zero prevents a loss in a weak one. That summary is true, and for deciding whether the product belongs in a plan at all, it is enough. For deciding which contract to buy, it is not. Now the contract itself comes open.

The reason the summary is not enough is that a fixed indexed annuity is not one product. It is a frame around a set of choices, and the choices determine almost everything about how the contract behaves. Two annuities sold under the same name, by the same company, in the same week, can credit very different amounts depending on which options the buyer selects and which rates the company happens to be offering. The most useful skill a buyer can bring is not picking the single best annuity, which does not exist, but knowing which questions decide the outcome. The rest are those questions.

How the Interest Is Calculated

Start with the engine, the part of the contract that turns an index’s movement into a number credited to the account. Three devices do that work, and a

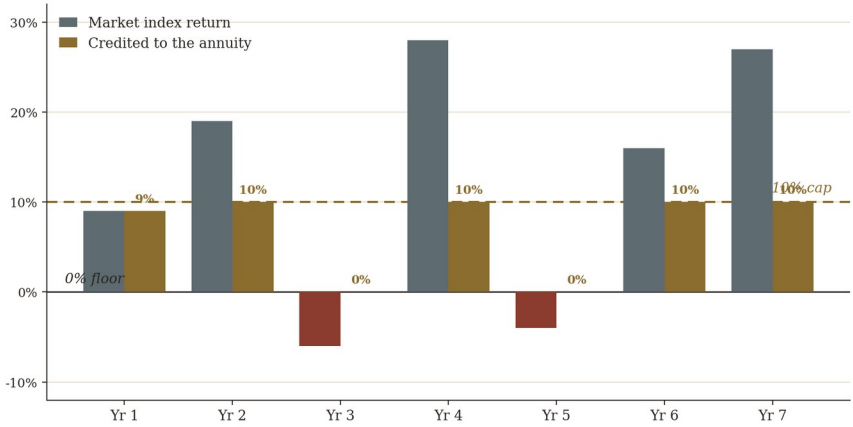
contract uses one of them or a combination. The first is the cap, already familiar: a ceiling on the credit for the period. If the cap is 10 percent and the index rises 19 percent, the account is credited 10 percent; if the index rises 6 percent, the account is credited 6 percent. The second is the participation rate, which credits a set percentage of the index’s gain rather than capping it. At a 60 percent participation rate, a 12 percent index gain becomes a credit of 7.2 percent, and there may be no ceiling above that. The third is the spread, which subtracts a fixed amount from the index’s return: with a 3 percent spread, a 10 percent index gain credits 7 percent. In every case the floor still holds. A losing index year credits zero, never less.

These devices can be layered, and the order matters. A contract might apply a participation rate first and then a cap, so a strong year is trimmed twice. Contracts also differ in when they measure the index. The most common method, annual point-to-point, compares the index on the contract’s anniversary with its value exactly one year later and credits the change, subject to the cap or participation rate. A monthly method instead adds up the index’s monthly moves, usually capping each positive month at around 2 or 3 percent while letting the down months count in full, an arrangement that can produce a small credit in a choppy year even when the index finished higher. The table below shows how the same index year lands differently under each annual approach.

Crediting method	What it does	Index up 12%	Index down 6%
Cap of 10%	Credits the gain up to a ceiling	10%	0%
Participation of 60%	Credits a set share of the gain	7.2%	0%
Spread of 3%	Credits the gain minus a set amount	9%	0%

Three ways a contract can turn the same index year into a credit. A 12 percent gain becomes a different number under each method, and any losing year credits zero. The rates shown are illustrative; real caps, participation rates, and spreads vary by contract and can reset over time.

Every crediting method shares one feature that the marketing rarely mentions. The index a fixed indexed annuity tracks is almost always the price index, which measures only the change in stock prices and leaves out dividends. An investor who owns the stocks directly collects those dividends; the annuity owner does not. Over the long run that gap has been worth roughly 1.5 percent a year. It is part of the cost of the protection, taken before any cap or participation rate is even applied, and it is one reason the credited return of an indexed annuity should never be set side by side with the headline return of the stock market as though the two were the same measure.



How a cap and a floor reshape a market index over time. In strong years the cap gives up the gain above its ceiling; in down years the floor credits zero instead of a loss. The result is steadier but lower than the index itself. A 10 percent cap and a 0 percent floor are shown against an illustrative seven-year index path.

The Rates Do Not Stay Put

A number on an illustration can give the impression that the cap is a permanent feature of the contract. For most contracts it is not. The cap, participation rate, or spread shown at purchase is typically guaranteed only for the first crediting period, often a single year. After that the company can reset it, up or down, at its discretion, for each period that follows. A contract that opens with a 10 percent cap might renew at 8 percent, or 6 percent, depending on conditions and on how hard the company chooses to compete for money it already holds. There is a backstop: every contract names a guaranteed minimum, a level the cap or rate can never fall below, such as a floor of 1 percent on the cap or a minimum fixed rate written into the contract. But the gap between the opening rate and the guaranteed minimum is often wide, and the opening rate is the one printed in the marketing.

Some contracts handle this differently, and the difference is worth hunting for. A minority guarantee the opening cap or participation rate for the entire surrender period rather than a single year, so a buyer who locks in a 10 percent cap keeps it for the full seven or ten years. That feature has real value, because it removes the risk that the rate is generous in year one to win the business and thin thereafter. The practical question for any contract is therefore not only what the opening rate is, but for how long it is guaranteed, and what the company's renewal rates have been over the past five to seven years. A company willing to show that history is making a different kind of promise than one that will not.

Why the Insurer Can Promise a Floor

A floor that keeps a bad market from ever reducing the account can sound too good to be true. It is not, and following the mechanics explains why caps move the way they do. When a premium comes in, the company puts most of it, very roughly 90 percent, into safe bonds. Over the surrender period those bonds earn enough interest to grow that portion back to the full premium, which is what funds the floor; even if the index does nothing for years, the bond interest rebuilds the account. The smaller remainder buys options on the index, contracts that pay off if the index rises and expire worthless if it does not, which is what funds the upside. The cap is set by how much upside that option budget can buy.

This structure explains a fact that surprises many people. Caps move in the same direction as interest rates, not against them. When rates are higher, the

bond portion rebuilds the principal faster, which leaves a larger budget for options and allows a higher cap. The rate increases of 2022, painful as they were for existing bondholders, pushed indexed-annuity caps upward, because new contracts could suddenly afford more upside. The same logic explains why caps were so thin in the years after 2008, when rates sat near zero and the option budget was small. A cap is not a number the company picks freely; it is the output of the bond market the company operates in.

Reading an Illustration

Most of these moving parts come together in the document a buyer is handed before purchase: the illustration. An illustration projects the account's value years into the future, and it almost always shows two sets of numbers side by side. The first is the guaranteed column. It assumes the index credits zero every single year and applies only the contract's guaranteed minimum, so it shows the worst the contract can legally do. On a typical contract that minimum is low, often crediting somewhere between 1 and 2 percent a year through a nonforfeiture guarantee, and the guaranteed column reflects it. This column is a real floor, and it is the only column the company is promising.

The second set is the non-guaranteed column, and it is the one that draws the eye, because the numbers are far larger. It is also the one most likely to mislead. The non-guaranteed column is built by taking a stretch of past index history, usually the most recent ten years, and assuming it repeats, while holding today's cap fixed for the entire projection and folding in any one-

time bonus. A column built that way can show an average return near 6 or 7 percent and a balance that more than doubles. None of it is promised. It assumes the next decade looks like the last, that a cap guaranteed for only one year stays put for ten, and that a bonus paid once keeps compounding. The carrier’s own disclosure, usually in small print on the same page, says as much: the values are not estimates of what the buyer should expect. A fair way to read an illustration is to treat the guaranteed column as the floor, the non-guaranteed column as a single optimistic scenario rather than a forecast, and the truth as sitting between them, closer to the 5 to 7 percent long-run range that history and the cap structure together suggest.

Riders, Bonuses, and What They Cost

A frequent question about annuities is where the fees hide. For a plain fixed indexed annuity, the surprising answer is often nowhere: many carry no explicit annual fee at all, because the company’s margin is built into the cap rather than charged as a separate line. That makes the plain version cheaper to hold than a typical variable annuity, which layers on yearly charges. The same advantage shows up against a bond fund, the usual home for conservative money: a bond fund charges an annual expense ratio that trims a yield often lower than the indexed annuity pays to begin with, while the plain contract takes no such ongoing cut. Costs enter mainly through optional riders, and the most common is an income rider, which provides a stream of lifetime income and typically costs around 1 percent of the account value each year. For a buyer

who wants the income guarantee, that fee may be worth paying; for one who does not, it is a charge to decline rather than to accept by default.

Other riders cost nothing and are worth knowing about. Many contracts waive their surrender charges if the owner enters a nursing home or is diagnosed with a terminal illness, which turns a locked-up contract into an accessible one at exactly the moment a household may need the money; these waivers commonly carry no charge, though they are not offered in every state. A different feature deserves more caution: the premium bonus, sometimes called an interest boost, which adds a one-time credit to the account at the start. A bonus is real money, but it is rarely free. The company funds it by setting a lower cap, a longer surrender period, or a schedule under which the bonus is forfeited if the contract is surrendered early. The way to weigh a bonus is to ignore the headline and ask what was given up to pay for it.

A bonus often appears for a particular reason. When a buyer moves money from one annuity into another, leaving the old contract early can trigger its surrender charge, and a new contract’s bonus is frequently sized to cover that charge, so the switch looks costless. Sometimes the move is sound, when the new contract genuinely offers better rates, a shorter remaining lockup, or terms that fit a changed situation. Other times the bonus only hands back, through a lower cap or a fresh surrender schedule, money the buyer is paying a penalty to free up in the first place. The transfer itself can usually be made without tax,

through what is called a 1035 exchange, but tax-free does not mean cost-free. The only way to know whether a switch makes sense is to run the numbers: weigh the surrender charge given up and the new surrender period taken on against the added value the new contract really delivers.

Getting Out: Surrender and the Market Value Adjustment

Liquidity is where a fixed indexed annuity asks the most of its owner, and the terms reward a close reading. The surrender charge is a penalty for taking out more than the contract allows during its early years; a typical schedule starts near 9 percent in the first year and steps down to zero over five, seven, or ten years, depending on the contract chosen. Most contracts soften this by letting the owner withdraw a set share each year, commonly 10 percent of the value, with no charge at all, which is enough for the same money to double as the volatility buffer. A second adjustment is less familiar and catches people off guard. The market value adjustment, or MVA, applies on top of the surrender charge to withdrawals above the free amount, and it moves with interest rates: if rates have risen since the contract was bought, the MVA reduces the amount paid out, and if rates have fallen, it increases it. The adjustment can cut either way, but its existence is one more reason an indexed annuity should hold only money the owner is genuinely willing to leave alone for the length of the surrender period.

Comparing Two Contracts

Put the pieces together and a short list of questions does most of the work of telling one contract from another. The point of the list is not to crown one winner but to make the differences visible, because the differences are where the money is. The questions below are the ones that change the outcome.

What to ask	Why it matters
Which crediting method, and at what cap, participation rate, or spread?	These set how much of an up year reaches the account. The opening rate is only a starting point.
How long is the opening rate guaranteed, and what is the guaranteed minimum?	A rate locked for a single year can reset much lower. The minimum is the true worst case.
What have the company's renewal rates been over the past five to seven years?	Past renewals show whether the opening rate was bait or a lasting offer.
Which index, and is it the plain S&P 500 or a proprietary one?	Proprietary volatility-controlled indexes pair high participation rates with lower built-in returns, so the two numbers can mislead.
How long is the surrender period, and how do the schedule and market value adjustment work?	This is the money's lockup. It should match cash the household will not need.
How much can be withdrawn free each year?	Usually about 10 percent. It sets everyday access and whether the contract can double as a buffer.
Which riders apply, which cost extra, and does any bonus carry strings?	Income riders run about 1 percent a year. A bonus is paid for elsewhere in the contract.

What to ask	Why it matters
How financially strong is the issuing company?	The guarantee rests on its ability to pay claims, backed only secondarily by a state limit.

The questions that separate one fixed indexed annuity from another. The opening rate is the most advertised number and rarely the most important one.

The last question on that list carries more weight than its place at the bottom suggests. An annuity's guarantee rests on the issuing company's ability to pay its claims rather than on federal deposit insurance. The practical version of that principle is to check the company's financial strength rating from an independent agency before buying, and to favor the stronger carriers for a commitment measured in years. It helps to keep that risk in proportion. Insurance companies fail far less often than banks, partly because they are held to strict reserve requirements and cannot be drained by a sudden run the way a bank can; 2008 made the point, when hundreds of banks failed while the insurance industry's annuity obligations came through largely intact. Outright losses to annuity owners are rarer still. The 1991 failure of Executive Life, the case usually cited, left some annuity holders waiting years and recovering only part of what they were owed. It failed for a specific reason: it had loaded its portfolio with junk bonds, the riskiest and lowest-rated kind, which is precisely the asset-quality weakness a financial strength rating is built to catch. Every annuity carrier is graded by independent rating agencies, so the simplest due diligence, checking that grade and

staying with the stronger companies, steers a buyer clear of the contracts most likely to run into trouble. For a highly rated carrier, an annuity owner losing money is close to unheard of.

A secondary backstop sits behind the company itself. Every state runs a guaranty association that covers annuity owners if an insurer fails, commonly up to $250,000 in present value per owner per company, though the exact figure is set by each state. That backstop is real, but it is narrower than federal deposit insurance and is not meant to be the first line of defense. The first line is a financially sound company, which is why the rating belongs on the checklist at all.

When a Simpler Annuity Will Do

Everything to this point describes a contract with moving parts. Some people want the protection and the tax deferral but none of the machinery, and for them a simpler cousin exists: the multi-year guaranteed annuity, usually shortened to MYGA. A MYGA strips out the index, the cap, the participation rate, and the options. In their place is a single feature: one fixed interest rate, set in advance, for a term the buyer chooses, commonly anywhere from two to ten years. There are no good years and bad years to track. The rate is the rate, the same every year of the term, known the day the contract is signed.

The closest familiar relative is the bank certificate of deposit. Where the indexed annuity stands as an alternative to bonds, the multi-year guaranteed annuity is best understood as an alternative to the certificate of deposit, and the comparison is exact

enough to be useful. Both lock up money for a set term in exchange for a stated rate, and both penalize early withdrawal. A MYGA differs in three ways that tend to favor it for retirement money. It usually pays more: in 2026 a multi-year guaranteed annuity from a well-rated company paid somewhere around 5 to 6 percent, against roughly 4 percent for a comparable certificate of deposit, a gap of a point or two that compounds over the term. Its growth is tax-deferred, so no tax is owed on the interest until the money is withdrawn, while a certificate's interest is taxed every year. And like the indexed annuity, it usually permits a free withdrawal of about 10 percent a year. What the certificate has that the MYGA does not is federal deposit insurance; the MYGA rests instead on the same insurer strength and state guaranty backstop described above.

Two further details make the MYGA useful in places the indexed annuity is not. It can usually hold inherited retirement money, which matters for a beneficiary who must keep funds in an inherited account while the law's ten-year withdrawal clock runs and wants a known return in the meantime. And at the end of its term a MYGA typically opens a short window, often 30 days, during which the owner can walk away with no surrender charge at all, move to a new term, or take the money, a clean exit the indexed annuity's longer schedule does not offer mid-contract. The MYGA fits the person who wants a known number and zero moving parts, who values certainty over the chance at index-linked upside, and who has a clear time horizon to match against the term.

Matching the Tool to the Job

Robert and Susan Whitfield, a fictional couple, show how these choices play out once they are concrete. They had been weighing whether to hold part of their conservative money in a fixed indexed annuity rather than entirely in bonds. With the contract opened up, their decision becomes more than a single yes or no. They might place one slice in an indexed annuity whose cap is guaranteed for the full term, for protected money that still reaches for index-linked growth, and another slice in a multi-year guaranteed annuity, for money whose job is simply to earn a known rate with no surprises. Two different tools, two different jobs, both inside the safe sleeve.

Linda Castellano, a fictional single saver, shows the opposite lesson, the one worth carrying away. Linda has an inherited account she must empty within a couple of years. A ten-year indexed annuity, whatever its cap, would be the wrong home for that money, because its surrender period would outlast the deadline and lock up funds she is required to withdraw. The mistake would not be the annuity; it would be the mismatch. The same product that fits the Whitfields' decades-long horizon fails Linda's two-year one. The lesson underneath every one of these questions is that a fixed indexed annuity is neither good nor bad on its own. It is a tool with a specific shape, and the only real question is whether that shape fits the job the money has to do.

It comes down to this. A fixed indexed annuity is not the miracle its marketing sometimes suggests or the trap its critics sometimes claim. It is a contract with a floor that does what a bond fund cannot, an upside that is real but limited, rates that move, terms that lock, and costs that sit mostly inside the cap rather than charged outright. A buyer who understands those features can tell a sound contract from a poor one and match either to the right job. A different tool, built on a cousin of the same machinery, comes next: indexed life insurance, with its particular jobs of tax-advantaged income and leaving a legacy, where it does its best work.

Chapter 13: Indexed Universal Life — Tax-Free Income and Legacy Planning

The crediting method that runs a fixed indexed annuity is also used in a type of life insurance. That product is indexed universal life, usually shortened to IUL, and the family resemblance is real. Its cash value grows on the same indexed crediting as the annuity, with a cap that limits the gain in a strong year and a floor of zero that prevents a loss in a weak year. But the resemblance stops at the cash value, because indexed universal life is first of all life insurance. It does what every life insurance policy does: it pays a sum of money to the people the insured names, at the insured's death. That payment, the death benefit, arrives income-tax-free to those beneficiaries, and it is the reason the policy exists at all. Everything else the product can do rests on that death benefit, chiefly the unusually generous and little-understood tax treatment of the cash value it builds, and the ability to build that cash value faster by paying more premium than the policy strictly requires. Both deserve a closer look.

What draws most of the attention, though, is what the policy can do while the insured is still alive. The cash value inside it grows without a yearly tax bill, and money can be taken out later without a tax bill as well, through a mechanism worth walking through carefully. That combination, lifetime access to money on tax-advantaged terms plus a tax-free payment to heirs, is what makes indexed universal life worth understanding on its own terms. It is also among the most aggressively marketed products in all of personal

finance, wrapped in illustrations that promise more than the contract can deliver. Both sides deserve a serious look: what the policy does well, and who is better served by something simpler.

The Third Bucket

To see where indexed life insurance fits, it helps to recall one idea: retirement savings fall into three tax buckets. The first is the taxable bucket. It holds money in an ordinary brokerage or bank account, where interest, dividends, and gains are taxed as they happen. The second is the tax-deferred bucket. It holds the traditional 401(k) and IRA, where nothing is taxed going in or along the way, but every dollar is taxed as ordinary income on the way out. The third is the tax-free bucket, where money is taxed going in but never again. The Roth account is the familiar example, and a health savings account spent on medical costs is another. The cash value of certain permanent life insurance, indexed universal life included, can also be reached on a tax-free basis within limits, which places it in this third bucket.

> Holding money across all three buckets means no single future tax rate can reach the whole pile. Most people arrive at retirement with nearly everything in the tax-deferred bucket, which leaves every withdrawal exposed to whatever tax rates Congress sets decades from now. Spreading savings across all three buckets blunts that risk. These fictional households illustrate the spread. Dave Thompson’s traditional IRA sits in the tax-deferred bucket; his wife Mary’s Roth and her

health savings account sit in the tax-free bucket. Linda Castellano's workplace account holds both treatments at once, a traditional portion and a Roth portion. The Whitfields hold a large taxable brokerage account alongside their tax-deferred IRAs. Indexed life insurance is one more way to build out the tax-free bucket, used by people who have a specific reason for it.

One common product resists a single bucket. A fixed indexed annuity grows tax-deferred, so its earnings are taxed as ordinary income when they are withdrawn. But when the annuity was bought with already-taxed dollars, the original principal comes back out tax-free, the way already-taxed money in the taxable bucket does. An annuity held inside an IRA sits squarely in the tax-deferred bucket; an annuity funded with after-tax savings straddles two at once. The three buckets are a guide to how money is taxed, not a rule every product obeys.



The three tax buckets. Money is taxed on the way in, along the way, or on the way out, depending on the bucket it lands in. Indexed universal life is one way to build out the third, the tax-free bucket.

Inside the Policy

Underneath the marketing, an indexed universal life policy works on a simple division of every premium dollar. Part of each payment covers the cost of the life insurance itself, the charge the company levies for promising the death benefit, along with administrative and policy fees. Whatever is left over goes into the cash value, where it is credited using the familiar indexed method: a cap, a participation rate, or a spread applied to the movement of a market index, with a floor of zero in a losing year. Recent caps on policies tied to the S&P 500 have run in a range of roughly 8 to 12 percent. The lessons from the annuity carry straight over: the tracked index is usually a price index and leaves out dividends, and the cap can be lowered by the company after the policy is issued.

One feature separates this product from the annuity and drives almost everything about how it behaves over time. The cost of insurance is not fixed. It rises every year as the insured ages, because the company’s risk of having to pay the death benefit climbs with age, and that rising charge is deducted from the cash value whether the index credited a gain that year or nothing. Early on, when the insured is young and the charge is small, the credits cover it easily and the cash value grows. Late in life the charge can become large, and in a flat year, when the floor credits zero, there is no gain to absorb it, so the charge comes straight out of the cash value. A policy not fed enough premium to stay ahead of that rising cost can begin to consume itself.

The figure traces how the charge climbs with the insured's age.

Funding Is Everything

A second feature sets indexed universal life apart, and it explains both its flexibility and its danger: the premium is not fixed either. Term and whole life insurance charge a level premium on a fixed schedule, but a universal life policy lets the owner vary what is paid from year to year, within limits. In a prosperous year the owner can pay in more, and in a lean year less. That freedom is useful, but it is not a license to improvise. A sound policy is built around a planned premium, chosen with an illustration that projects the likely path of the cash value, and it works only if that planned premium is actually paid. Because the cost of life insurance climbs every year, how generously the policy is funded matters more than any other single choice.

Two people can buy the same policy and end up in opposite places. The buyer who funds it generously in the early years builds a cash value large enough that its own growth outruns the rising charges. The buyer who pays the bare minimum may watch the charges hollow the policy out and, in a bad stretch, push it to lapse. For the cash-value purpose, the aim is usually to pay in as much premium as the tax law allows relative to the death benefit, which front-loads the account and spreads the fixed charges over a larger base. There is a ceiling on this, set by the tax code, and crossing it carries consequences worth knowing.

That ceiling marks the line between life insurance and what the tax code calls a modified endowment contract, or MEC. Pay in too much too fast, relative to the death benefit, and the policy crosses that line and loses the tax treatment that made it attractive: withdrawals and loans become taxable as the gain comes out first, with a penalty on top before age fifty-nine and a half. The practical effect is a narrow band. A policy aimed at cash value wants to be funded close to the maximum the law permits but not over it. The reassuring part is that this line is well marked and hard to cross by accident. The federal limit is fixed by a formula known as the seven-pay test, and the insurer measures every premium against it, flagging or refusing a payment that would tip the policy over. A policy can also be designed with a larger death benefit, which raises the limit and makes room for more premium. Building a policy to ride just under that line, rather than to show the largest death benefit for the smallest premium, is one of the clearest signs of a contract meant to accumulate cash value rather than to look impressive on an illustration.

Taking Tax-Free Income

Income that gives indexed life insurance its reputation comes out through a specific door, and understanding that door is the whole point. Money is not withdrawn from the cash value; it is borrowed against it. The policy owner takes a loan, with the cash value standing as collateral, and the insurance company advances the money from its own funds. Here is the feature people most often misunderstand: the cash value is not reduced by the loan. It stays in the policy in full and

keeps earning interest on the whole balance, as though nothing had been taken out, even while the borrowed money is in the owner's hands. Because a loan is not income, the tax code does not treat it as a taxable event, which is the entire basis for calling the proceeds tax-free. What keeps this from being a trick is that the loan is never required to be repaid during the insured's life. It is settled at the end, out of the death benefit.

That settlement is the catch, and it is the part the illustrations tend to mute. Every dollar borrowed, plus the interest that accrues on it, reduces the death benefit that eventually reaches the beneficiaries. Income taken during life and legacy left at death draw on the same pool, so taking more income now leaves a smaller benefit later. There is a sharper danger as well. If the loans and their interest grow faster than the cash value behind them, the policy can collapse under its own debt, a situation known as overloan. A lapse while a large loan is outstanding does not only end the coverage; it can convert all the gain that was borrowed tax-free into taxable income in a single year, turning a tax-free plan into a tax bill at the worst possible moment. Many modern policies include an overloan protection feature that guards against this by capping the borrowing and locking in a reduced but permanent death benefit, and on a policy built to supply income, that protection earns its place.

Loan mechanics also vary, and the variation changes the risk. A fixed loan charges a set interest rate and moves the borrowed amount to a stable account, so the arithmetic is predictable. A participating loan leaves

the borrowed amount in the indexed account, where it keeps earning credits, while charging interest on the loan: when the credits beat the loan rate the owner comes out ahead, and when they trail it the gap compounds against the policy. The second type is what produces the eye-catching income figures in an aggressive illustration, and it is also what can hollow a policy out if the index disappoints. A conservative design leans on the predictable loan and treats the participating loan's upside as a possibility, not a plan. The sound practice is for the professional to illustrate the plan at a conservative rate, and for the owner to understand and agree to that plan before the funding begins.

A Second Volatility Buffer

Set against retirement income, the cash value can do one more job. Because it does not fall when the market falls, it can serve as a second volatility buffer, a pool of money to draw on in a year when stocks are down so that investments are not sold at a loss to pay the bills. A retiree who can borrow a year of spending from a policy after a market drop gives the rest of the portfolio time to recover, the same sequence-of-returns protection a fixed indexed annuity offers, with the added feature that the borrowed money arrives tax-free. The value is real, but it depends entirely on the cash value being there, which loops back to funding: only a policy funded consistently for years holds enough inside it to play this role.

Reading the Illustration

No product leans on its illustration the way indexed universal life does, and no illustration deserves more skepticism. The document a buyer is handed projects decades of smooth growth from a single assumed rate, holds a cap steady for forty years that the company can change next year, and assumes a cost of insurance that in reality keeps climbing. Regulators have noticed. A series of national rules, tightened twice in recent years, now limits how rosy an illustration may look, capping the rate a carrier can show and curbing the bonus features companies once used to inflate the numbers. Those rules govern the illustration, not the policy, so a contract can still perform better or worse than the page suggests. Because any single illustration rests on one set of assumptions, a careful professional shows several, run at different rates, so the buyer sees a range of outcomes rather than a single hopeful projection. Even at the capped rate, any illustration is just one optimistic scenario, not a forecast.

Reading an illustration takes the same discipline the annuity did. The guaranteed column, which assumes the worst the contract legally allows, is the only set of numbers the company is committed to, and on a policy built for cash value it can look alarming. The illustrated column beside it is a single hopeful scenario. There is a realistic middle between the two. After caps, fees, and the rising cost of insurance are all accounted for, the cash value of a well-funded policy tends to grow over a long horizon in the range of 4 to 7 percent a year. That is closer to the result of a fixed indexed annuity than to the stock market the index name evokes. A buyer can pressure-test any

illustration by asking to see it run at a lower rate, and by asking what happens to the policy if the cost of insurance climbs faster than assumed once the insured reaches their seventies.

The Death Benefit Is the Point

For all the attention the living benefits attract, the reason to own indexed universal life remains the death benefit. The policy pays a sum to the insured's beneficiaries at death, and that sum arrives free of income tax. For a household that wants to leave money behind, that is a powerful and simple thing: a known, income-tax-free amount delivered to heirs outside the delays of probate, funded by premiums smaller than the benefit they buy. Where a traditional IRA left to children lands in their hands as fully taxable income they must drain within ten years, a life insurance death benefit lands tax-free. That contrast is the heart of the legacy case.

Robert and Susan Whitfield, a fictional couple, show how that case can fit a real plan. They have more than they expect to spend and a clear wish to pass wealth to their three children efficiently. For them, indexed universal life offers a direct route. They can move a slice of taxable IRA or brokerage money into a policy over time, at a controlled tax cost, and turn it into an income-tax-free benefit for the next generation. The policy becomes the tax-free container for the legacy, working alongside Roth conversions rather than instead of them. For a household at their level, with the other buckets already full, the death

benefit is doing exactly the job the product does best.

Where It Fits, and Where It Does Not

The case against the product matters as much as the case for it. The clearest example is pure protection during the working years, where term life insurance paired with ordinary investing wins outright. A young family that needs a safety net if a breadwinner dies can buy term life insurance, pure death benefit for a set number of years, for a small fraction of an indexed policy's premium, and invest the difference in a Roth IRA or a workplace retirement plan. For that family, with a temporary need, that combination is cheaper and usually ends with more money. But term and permanent insurance answer different questions, and the comparison is often made unfairly. Term insurance expires, and a policy that expires can leave nothing behind. It cannot move wealth to the next generation, because by design it is gone before the next generation inherits. A household whose goal is a permanent, income-tax-free transfer cannot reach that goal with term insurance at any price.

Two cautions still apply, each with a flip side. The first is sequence. The cheaper corners of the tax-free bucket should usually be filled before a policy is considered, because a Roth IRA, a Roth workplace account, and a health savings account all reach tax-free territory with none of the cost of insurance and far less complexity. Indexed universal life makes sense mainly once those are full and the household still wants more tax-free room. For higher earners that

point arrives sooner than it sounds, because the Roth and the health savings account carry low annual limits that substantial income exhausts quickly. This is exactly why the affluent so often turn to indexed universal life. It is one of the few remaining places to build a large, tax-free pool once the smaller accounts are maxed, and it doubles as a tax-free source of income in retirement, which matters most to those whose tax rates are high now and may climb higher later.

The second caution is age, and it is the most often misread. Because the cost of insurance rises with the insured's age and the cash value needs a long runway, a policy that insures someone close to retirement has little time to work and a steep cost curve just ahead. But the insured and the owner need not be the same person. A grandparent can own and fund a policy on the life of a young child or grandchild, where the cost of insurance starts low and the cash value has decades to compound, and can draw tax-free income from it as the owner while it is still building, then pass it on. Read that way, the age question is about whose life the policy covers, not about who may benefit from it.

> Dave and Mary Thompson, in their sixties, would gain little from a new policy on their own lives, but a policy on a young grandchild, funded by them and built into their plan, could serve both their income and their legacy.

Matching the Tool to the Job

Indexed universal life fits a specific profile, though a wider profile than the skeptics allow. At its center is a long-term horizon, good health so the cost of insurance starts low, the cheaper tax-free accounts already filled, a lasting reason to carry permanent life insurance, and a willingness to fund and mind the policy for decades. A person who buys young, or an older person who funds a policy on a young life, gives the cash value the years it needs to outgrow the charges, and the policy can then supply both tax-free income and a legacy. Linda Castellano’s daughter Maya, not yet thirty and decades from retirement, is a natural fit on her own life. Linda, in her late fifties, is a weaker fit on her own life, yet the same goal is within reach if she funds a policy on Maya and draws on it later. The product can serve more than one generation at once, which is much of its appeal to families thinking past a single lifetime.

Pulling the threads together, a short list separates the situations where indexed life insurance shines from the ones where a simpler tool will serve. The questions are about the buyer and the goal, not the policy, because the policy is only ever as good as its match to the job.

Indexed universal life fits when	A simpler tool fits better when
The goal is a permanent, income-tax-free transfer to the next generation	The need is temporary protection during the working years, which term covers cheaply
The cheaper tax-free accounts, Roth, workplace Roth, and HSA, are already full	Roth, workplace, and HSA room still sits unused
A long horizon is available, whether the insured is the buyer or a young child or grandchild	Neither the buyer nor a younger insured has decades for the cash value to build
The owner will fund the policy generously and mind it for years	There is no desire to fund and manage a long-term contract
Tax-free retirement income matters, especially at high or rising tax rates	Current and expected tax rates are low enough that the tax-free feature adds little benefit

Indexed universal life is a specialist's tool: powerful for permanent transfer and tax-free income, and easy to misuse for needs that simpler accounts serve better. The deciding factors are the goal, the horizon, and the willingness to fund and manage the policy.

That is the balanced verdict on indexed universal life. It is neither the tax-free miracle its boosters claim nor the trap its critics describe. It is a life insurance policy first, paying an income-tax-free benefit to the people who outlive the insured, with a cash value that can grow on tax-advantaged terms and be borrowed against for tax-free income, all of it resting on steady

funding and realistic expectations. For the household with a long horizon, a lasting need for permanent coverage, and the cheaper accounts already full, it does work that few other tools can match, above all the work of moving wealth to the next generation and drawing tax-free income along the way. For a family with a shorter horizon or a simpler need, the cheaper tools come first. Used for the right purpose, by the right buyer, it is one of the most powerful instruments available. The larger task all of these products serve comes next: weaving the taxable, tax-deferred, and tax-free buckets into one coordinated plan, so that required distributions, Roth conversions, and Medicare’s income tests work together across the years rather than colliding.

Chapter 14: Tax Mitigation in Retirement — Building a Multi-Year Plan

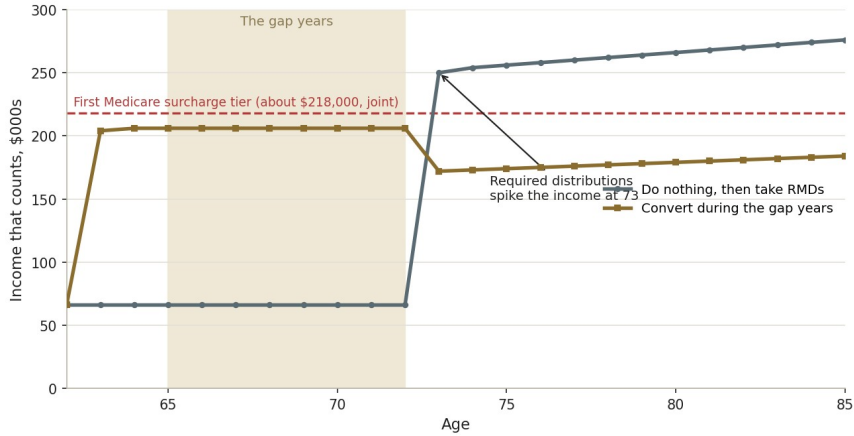
Several separate pieces make up the retirement-tax picture. Required minimum distributions are the tax the law eventually forces out of a traditional account. The Roth conversion is the deliberate move of money into the tax-free bucket. The Medicare surcharge is triggered by higher income, and savings sort into three tax buckets. Each part matters on its own. Coordinated across years rather than handled separately, the parts do far more, and that coordination is the real work. A retiree holds something most taxpayers do not: real influence over how much income lands in a given year, and therefore over the rate at which that income is taxed.

That control is the whole opportunity, and using it well means changing the question. The instinct most people carry out of their working years is to keep this year’s tax bill as low as possible. In retirement that instinct can be expensive, because a dollar of tax avoided now at a low rate can force several dollars of tax later at a higher rate. The better aim is to smooth the tax bill across the whole of retirement, paying a little more in the lean years to pay far less in the heavy years. A low-income year is more than a break; it is an opportunity that expires. A plan that sees the years as a set, rather than year by year, is what turns scattered tactics into real savings.

The Shape of the Problem

Retirement income rarely arrives in a level stream, and its uneven shape is what creates both the opening and the trap. For many people the income drops sharply once the paycheck stops, then drops further if Social Security is delayed, leaving a stretch of unusually low-income years. Then, often in the same year that delayed Social Security and required distributions both switch on, it jumps back up and stays there. The years in the dip are taxed at low rates almost by accident. The years after the jump can be taxed at high rates, and the jump itself can push a household through tax brackets and Medicare thresholds it never crossed while working.

In one sentence, the plan is to fill the low years on purpose so the high years are not forced higher than they have to be. Money pulled from a traditional account during the lean stretch, or converted to a Roth, is taxed at the low rate that stretch offers and is no longer sitting in the account to be forced out later at a higher rate. The same total is taxed either way; the difference is the rate, and the rate is what the timing controls. The figure below shows the shape of the choice for a single household.



Two paths for the same household. Doing nothing keeps income low in the gap years, then lets required distributions spike it past a costly Medicare threshold at seventy-three. Converting during the gap years lifts income deliberately, but keeps it under the line in every year. The figures are illustrative.

The Three Phases

Across three phases of retirement, the opportunity changes character, each phase defined by what has and has not started yet. The table sets them side by side. They are not sharp lines, and a given household may move through them faster or slower, but the sequence holds: a stretch before Medicare, a stretch on Medicare but before required distributions, and the long stretch once those distributions begin.

Phase	Typical ages	What has and has not started	The main move
Before Medicare	Up to 64	No required distributions; often low income; individual-market health coverage in play	Convert or realize gains in a low bracket, watching the subsidy cliff
The gap years	65 to 72	On Medicare, so surcharges now apply; Social Security perhaps delayed; distributions not yet forced	Fill a bracket with conversions, staying under the next surcharge tier
Required distributions	73 and up	Distributions force income; Social Security is on; the survivor's higher rates loom	Use charitable distributions, draw the order deliberately, live on a base already lowered

The three phases of retirement tax planning. The lines are approximate and a household may move faster or slower, but the sequence holds: a window before Medicare, a wider window before required distributions, and the long stretch once they begin.

One thread runs through all three phases: early work shrinks the later problem. Every dollar moved out of

the traditional account before age seventy-three, whether spent or converted, is a dollar that will not be in the year-end account balance that sets the required distribution, and so is a dollar not forced out at whatever rate applies then. The household that does nothing in the first two phases does not avoid the tax; it merely defers the whole bill into the years when it has the least control left, and often the highest rates.

The Gap Years

The middle phase, the years on Medicare but before required distributions begin, is where the most valuable work usually happens, and it has a name worth knowing: the gap years. Social Security may not have started, especially for someone delaying it to seventy, and the required distributions have not begun, so taxable income can be at its lowest point in decades. That low floor is exactly what makes room for deliberate income. A household can convert traditional money to a Roth, or simply draw it down, filling the low brackets that would otherwise go unused, and pay tax at a rate it may never see again once the distributions and the full Social Security check arrive together.

Robert and Susan Whitfield, a fictional couple, sit squarely in this window. Robert is delaying his own Social Security to seventy, which leaves the years from his late sixties to seventy-two unusually light, and the couple has a large traditional balance waiting to become required distributions. Their plan is to convert a measured amount in each of those years, enough to fill the

bracket they have chosen as a ceiling, paying the tax from their brokerage account so the whole converted sum lands in the Roth. The fictional Dave and Mary Thompson have a smaller version of the same window. Mary’s pension and the couple’s Social Security cover most of their needs, so their lean years are less dramatic, but the principle is identical: a modest conversion at the 12 percent rate now is cheaper than a forced distribution at a higher rate later.

Watching the Guardrails

Filling the low brackets is not a blank check, because the tax code is full of thresholds where one more dollar of income costs far more than the dollar before it. These are the guardrails of any conversion plan, the lines to fill up to but not cross without a reason. Some are ordinary bracket edges, where the next dollar is simply taxed at a higher rate. Others are cliffs, where crossing the line by a single dollar triggers a cost on income well below it. The table gathers the thresholds that matter most in retirement.

Threshold (2026)	Roughly where it sits	Why it matters
Top of the 12% bracket	About $100,800 taxable, joint	Income above it is taxed at 22 percent and can begin taxing long-term gains
Top of the 22% and 24% brackets	About $211,400 and $403,550 taxable, joint	Common ceilings a conversion is filled up to, not past
Zero-rate capital gains	Up to about $98,900 taxable, joint	Below it, long-term gains and qualified dividends are taxed at zero
First Medicare surcharge tier	$218,000 joint, or $109,000 single, on income from two years earlier	One dollar over adds the full surcharge for the year
Net investment income tax	$250,000, joint	An extra 3.8 percent on investment income above the line
Health-subsidy cliff, before 65	Around four times the poverty line	Where individual-market premium help can drop sharply

Income thresholds worth watching in 2026, for a married couple unless noted. Bracket tops are measured on taxable income, while the surcharge, the investment-income tax, and the health subsidy use a

broader measure of income. Every figure should be re-verified before it is relied on.

Medicare’s income surcharge deserves particular attention, because it is a cliff and it runs on a delay. Its basics are familiar; the planning point here is its timing. The surcharge for a given year is set by income from two years earlier, so a large conversion at sixty-three can raise Medicare premiums at sixty-five, and a conversion at sixty-six can raise them at sixty-eight. In 2026, the first surcharge tier begins at $109,000 of income for a single filer and $218,000 for a couple, and crossing either by a single dollar adds the full surcharge for the year, not a little on the dollar over. Two rules follow: do the largest conversions before sixty-three, if possible, where they cannot reach back into Medicare premiums, and once on Medicare, size each conversion to stop just short of the next tier.

Other guardrails matter at different moments. Before sixty-five, a retiree who buys health coverage on the individual market may find that income drives the size of the premium subsidy, and, depending on the rules in force, a sharp cliff can appear around four times the federal poverty line, which can make a large conversion before Medicare costly in a way that has nothing to do with income tax. At the other end of the income scale sits a less obvious opportunity: a household whose taxable income is low enough can have its long-term capital gains and qualified dividends taxed at zero, which makes a lean year a chance to reset the cost basis on an investment rather than to convert. The guardrails interact, and a plan that watches only the income-tax brackets can walk straight into a surcharge or a lost subsidy it never saw.

The Order of Withdrawals

Alongside the question of how much income to create sits the question of which account to draw it from, and here a rough order of operations helps. The conventional sequence spends the taxable bucket first, then the tax-deferred bucket, and leaves the tax-free bucket for last. The logic is that the taxable account is already taxed as it grows, so spending it first stops that drag, while the Roth grows entirely tax-free and is worth protecting as long as possible, especially since it carries no required distributions and passes to heirs tax-free. Spending the taxed money first lets the sheltered money compound the longest.

The order is a starting point, not a rule to follow off a cliff. A household that empties its taxable account first and then takes nothing from the traditional account is wasting the low-bracket years the plan is built around, only to have the whole traditional balance forced out later at a higher rate. The better practice blends the two ideas: spend from the taxable bucket for daily needs while also pulling or converting traditional money each year up to a sensible bracket ceiling, so the tax-deferred balance shrinks steadily during the low years rather than waiting to erupt. The buckets are not a strict sequence; they are a set of taps to open in whatever mix keeps the lifetime rate low.

Coordinating with Social Security

When to claim Social Security is not separate from any of this; it sets the size of the gap-years window. Delaying a benefit to seventy raises it by about 8 percent for each year of waiting, and for a healthy person this is among the most efficient purchases in

retirement. It has a second effect that belongs here: every year Social Security is delayed is a year of lower income, and therefore a wider opening for conversions at low rates. The two strategies reinforce each other, with one caveat. The gap has to be funded from somewhere, usually the same accounts being drawn down or converted, so the claiming decision and the conversion plan have to be built together rather than separately.

When the Plan Meets Real Life

Not every household has a clean set of gap years to work with. Linda Castellano, a fictional single saver, faces two complications at once. She is younger than sixty-five and buying her own health coverage, so a large conversion could cost her a premium subsidy at the same time it costs her income tax, narrowing how much she can convert in any year. She also carries an inherited account she must empty within two years, which forces a stretch of high income at single-filer rates whether she wants it or not. For Linda the plan is less about creating income in lean years and more about not stacking avoidable income on top of the income she cannot avoid, which may mean holding her own conversions until the inherited account is finally empty.

Walter Higgins, another fictional retiree, shows the other edge, the cost of having no plan at all. At eighty-four, well into his required distributions, his low-income years are long behind him; the conversions that would have helped were

available decades ago and were not taken. His traditional balance is now forced out a slice at a time and taxed at whatever rate his other income sets, with little he can do to change it. One lever remains. After age seventy and a half, he can send up to a set amount each year directly from his traditional account to a charity, a qualified charitable distribution, which counts toward his required distribution but never appears in his income at all. For a retiree who gives anyway, it is the rare move that helps even this late.

The Limits of Optimizing

A plan this involved carries its own temptation, which is to optimize the taxes at the expense of everything else. The tax bill is one cost among several, and not always the largest. Money converted to a Roth is money that has left the spending account and cannot be reached without undoing the plan, so a household that converts aggressively can find itself cash-poor while technically tax-efficient. The future is unknown in ways no projection can fix: tax rates may rise or fall, health may turn, a spouse may die and shift the survivor into higher single rates. The aim is a plan robust across a range of futures, not a plan tuned perfectly to a single guess about what Congress will do.

Coordination is also genuinely hard, harder than any single piece of it. A conversion sized to fill a bracket can trip a Medicare surcharge two years later; a move that lowers this year’s tax can raise a lifetime’s; the right answer for one household is the wrong answer

for the household next door. This is the kind of problem where a multi-year projection, run by someone who can see the income, the brackets, the surcharges, and the household's actual spending all at once, often pays for itself. The point is not that the arithmetic is beyond anyone, but that the interactions are easy to miss, and the mistakes are the quiet, permanent kind that keep recurring.

Putting It Together

Pulled together, the Whitfields' plan becomes a decade-long sequence rather than a single move. In their late sixties and early seventies, with Robert's Social Security delayed, they convert traditional money each year up to the top of a chosen bracket, watching the Medicare tiers so a conversion never tips them into a higher surcharge, and paying the tax from their brokerage account so the full amount compounds in the Roth. They spend the brokerage account for living costs in the same years, which lowers its future taxable growth. Once required distributions begin, the balance left to be forced out is smaller than it would have been, and the Roth they built sits ready to supply income in any year when adding more would cross a costly line. After seventy and a half, the charitable gifts they were going to make anyway go straight from the traditional account, trimming the required distribution without adding to income.

Dave and Mary Thompson run the same logic at a smaller scale, which is the point worth taking

from the contrast. Their pension and Social Security cover the essentials, their traditional balance is modest, and their lean years are gentle, so their plan is a few measured conversions at the 12 percent rate and an eye on not letting a good year tip them into the next bracket. They will never face the Whitfields' surcharge tiers, and they do not need to. The framework is the same for both households, the levers are the same, and only the dials are set differently. A coordinated tax plan is not a luxury for the wealthy; it is a habit of looking at the years together, available to anyone with more than one bucket to draw from.

That habit is what all of this is really about. The pieces are not tricks. Required distributions, Roth conversions, the order of withdrawals, the Medicare and subsidy thresholds, the timing of Social Security: each is ordinary on its own. What makes them powerful is sequence and coordination, the decision to treat a thirty-year retirement as one connected tax problem instead of thirty separate problems. Done that way, the levers can save years' worth of taxes and a stream of avoidable surcharges. Left uncoordinated, they collide, and the collision is expensive.

Every plan here has worked within a single lifetime, arranging income so that one household keeps more of what it saved. The same care extends past the end of a life, to the money that outlasts its owner. Required distributions, Roth conversions, and tax-free death benefits all leave something behind, and how that something passes to the next generation, with or

without a tax bill of its own, is a planning problem in its own right. That is where the planning turns next.

Chapter 15: Generational Wealth Transfer

Every strategy so far has worked toward one goal: making a person's own money last as long as they do. What happens to whatever is left, once it no longer has to last, is a separate question. Passing money to the next generation is its own planning problem, with its own rules, its own taxes, and its own ways to go wrong, and it rewards the same habit good planning always rewards: deciding on purpose rather than by default. The good news is that for most families the hardest part is not the tax. It is making sure the money goes where it is meant to, in a form the people receiving it can use.

Two facts set the shape of everything that follows. The first is that the federal estate tax, the one most people picture when they hear the word inheritance, will never touch the great majority of families; in 2026 it applies only to estates above $15 million for an individual or $30 million for a couple. The tax that actually reaches ordinary inheritances is the income tax, and it falls on a specific kind of money: the traditional retirement account, which arrives at the heir's door fully taxable. The second fact is that the document controlling where much of an estate goes is not the will at all. For a large share of what most people own, a separate form decides, and a will cannot override it.

The Form That Outranks the Will

That form is the beneficiary designation. Retirement accounts, life insurance policies, and accounts marked

payable on death or transfer on death all pass directly to whoever is named on their beneficiary forms, no matter what the will says. A will that leaves everything to a current spouse cannot redirect an IRA whose form still names an ex-spouse from a marriage two decades earlier. This is not a rare mistake; it is one of the most common and most painful in all of estate planning, because it stays invisible until the moment it can no longer be fixed. No hour of planning pays off better than the time spent pulling up every retirement account and insurance policy to confirm, in writing, who is named on each.

Those beneficiary designations, the forms on file with each account and policy, also decide how smoothly the money moves. Assets that pass by beneficiary designation skip probate, the court process that validates a will and settles an estate. Probate can take months, is a matter of public record, and can carry court fees, executor commissions, and in some states a probate tax. Assets that pass only through the will go through probate first. Naming a beneficiary on every retirement account, insurance policy, and payable-on-death account, and a contingent beneficiary in case the first has died, keeps those assets private and quick. The will still matters; it governs everything without its own beneficiary form, names the executor who carries out the estate, and, for parents of minor children, names a guardian. But for the typical retirement saver, the beneficiary forms move more of the money than the will ever will.

What Heirs Take Home

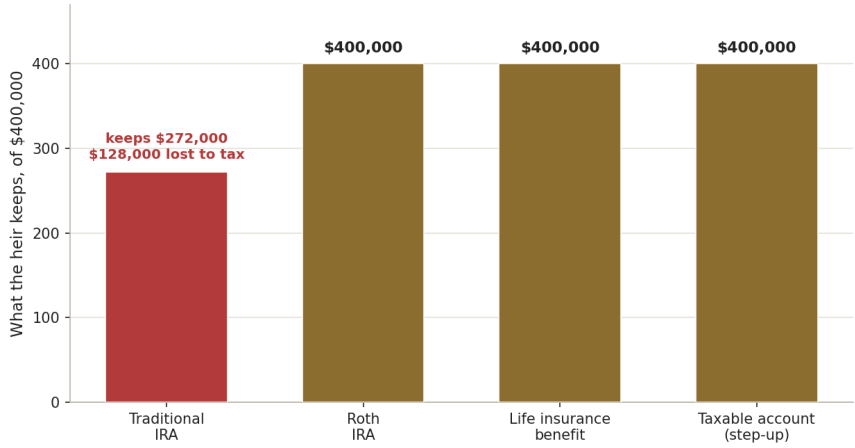
A second idea matters as much as the first: a dollar left to an heir is not simply a dollar, because what the heir owes in tax depends entirely on which account it came from. Two people can leave estates of identical size and hand their children wildly different amounts of spendable money, purely from the mix of accounts the wealth sits in. Sorting an estate by tax character, the same way savings sort into taxable, tax-deferred, and tax-free buckets, is the heart of leaving money well. The table lays out what an heir keeps from each kind of account.

Account left to heirs	What the heir owes in tax	The catch, or the gift
Traditional IRA or 401(k)	Ordinary income tax on every dollar	Must be emptied within ten years, often in the heir's peak earning years
Roth IRA	Nothing, if the account was open five years	Still emptied within ten years, but the growth stays tax-free
Life insurance death benefit	Nothing	Passes income-tax-free, outside probate, to the named beneficiary
Taxable brokerage account	Tax only on growth after the date of death	The step-up erases the gain built up during the owner's life
Cash and bank accounts	Nothing on the principal	Already taxed; it simply passes to the heir

What an heir keeps, by the kind of account left to them. A traditional retirement account is the most heavily taxed inheritance; a Roth, a life insurance benefit, and a stepped-up taxable account are the lightest.

One row of that table is the problem at the center of modern inheritance. The traditional IRA or 401(k), the account most people accumulate the most in, is the worst kind of money to leave, because the heir owes ordinary income tax on every dollar withdrawn, and the law now forces it all out within ten years. An adult child inheriting a large traditional account often

receives it squarely in their own peak earning years, stacking the inherited income on top of a salary already taxed at a high rate. The best kinds of money to leave are the mirror image: a Roth account and a life insurance death benefit both reach the heir entirely free of income tax, and an ordinary taxable account carries a hidden gift of its own. The figure shows the gap.



The same $400,000, left four ways. A traditional retirement account reaches an heir in their peak earning years fully taxable, while a Roth, a life insurance benefit, and a taxable account with a stepped-up basis arrive with little or no tax. The traditional-account figure assumes an heir in a high bracket; all figures are illustrative.

The Ten-Year Rule

The ten-year rule is recent enough that many people's plans were made before it existed and do not account for this major change in how an inherited traditional account is taxed. For decades an heir could stretch withdrawals from an inherited retirement account across their own lifetime, letting it grow tax-deferred for years and pulling out only a little annually. A 2019 law ended that for most heirs. Today a non-spouse beneficiary, which means most adult children, must

empty an inherited retirement account by the end of the tenth year after the owner’s death. The tax that the original owner deferred for a lifetime now comes due inside a single decade. The rule and its exceptions appear in IRS Publication 590-B, Distributions from Individual Retirement Arrangements.

How that decade is handled makes a large difference, and there is a wrinkle the regulators were slow to settle. If the person who died was already taking their own required distributions, the heir must keep taking a required amount in each of those ten years, rather than simply empty the account at the end. Either way, the worst approach is usually to wait and take the whole balance in year ten, which can pile a fortune of taxable income into a single year and push the heir into the highest brackets. Spreading the withdrawals across the ten years, and leaning them toward the heir’s own lower-income years, is what keeps the rate down. The inherited account becomes the multi-year tax problem in miniature, handed to the next generation.

The heir	How long to withdraw	What it means
A surviving spouse	Can treat the account as their own	The most flexible option; withdrawals can wait for years
An eligible designated beneficiary	Over their own lifetime	A minor child of the owner, a disabled or chronically ill heir, or one not more than ten years younger
Most other heirs	Within ten years	The standard rule for adult children and other non-spouse heirs
A successor beneficiary	Only the time left on the first clock	Inherits an already-inherited account, with no fresh ten years

Who must empty an inherited retirement account within ten years, and who can take longer. The category is set by the heir's relationship to the person who died.

Not every heir lives under the ten-year clock, and the exceptions are worth knowing. A surviving spouse has the most freedom, able to treat an inherited account as their own and put off withdrawals for years. A handful of others, called eligible designated beneficiaries, can still stretch withdrawals over a lifetime: a minor child of the owner, until adulthood; someone disabled or chronically ill; and anyone not more than ten years younger than the person who died. Everyone else lives with the ten years. And a successor beneficiary, the person who inherits an account that was already

inherited once, gets only whatever time is left on the original clock, never a fresh ten.

> Linda Castellano is in exactly that position, with two years left on a clock her mother started, which is why those two years press on every other decision she makes.

The Step-Up That Rewards Patience

The hidden gift inside an ordinary taxable account is called the step-up in basis, and it reverses the usual advice about which money to spend first. During a person’s life, an investment sold for more than its purchase price owes tax on the gain. At death, that slate is wiped clean: the heir inherits the investment as though they had bought it at its value on the day the owner died, and the lifetime of gain that built up underneath simply disappears for tax purposes. A stock bought for $40,000 and worth $200,000 at death passes to the heir with a fresh cost basis of $200,000, and the heir can sell it the next day with no tax owed.

Leaving a legacy changes the usual order-of-withdrawals logic. The conventional advice spends the taxable account first, to let the sheltered accounts compound. But an appreciated taxable account held until death hands the heir its gains tax-free, while a traditional account held until death hands the heir a tax bill. For a household that expects to leave money behind, the calculus can flip: it may be better to spend down or convert the traditional account and let the appreciated taxable account ride, precisely so the step-up can do its work. The right answer depends on the

specific accounts and the specific heirs, which is part of why this kind of planning rewards a careful look rather than a rule of thumb.

Turning a Taxable Legacy Into a Tax-Free One

Seen from the heir's side, two familiar tools change character. The Roth conversion, known first as a way to control a person's own tax rate, is also one of the most effective legacy moves available, because it transforms the worst kind of money to inherit into the best. A traditional dollar converted to a Roth during the owner's low-income years, taxed once at a modest rate, passes to the children entirely tax-free, sparing them the ten years of taxable withdrawals they would otherwise face. The owner, in effect, prepays the heirs' tax bill at their own lower rate.

Life insurance does something similar by a different route. The death benefit from a permanent policy arrives to beneficiaries free of income tax, which makes it a way to convert taxable assets into a tax-free inheritance: money is drawn from a traditional account during life, used to fund the policy, and reappears at death as an untaxed benefit.

> Robert and Susan Whitfield, a fictional couple, are using both tools at once, converting traditional balances toward a Roth in their gap years and directing part of their estate into a permanent policy, so that what reaches their three children arrives with as little tax attached as the law allows. For them, with an estate far below the federal threshold, the entire effort is aimed at

income tax on the inheritance, not estate tax, which they will never owe.

The Estate Tax Most Families Will Never Pay

Worry about the estate tax runs well ahead of the facts. In 2026 the federal exemption stands at $15 million for an individual and $30 million for a married couple, a level a 2025 law made permanent and indexed to inflation. Only a tiny fraction of estates, well under one percent, owe any federal estate tax at all. For everyone else, the federal estate tax is simply not the threat; the income tax on inherited retirement accounts is the one that bites, and it bites at every level of wealth.

A pair of cautions rounds out the picture. Some states levy their own estate or inheritance tax with far lower thresholds than the federal one, a few starting as low as one or two million dollars, so where a person lives can matter even when the federal tax does not, which is worth checking against the rules of the specific state. And there is a simple tool for those who want to give during life rather than at death: each year a person can give a set amount, $19,000 in 2026, to any number of recipients without touching the lifetime exemption or filing anything, a simple way to move money to children or grandchildren a little at a time.

The Documents

Underneath the tax questions sits a layer of plain paperwork that does more to protect a family than any elaborate strategy. A will is the foundation, directing whatever does not pass by beneficiary form and

naming the people who will carry out the estate and raise minor children. Many households add a revocable living trust, an arrangement that holds assets during life and passes them at death without probate, keeping the transfer private and avoiding the court delay; it is especially useful for those who own property in more than one state or simply value the privacy. Neither document is exotic, and both are far cheaper than the confusion their absence creates.

Another two documents matter as much for life as for death, because they take effect while a person is still living but unable to act. A durable power of attorney names someone to manage money and sign documents if the owner cannot, and a healthcare power of attorney, paired with an advance directive, names someone to make medical decisions and records the person’s own wishes about care. Without them, a family may have to go to court to gain the authority to help, at the worst possible moment to be filing paperwork. These are not estate planning in the glamorous sense, but they are the documents whose absence causes the most avoidable grief.

When an Heir Needs Protecting

Sometimes the goal is not only to pass money but to protect the person receiving it, and a direct inheritance can do more harm than good. An heir who relies on means-tested public benefits, such as a person with a disability, can be disqualified from those benefits by a lump sum landing in their name. The tool for this is a special-needs trust, which holds the inheritance for the person’s benefit without counting as their own resource, preserving both the money and

the benefits. A companion account, known as an ABLE account, lets a person with a qualifying disability hold a limited amount of savings on the same protected footing. The Whitfields face exactly this question with a grandchild, and for them a special-needs trust is the difference between an inheritance that helps and one that quietly strips away the support the child depends on.

Minor children raise a smaller version of the same issue, since they cannot legally receive a large inheritance outright, and money left to them without a structure can land under court control until they come of age. Naming a trust, or a custodial arrangement, gives a trusted adult the job of managing the money until the child is ready. The same caution applies to naming a trust as the beneficiary of a retirement account, which can be done but interacts with the ten-year rule in technical ways that reward professional drafting. The theme across all of these is that the right structure protects the gift, and the wrong one, or none at all, can undo it.

Putting It Together

The Whitfields' legacy plan pulls these threads into one effort. They keep every beneficiary form current, so the bulk of their accounts will pass without probate to the people they intend. They are converting traditional money to a Roth and funding a permanent policy, so that what their children inherit arrives with little or no income tax. They plan to spend down the traditional account and let an appreciated brokerage account

ride, so the step-up can erase its gains for the heirs. And they have set up a special-needs trust for their grandchild, so that one child's share helps rather than harms. None of it is exotic, and none of it requires an estate large enough to owe estate tax. It is the same habit of deciding on purpose, carried one generation further.

Not every story ends with a legacy, and pretending otherwise helps no one. Helen, the mother of the fictional Mary Thompson, may need years of paid care that her savings were built to fund, and if she does, much of what she might have left could go to that care instead. That is not a planning failure; it is what the money is for. A plan that protects a person's own security first, and treats whatever remains as a bonus to pass along, has its priorities in the right order. The goal is not to maximize an inheritance at the expense of the life that earned it, but to make sure that whatever is meant to pass along does, with as little lost to tax and friction as possible.

Legacy planning is lifetime tax planning carried one step further, past the point where the saver is still alive to steer it. The same levers that lower the tax on a withdrawal, a conversion, or a single year's income also decide how much of what remains reaches the next generation intact. What is left is to bring those levers together into one plan for one household, rather than a drawer full of separate tactics.

Chapter 16: Putting It All Together

Each chapter of this book has handed over a tool. How to time a Social Security claim. How to enroll in Medicare without tripping a lifelong penalty. How to sort savings into buckets, convert from one to another, and draw them down in an order that keeps the lifetime tax bill low. How to judge whether a particular insurance product earns its place. How to pass what is left to the next generation with as little lost to tax as possible. Each tool is useful. None of them is a plan. A drawer full of good tools is not the same as a house, and the work of this final chapter is assembly: fitting the pieces into a single structure that stands.

Four fears opened this book, and they are the right frame for the assembly. The fear of running out of money. The fear of paying more tax than necessary. The fear of losing savings to a bad market at the wrong time. The fear of being blindsided by the true cost of healthcare in retirement. Every tool in the book exists to answer one of those four, and a real plan is what makes the answers work together instead of at cross-purposes. The chapters took the fears one at a time; a household lives all four at once.

From Tools to a Plan

What separates a tactic from a plan is coordination. A Roth conversion sized without a glance at the Medicare thresholds can lower one tax and raise another. A Social Security claim made for the comfort of an early check can shrink a survivor’s income for decades and waste the lowest-tax years a household will ever see. An insurance product bought because it

sounded safe, without a defined role, can lock away money that was needed elsewhere. Each of these is a sound tool used in isolation, and isolation is exactly what turns a tool into a mistake. The plan is the thing that holds them in relation to one another.

A plan also has an order. Not every decision carries equal weight, and some have to be made before others can be judged. A household that knows how its essential bills will be paid can afford to take more risk with the rest of its money; one that has not settled that question cannot sensibly answer any of the others. So the assembly proceeds in layers, from the most load-bearing decision to the most refined, each resting on the one beneath it. The table below sets out the order.

Build the Plan in Layers

These layers move from security to optimization, which is roughly the order of the four fears by consequence. The foundation is reliable income for the essentials. On top of it sits the healthcare decision, with its unforgiving deadlines. Above that, the multi-year tax plan that coordinates the buckets. And at the top, the protection against rare catastrophes and the arrangements for what outlives the household. A plan built in this order is stable; one built in reverse, chasing tax savings before the floor is secure, is a roof with no walls.

The layer	The task	The fear it answers
One: the floor	Cover the essential bills with income that cannot run out, from Social Security, a pension, and if needed a lifetime-income product	Running out; market loss
Two: healthcare	Enroll in Medicare on time, choose the path with an eye on the later lock-in, and plan around the income surcharge	A Medicare surprise
Three: the tax years	Coordinate the three buckets, convert in the gap years, and draw down in a low-rate order	Paying too much tax
Four: the tail and the transfer	Plan for long-term care, and arrange a tax-efficient legacy with current documents	Running out late in life; protecting heirs

The plan in four layers, built from the bottom up. Each layer rests on the one below it, and the layers answer the four fears roughly in order of consequence.

Order matters here because the layers depend on one another in one direction only. The tax plan in the third layer cannot be designed until the income floor in the first is set, because the floor determines how much of the portfolio is free to be converted, spent, or left alone. The transfer plan at the top assumes the

household's own security is already handled, which is why it comes last and not first. Building from the bottom keeps every decision resting on settled ground.

Layer One: A Floor Under the Essentials

The first task is to make certain that the bills which must be paid every month, regardless of what the market does, are covered by income that cannot run out. The bedrock of that floor is Social Security, which an earlier chapter argued is best understood as a government-backed, inflation-adjusted lifetime annuity, and which grows by about 8 percent for each year a healthy person can delay it toward seventy. A pension, where one exists, adds to the floor. For a household whose dependable income still falls short of its essential bills, a portion of savings can be turned into a contractual stream of lifetime income to close the gap. The point of the floor is not to maximize anything; it is to make the essentials immune to a bad market and a long life at once.

Among the cast, the floor looks different from one household to the next. Dave and Mary Thompson, a fictional couple, have it easy: Mary's teacher pension and the couple's Social Security together cover their essential spending, so the rest of their savings is a cushion rather than a lifeline, and they can invest it for growth without fear of a market dip reaching the grocery budget. The Whitfields, also fictional, build their floor mainly by delaying Robert's Social Security to seventy and drawing on a large portfolio. Linda Castellano, a fictional single saver, leans on

delaying her own benefit and is building the floor she will one day stand on. In each case the floor comes first, and everything else is arranged above it.

Layer Two: Settle the Healthcare Question

The second layer is Medicare, and it earns its place this high because its mistakes are both permanent and easy to make. The tasks are concrete: enroll during the seven-month window around the sixty-fifth birthday to avoid penalties that attach for life, choose between Original Medicare with a supplement and a Medicare Advantage plan with the later lock-in fully in mind, and understand that income drives the Medicare surcharge two years after the fact. That last point is the hinge between this layer and the next, because the same income that fills a tax bracket can also trip a Medicare premium, and the two have to be planned together.

Walter Higgins, another of the book's fictional retirees, is the cautionary figure here. He chose a Medicare Advantage plan at sixty-five for its low premium, stayed healthy for years, and later found he could not move to a supplement when his health changed, because the early window had closed. His plan, in the one layer that most rewards getting it right the first time, was made by default rather than on purpose. The lesson is not that he chose wrong, but that he chose without weighing the lock-in, which is the kind of decision this layer exists to make deliberately.

Layer Three: Coordinate the Buckets and the Years

The third layer is the multi-year tax plan, and it is where the most detailed work lives. Its raw material is the three buckets, the taxable, the tax-deferred, and the tax-free, and its method is to move income deliberately across the years so that the lifetime rate, not any single year's rate, is what comes out low. The gap years between the end of work and the start of required distributions are filled on purpose with conversions; the withdrawal order is blended to keep each year inside a sensible bracket and under the Medicare tiers; required distributions are shrunk in advance rather than absorbed at full force. This is the layer that answers the fear of paying too much tax.

Inside this layer sit the insurance products from earlier chapters, each a tool with a defined role rather than a plan unto itself. A fixed indexed annuity can serve as a volatility buffer, a place for money that needs to avoid a market loss without sitting in cash. A multi-year guaranteed annuity can act as an alternative to a certificate of deposit for a fixed holding period. Indexed universal life can provide a tax-advantaged stream of income and a tax-free benefit to heirs, for the right buyer with the right time horizon. The discipline is the same in every case: the product follows the role, and a product without a clear role does not belong in the plan.

> Here too the households diverge. The Whitfields run a full decade-long sequence, converting toward the top of a chosen bracket each year,

watching the surcharge tiers, paying the tax from their brokerage account, and giving from their traditional account once charitable distributions become available to them. Dave and Mary run the same logic in miniature, a few modest conversions at a low rate, with far less to coordinate. The framework does not change from one household to the other; only the size of the dials does. A coordinated tax plan rewards anyone with more than one kind of account, not only the wealthy.

Layer Four: Guard the Tail and Pass It On

The top layer holds the risks that are unlikely in any given year but ruinous if they land, and the arrangements for what remains at the end. The chief tail risk is long-term care, a cost that can run for years and that ordinary Medicare barely touches; planning for it may mean insurance, an earmarked share of savings, or a clear-eyed decision to rely on family and public programs, but it should be a decision and not a surprise. Alongside it sits the transfer: keeping beneficiary forms current, using conversions and life insurance to hand heirs money in its least-taxed form, and putting the basic documents, a will, the powers of attorney, perhaps a trust, in place.

This layer asks a household to plan for the outcomes it would rather not picture. Helen, the mother of the fictional Mary Thompson, may spend much of her savings on the care this layer plans for, leaving less behind than she once imagined. That is not a failure of planning; it is

planning working as intended, with a person's own security ranking ahead of an inheritance. The Whitfields, at the other end, arrange a special-needs trust so that a grandchild's share helps without stripping away support. The tail and the transfer are where a plan admits what it cannot control and arranges the controllable parts around it.

A Plan Is Never Finished

A plan assembled once and filed away will drift out of date, because the things it depends on keep moving. Tax brackets and Medicare thresholds change each year. Markets rise and fall. Health turns. Laws are rewritten, as the recent changes to estate and retirement rules show. A plan is better understood as a habit than a document: a yearly pass through a short list of moving parts, made at the natural moments the calendar already provides. The table below sets out that review.

When	What to check
Each fall, October 15 to December 7	Re-shop the Medicare drug and Advantage coverage during open enrollment, since plans change every year
Before year-end	Size the year's Roth conversion to a bracket and under the next surcharge tier, take any required distribution, and give from the IRA if making charitable gifts
After any large life or income change	Re-check the plan; a marriage, a death, a windfall, or a new law can move several layers at once
Once a year, in any case	Confirm the beneficiary form on every account still names the right people

A short annual review keeps the plan current, made at the moments the calendar already provides. Dates and thresholds should be confirmed each year.

And a plan should know its own limits. The future cannot be forecast, so the aim is a plan that holds up across a range of futures rather than one tuned perfectly to a single guess. The tax tail should not wag the dog: liquidity, peace of mind, and the freedom to live now all count alongside the lifetime tax bill. Where the interactions grow genuinely tangled, and they do, the coordination is the kind of work that often repays a professional who can see the whole board at once. The goal is not a perfect plan, which does not exist, but a sound one that lets a person stop worrying and live.

The Whole Picture, One Household at a Time

Seen whole, the Whitfields' plan is all four layers at once. Their floor is Robert's delayed Social Security atop a large portfolio. Their healthcare layer is Original Medicare with a supplement, chosen with eyes open and managed around the surcharge tiers. Their tax layer is the decade of gap-year conversions, the blended withdrawal order, and the charitable distributions to come. Their top layer is the long-term-care decision, the conversions and policy that will pass money to their children tax-free, and the special-needs trust for their grandchild. No single move in that list is remarkable. Together they are a plan that meets all four fears for one specific family.

Dave and Mary's plan is quieter, and that is its strength. Because their pension and Social Security cover the essentials, the floor does most of the work, and the layers above it can stay modest. Medicare on time, a low-premium drug plan funded from a health savings account, a few small conversions, a long-term-care conversation, and current beneficiary forms: that is nearly the whole of it. They will never face the Whitfields' surcharge tiers or their estate questions, and they do not need to. A good plan is sized to the household, and a simple plan fully carried out beats an elaborate one half-built.

Linda's plan carries the most tension, because her layers pull against one another. She is delaying Social Security to strengthen a floor she has not

yet reached, while a two-year deadline on an inherited account forces high income at single-filer rates, and her age below sixty-five ties any conversion to the cost of her own health coverage. Her plan is less about adding moves than about timing the unavoidable ones so they do not collide, holding her own conversions until the inherited account is empty and the path to Medicare is clear. A constrained plan is still a plan, and naming the constraints is the first step to working within them.

Closing

The four fears that opened the book are real, and none of them is conquered by a single product, a trick, or a lucky year. Each is met, instead, by a plan: dependable income that answers the fear of running out, a coordinated tax strategy that answers the fear of paying too much, a protected floor and a volatility buffer that answer the fear of a bad market, and an early, deliberate command of Medicare that answers the fear of a healthcare surprise. Faced one at a time, the fears are daunting. Faced by a plan that holds them together, they lose their grip.

That is the whole of it. A lifetime of careful saving deserves a retirement arranged with the same care, so that the money becomes income that lasts, the tax is no more than it has to be, the market cannot reach the essentials, and what remains passes cleanly to the people who come next. The particular plan belongs to each household and looks a little different for every one; the principles behind it are shared, and they have

been the subject of this book. The title was a promise and an instruction at once. With a plan in place, it becomes the likeliest outcome of all: that a person does not run out.

Important Notice

This book is educational and informational. It explains how Social Security, Medicare, taxes, and certain insurance products generally work, using simplified examples. It is not financial, tax, legal, investment, or insurance advice for any particular person, and it is not a recommendation to buy or to refrain from buying any product or security.

Every household's situation is different. The right decision for one of the fictional families in these pages may be the wrong decision for a real reader whose income, health, family, and goals differ in ways no book can anticipate. Before acting on anything described here, a reader should consult professionals qualified for the specific question: a tax professional for tax matters, an attorney for legal and estate matters, and a licensed advisor for investment and insurance matters.

The author is licensed to discuss and offer life, health, annuity, and Medicare-related insurance products. The author is not, and does not act in this book as, an attorney, an accountant, or a registered investment adviser, and nothing in this book creates an advisory or client relationship of any kind.

Reading this book creates no ongoing relationship, obligation, or duty between the reader and the author. Any decision a reader makes after reading remains the reader's own.

The figures in this book, including tax brackets, standard deductions, contribution limits, Social Security amounts, Medicare premiums and

surcharges, and estate and gift thresholds, are stated for 2026 or as otherwise noted. These numbers change frequently, often every year, and the laws behind them can change at any time. Any figure should be confirmed against current sources before it is relied on.

The people in this book, including the Whitfields, the Thompsons, Linda Castellano, Walter Higgins, Helen, and every other named individual, are fictional. They were created to illustrate ideas. Any resemblance to real persons, living or deceased, is coincidental.

Insurance and annuity products involve costs, limitations, holding periods, and surrender charges, and any product feature described as guaranteed depends on the claims-paying ability of the issuing insurer. Product illustrations show how a contract might behave under stated assumptions; they are not promises of future results. The specific terms of any product are governed by its own contract and disclosures, which control over any general description in this book.

Charts and illustrations in this book are simplified to make a point. They use rounded, illustrative figures rather than the precise output of any specific product, plan, or tax return.

Glossary

Adjusted gross income (AGI). A person's total income minus certain adjustments, calculated on the federal tax return. It is the starting point for many other figures, including the income measures that drive Medicare surcharges and health subsidies.

Annual gift exclusion. The amount one person can give to another in a year without using any lifetime gift and estate exemption or filing a gift tax return. It was $19,000 per recipient in 2026.

Cap. In a fixed indexed annuity or indexed universal life policy, the highest interest rate that can be credited in a given period, no matter how much the underlying index rises.

Cost of insurance. The internal charge inside a permanent life insurance policy that pays for the death benefit. It rises as the insured ages and is deducted whether the policy's index gains or sits flat.

Delayed retirement credits. The increase, about 8 percent for each year, that a Social Security benefit earns for each year a person waits to claim past full retirement age, up to age seventy.

Eligible designated beneficiary. A category of heir who may stretch withdrawals from an inherited retirement account over a lifetime rather than emptying it within ten years. It includes a surviving spouse, a minor child of the owner, a disabled or chronically ill heir, and an heir not more than ten years younger than the owner.

Fixed indexed annuity. An insurance contract whose interest is tied to a market index, with a floor that

prevents a loss in a down year and a cap or other limit on the gain in an up year. It is often used as a volatility buffer.

Full retirement age. The age at which a person can claim an unreduced Social Security benefit, ranging from sixty-six to sixty-seven depending on the year of birth.

Guaranteed minimum. The floor an annuity contract promises, such as a minimum interest rate or a minimum lifetime income, regardless of how the index performs.

Indexed universal life. A permanent life insurance policy whose cash value earns interest tied to a market index, with a floor and a cap, and whose death benefit passes to heirs free of income tax.

IRMAA. The income-related monthly adjustment amount, a surcharge added to Medicare Part B and Part D premiums for higher-income beneficiaries, based on income from a tax return two years earlier.

Market value adjustment. An increase or decrease applied to a larger-than-allowed annuity withdrawal during the surrender period, reflecting how interest rates have moved since the contract began.

Medicare Advantage. A privately run alternative to Original Medicare that bundles hospital, medical, and usually drug coverage into one plan, often with a lower premium but a fixed network and a later limit on switching to a supplement.

Medigap. A Medicare Supplement policy that pays many of the out-of-pocket costs Original Medicare

leaves behind. Outside an early enrollment window, switching into one can require medical underwriting.

Modified adjusted gross income (MAGI). A person’s adjusted gross income with certain items added back. Different programs define it slightly differently; the version used for Medicare surcharges adds back tax-exempt interest.

Modified endowment contract (MEC). A life insurance policy funded so quickly that the tax code reclassifies it, removing the tax advantages of its withdrawals and loans. Careful funding keeps a policy under this line.

Multi-year guaranteed annuity (MYGA). An annuity that pays a fixed interest rate for a set number of years, often used as an alternative to a certificate of deposit.

Participation rate. In an indexed product, the share of an index’s gain that is credited to the account, such as 80 percent of the index’s rise.

Point-to-point crediting. A method that measures an index’s change from the start of a period to its end, ignoring the movement in between, to decide the interest credited.

Probate. The court process that validates a will and settles an estate. It can take months and is a matter of public record; assets that pass by beneficiary designation skip it.

Qualified charitable distribution (QCD). A gift sent directly from a traditional IRA to a charity, available after age seventy and a half, that counts

toward the required distribution but is left out of taxable income.

Required beginning date. The point at which required minimum distributions must start, generally April 1 of the year after a person reaches the required age.

Required minimum distribution (RMD). The amount the law forces out of a traditional retirement account each year once the owner reaches the required age, currently seventy-three and rising to seventy-five for those born in 1960 or later.

Revocable living trust. An arrangement that holds assets during life and passes them at death without probate, while remaining changeable by the person who created it.

Roth conversion. Moving money from a traditional retirement account to a Roth account, paying income tax on the amount moved in exchange for tax-free growth and withdrawals afterward.

Spread. In an indexed product, an amount subtracted from the index's gain before interest is credited, such as crediting the index's rise minus 2 percent.

Step-up in basis. The reset of an asset's cost to its value on the owner's date of death, which erases the capital gain that built up during the owner's life for the heirs who inherit it.

Successor beneficiary. The person who inherits a retirement account that had already been inherited once. A successor receives only the time left on the original ten-year clock, not a fresh ten years.

Surrender charge. A fee for withdrawing more than the allowed amount from an annuity during its early years, usually declining to zero by the end of the surrender period.

Tax buckets. A way of sorting savings by how they are taxed: taxable accounts, tax-deferred accounts such as a traditional IRA, and tax-free accounts such as a Roth.

Volatility buffer. A pool of money that does not fall when markets fall, drawn on during a downturn so that investments are not sold at a loss, giving the rest of a portfolio time to recover.

Sources and Where to Verify the Numbers

The dollar figures in this book come from the agencies that set them. Those numbers are revised every year, and the laws behind them can change at any time, so the most useful thing this page can do is point to the offices that publish the current figures. Any amount in these pages can then be checked against the year in which it is read. The figures throughout are stated for 2026.

Medicare. The Centers for Medicare and Medicaid Services publish the annual Medicare costs. The yearly Medicare Parts A and B Premiums and Deductibles fact sheet is the source for the Part B premium, the Part A and Part B deductibles, and the income-related surcharge tiers. The Part D figures, including the out-of-pocket cap, the maximum deductible, and the base beneficiary premium, come from the annual Part D redesign program instructions. Plan-level costs and the current-year basics are summarized for the public at Medicare.gov.

Social Security. The Social Security Administration publishes the annual benefit and earnings figures at SSA.gov. The yearly cost-of-living adjustment fact sheet is the source for the earnings-test limits and the taxable maximum, and the retirement pages of its Benefits Planner explain the claiming ages, the delayed-retirement credits, and how the earnings test is applied.

Federal income, capital gains, and estate taxes. The Internal Revenue Service publishes the annual

inflation adjustments as a Revenue Procedure, summarized in a companion news release at IRS.gov. That release is the source for the tax brackets, the standard deduction, the long-term capital gains thresholds, and the estate and gift exemption. A separate annual notice on retirement-plan limits is the source for the contribution amounts.

Further reading. For plain-language explanations that track these releases each year, the Tax Foundation (taxfoundation.org) is a reliable nonpartisan source on the tax figures, and the Medicare Rights Center (medicarerights.org) and the Center for Medicare Advocacy (medicareadvocacy.org) are reliable on Medicare.

Every figure in this book should be confirmed against the current-year release before it is relied on. Where a number in these pages differs from an agency's current figure, the agency's figure governs.

Remarks attributed to named public figures are widely published statements, quoted to illustrate a point.

About the Author

Jeff Batson, RFC®, is a Registered Financial Consultant® and the principal of Goat Planning, a retirement planning agency in Virginia. He has spent his career helping households turn what they saved into retirement income, with particular attention to Social Security timing, Medicare choices, and tax-aware withdrawal planning.

More information on the topics covered in this book, along with a way to reach the author, is available at goatplanning.com.