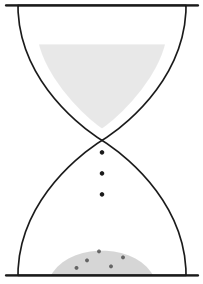
The Hundred-and-Twenty Question

THE REAL SCIENCE OF AGING

from the astonishing to the oversold



Meira Davidson

The Hundred-and-Twenty Question

THE REAL SCIENCE OF AGING

Meira Davidson

The Hundred-and-Twenty Question

THE REAL SCIENCE OF AGING

Copyright © 2026 Meira Davidson

All rights reserved.

No part of this book may be reproduced in any form without permission from the author.

First edition, 2026.

For the Jessica family kids —

be curious, stay open, be kind.

Contents

The Question 6

Chapter 1 · The Counting Problem 16

Chapter 2 · The Clock and the Lever 41

Chapter 3 · The Cells That Won't Die 64

Chapter 4 · The Hunger Switch 88

Chapter 5 · The Beautiful Wrong Idea 111

Chapter 6 · Folding and Falling Apart 132

Chapter 7 · The Reset 155

Chapter 8 · The Reckoning 174

Sources and Further Reading 185

The Question

WHETHER YOU WILL LIVE TO A HUNDRED AND TWENTY — AND *how to tell the science from the slogans.*

You have seen the headlines. A pill that turns back your biological clock. A billionaire pouring a fortune into reversing aging. A supplement that promises to restore your cells to a youthful state. A famous scientist predicting that the first person to live to a hundred and fifty has already been born. And somewhere underneath all of it, a question you may have asked yourself — half hopeful, half skeptical — about whether any of this is real, and whether *you* might live to a hundred and twenty.

This book is about that question. But it is not going to hand you a slogan, in either direction. It will not tell you that you'll live forever if you buy the right supplement, and it will not tell you that aging research is all hype and that nothing works. Both of those are false, and the truth is far more interesting than either. What this book will do is give you something better than an answer: the actual biology of why we age, the actual evidence for what can and cannot be done

about it, and — most of all — the *judgment* to tell the astonishing from the oversold, for the rest of your life. By the last page you will be able to answer the hundred-and-twenty question for yourself, which is worth a great deal more than being told.

From fate to a problem

To see why the question is even worth asking, you have to appreciate how recently it became one. For almost all of human history, aging was simply *fate* — the unalterable winding-down of the body, no more a subject for intervention than the weather or the tides. It looked like pure entropy, the slow rusting of a complicated machine. There were no genes "for" aging, the thinking went, any more than there are genes for a car rusting; you might slow the process a little by living carefully, but the process itself was beyond reach.

Then, only a few decades ago, that changed — and it changed because of a worm. In a discovery you'll meet properly in Chapter 4, a single genetic change was shown to *double* the lifespan of a laboratory roundworm, and to keep it youthful and vigorous the whole way. Aging, it turned out, was not unalterable fate but a *regulated process* — something with molecular switches and levers, something that could in principle be adjusted. That one result, and the thousands that followed it, converted aging from a fact of life into a *scientific problem* — and a scientific problem can be worked on. This

book is a report from that revolution. It is also, not coincidentally, the source of all the noise: the moment aging became something we might be able to change, an enormous industry sprang up to promise that we already can.

Two words the headlines blur

Before we go further, two distinctions that the headlines constantly blur together, and that you will need to keep straight to think clearly about any of this — including the hundred-and-twenty question itself.

The first is the difference between **life expectancy** and **maximum lifespan.** Life expectancy is the *average* — how long a typical person can expect to live. Maximum lifespan is the *ceiling* — how long the longest-lived human being makes it. These are completely different quantities, they respond to completely different things, and they have very different histories. Over the past century and a half, life *expectancy* in wealthy countries roughly doubled — but almost all of that gain came from preventing *early* death: clean water, vaccines, antibiotics, safer childbirth. The *ceiling* has barely moved. The longest verified human life, a French woman who reached 122, ended in 1997, and no one has clearly beaten it in the nearly thirty years since. So whenever someone says "we're going to live to 120," you must immediately ask *which number they mean* — the average creeping upward, or the ceil-

ing being broken — because the evidence for those two things could hardly be more different.

The second distinction is between **lifespan** and **healthspan.** Lifespan is how long you live. Healthspan is how long you live *in good health* — before the years of frailty and disease at the end. They are not the same, and which one you care about changes everything. A treatment that adds frail, sick years is a very different thing from one that adds healthy ones — and, as you'll see, the most defensible goal of aging science may not be to make us live *longer* at all, but to keep us *healthy* longer, compressing the sick years at the end into a shorter span. Hold both distinctions as you read; the whole book, and especially its final reckoning, turns on them.

The mouse and the human

Now the single most important idea in this book — the thread to carry through every chapter, the question to put to every claim.

Aging research is full of genuinely astonishing results: lifespans doubled, organs regrown, biological clocks wound backward, the visible reversal of aging. And almost all of those results are in *animals* — in worms, in flies, and above all in mice. Here is the hard truth that organizes everything that follows: **some of those astonishing animal results survive the jump to a human being, and many of them quietly do not.** A mouse is not a person. Again and again in these

pages you will watch a beautiful, well-supported result in mice fail to transfer to humans — sometimes after decades and billions of dollars — and fail for a specific, instructive reason: the mouse differs from us in exactly the way that made the experiment work in the first place.

So the question to carry, the spine of this entire book, is simple: *for any claim about slowing or reversing aging, does it actually survive the jump from the mouse to the human?* Most of the distance between the headlines and reality lives inside that one gap.

The marker and the lever

A close cousin of that question will recur so often it's worth planting now. Aging science is full of *markers* — things you can measure that change with age: the length of the caps on your chromosomes, the reading of a chemical "clock" in your cells, the amount of a particular protein in your brain. It is endlessly tempting to assume that if you can *move* one of these markers, you have *changed aging itself.* But a marker that *predicts* aging is not the same as a *lever* that *controls* it — and one of the most important and most consistently ignored facts in this whole field is that you can very often move the marker without changing the thing it was measuring. "We can measure it" is not "you should act on it." A thermometer is not a thermostat. You will meet this trap in its mildest form in Chapter 1 and its sharpest form in Chapter 2, and it

will keep returning, because it is the single most common way that people — including very intelligent, very well-funded people — fool themselves about aging.

This is what I mean by giving you *judgment* rather than an answer. The book refuses the hype in both directions. It gives full, honest credit to what genuinely works and what is genuinely exciting — and there is real science here, some of it breathtaking. And it names, just as honestly, what is unproven, what is oversold, and what is flatly wrong — and there is a great deal of that too, often dressed in the costume of science and backed by enormous commercial interest. The discipline is to tell the two apart, and the discipline is teachable. By the end, you'll have it.

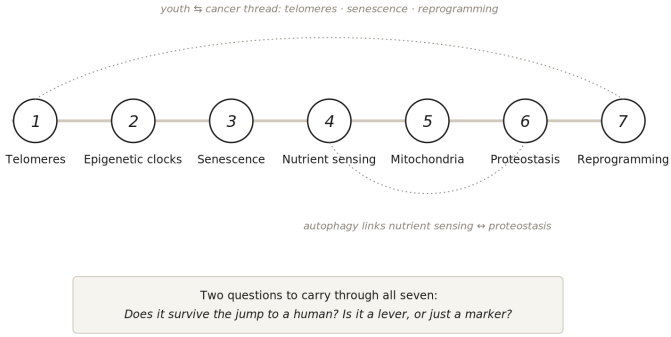
Seven hallmarks, done properly

To organize the biology, this book uses a framework the field itself uses: the **hallmarks of aging.** The core idea is that aging is not one single thing but a collection of distinct, identifiable, *interacting* kinds of damage and decline — a dozen or so fundamental processes that, together, add up to growing old. Rather than survey all of them shallowly, this book takes seven of the most important and most illuminating and does each one *properly* — because understanding seven mechanisms deeply will serve you far better than memorizing twelve names. Each chapter takes one hallmark, tells the story of how we discovered it, builds the actual biology at real depth,

examines what we've tried to do about it, and — without fail — names exactly what we do and do not know.

The seven, in order: the shortening of the protective caps on our chromosomes, and the discovery that a cell counts its own divisions (Chapter 1, *telomeres*); the drift of the chemical marks that switch our genes on and off, and the machine-learning "clocks" built to read it (Chapter 2, *epigenetic alterations*); the cells that stop dividing but refuse to die and poison the tissue around them (Chapter 3, *cellular senescence*); the ancient circuits that sense how much food is around and flip cells between growth and repair (Chapter 4, *nutrient sensing*) — the most reliable longevity lever we have, and the most oversold; the most famous *wrong* theory of aging, that we simply rust from the inside, and the expensive way we learned it was wrong (Chapter 5, *mitochondria*); the proteins that must fold into exact shapes to work, and the diseases that follow when folding and cleanup fail (Chapter 6, *proteostasis*); and finally the most radical idea of all — not slowing the damage but *erasing* it, resetting an old cell to a young one (Chapter 7, *reprogramming*), now entering its very first human trial.

Figure 0.1 — A map of the book



Where the machines fit

One more thread runs through all seven chapters, because this book is part of a series about how *artificial intelligence* is reshaping the search for medicines. Aging is one of the places where AI has made its most spectacular marks. The "clocks" of Chapter 2 are themselves machine-learning models. The prediction of protein structure in Chapter 6 was a fifty-year grand challenge that an AI essentially solved, winning a Nobel Prize. And the search for the molecules that might safely reset a cell's age, in Chapter 7, is exactly the kind of problem the newest AI tools are built to attack.

So as we go, we'll keep a second question running alongside the first: *where can artificial intelligence genuinely help here, and where can't it?* You'll find the honest answer is neither "everywhere" nor "nowhere." AI is extraordinary at certain pieces of this puzzle — finding patterns in vast data,

predicting structures, searching enormous spaces of molecules — and powerless over others: it cannot, by itself, tell you whether something will work in a living human being, and it cannot run the years-long trials that are the only thing that ever settles that question. A dazzling tool that solves one cleanly posed problem is not the same as a solution to the messy one you actually care about — and learning to see that boundary is part of the judgment this book is trying to hand you.

How to read this book

A last word about who this is for, and what it asks of you. It is written for anyone who wants to genuinely *understand* this: the curious teenager who might one day work in a lab like the ones described in these pages, the adult who wants to make sense of the headlines and the choices, the reader who needs to evaluate claims rather than just absorb conclusions. It does not assume you already know the biology — but it does not condescend, either. Where there is real mechanism, you'll get the real mechanism. Where an equation genuinely illuminates, the equation is right here on the page and explained in plain sight. Where the science is uncertain or contested, you'll be told so directly, because the contested places are exactly where you most need the truth and are most often denied it. The book is meant to reward you now, and to reward you again on a second reading years from now.

So carry the two questions — *does it survive the jump to a human?* and *is this a real lever, or just a marker someone moved?* — and we'll return, at the very end, to the one you came in with: whether you can expect to live to a hundred and twenty. By then, you'll be able to answer it yourself.

Let's begin where the field began — at the ends of your chromosomes, with the discovery that a cell can count.

Chapter 1 · The Counting Problem

TELOMERES, THE HAYFLICK LIMIT, AND THE FIRST THING *almost everyone gets wrong about aging.*

> **Companion notebook · Chapter 1.** *A Real Number That Can't Tell Your Age* — open it in Google Colab and build it yourself: Open in Colab →

A CELL THAT WOULD NOT DIE

In January 1912, a surgeon named Alexis Carrel cut a sliver of tissue from the heart of a chick embryo, dropped it into a flask of nutrient broth, and started a clock he believed would never stop. The cells lived. They divided; he fed them; he split the culture; they divided again. Carrel was already famous — he had worked out how to sew blood vessels back together end to end, the technique that made organ transplantation thinkable, and the same year he started his chicken-heart culture he won the Nobel Prize for it. When such a man announced that he had captured a piece of immortality

in a bottle, the world listened. His laboratory tended the culture for thirty-four years, outliving the chicken it came from by three decades. Every January, newspapers reported its birthday.

The lesson the public absorbed, and that biologists absorbed with them, was simple and wrong: a cell, given clean food and a clean home, is intrinsically immortal. If your cells die, the reasoning went, it is because the body wore out around them — the cells themselves could go on forever. For nearly fifty years this was the textbook position.

It was also almost certainly an artifact. The broth Carrel used to feed the culture was an extract made from fresh chick embryos, and there is good reason to think it quietly reintroduced living cells at every feeding. The "immortal" line was not living forever; it was being topped up. The deeper lesson — worth keeping for the whole of this book — is that a charismatic result announced by an authority can hold an entire field captive for decades. Authority is not evidence. The experiment is.

Then, in 1961, Leonard Hayflick counted.

Hayflick and his colleague Paul Moorhead grew normal human cells — fibroblasts, the workhorse cells of connective tissue — and did the unglamorous thing of keeping track of how many times they divided. The cells did not divide forever. After roughly forty to sixty doublings they stopped — not dying, exactly, but settling into a state where they sat alive and metabolizing yet would never divide again. And

Hayflick ran the controls that ruled out neglect. He mixed young cells and old cells in the same dish; the old ones stopped on schedule regardless of the company they kept. It was not the broth, not the technique, not the flask. The limit was coming from inside the cells.

The paper was initially rejected, in part because its conclusions conflicted with the long-standing view Carrel had established. He pushed it through anyway. The phenomenon he had found is now called the **Hayflick limit**: a normal human cell can divide only a fixed number of times — around fifty for these fibroblasts — and then it retires. Somewhere inside it, a cell keeps a tally of how many times it has divided, and when the tally runs out, it stops.

Hayflick had found the counter. He could not yet say what it was made of. That question — what physical thing inside a cell could possibly keep count — is where the molecular story begins, and it leads straight to the ends of your chromosomes.

Why the ends are a problem

To see what the counter is, we have to look at how a cell copies its DNA, and at one specific place where the copying machine fails.

Start with the molecule. DNA is a double helix of two strands running in opposite directions — *antiparallel* — and by convention we label the ends of each strand 5′ and 3′.

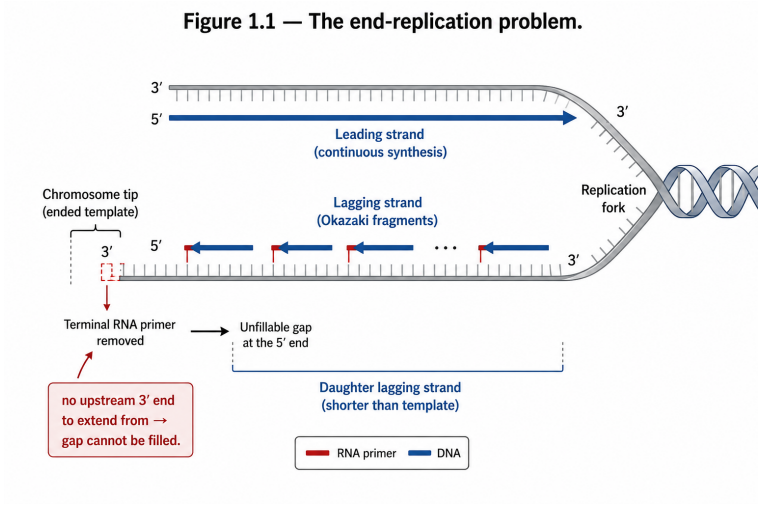
When a cell divides, it copies its DNA by *semiconservative replication*: the two strands separate, and each serves as the template for a new partner. The copying is done by an enzyme, DNA polymerase, which has two stubborn constraints. First, it can only add new building blocks onto an existing 3′ end — it cannot start a strand from nothing, so every new strand has to be kicked off by a short *primer* made of RNA. Second, it can only build in one direction, 5′ to 3′.

Those two constraints collide at a replication fork. Because the strands are antiparallel, one new strand can be built smoothly and continuously in the same direction the fork is opening — the **leading strand**. The other has to be built backwards, away from the fork, in short pieces — **Okazaki fragments** — each one started by its own little RNA primer, which is later removed, filled in with DNA, and stitched to its neighbor. For a chromosome's interior this works fine: every gap left by a removed primer is filled in from the fragment upstream of it.

But look at the very end of a linear chromosome. On the lagging strand, the last Okazaki fragment is started by an RNA primer sitting at the extreme 5′ tip of the new strand. When that terminal primer is removed, there is no fragment further out to fill the gap from — nothing provides the 3′ end the polymerase needs to extend. The gap simply cannot be filled. So every time a linear chromosome is copied, the new lagging-strand copy comes out a little shorter at its 5′ end than the template it was read from. **A linear DNA molecule**

cannot fully copy its own ends. This is the *end-replication problem*, recognized independently by Alexey Olovnikov in Moscow in the early 1970s and by James Watson in 1972.

Figure 1.1 — The end-replication problem.



How fast does the end erode? Measurements put the loss at roughly fifty to a hundred base pairs per division in human cells — it varies with cell type and method; call the per-division loss $\delta \approx 50\text{–}100$ bp. Human chromosome ends carry a buffer of repetitive DNA roughly eight to fifteen thousand base pairs long at birth — though it varies considerably from person to person and tissue to tissue. So the back-of-the-envelope number of divisions before that buffer is gone is

$$\frac{L_0}{\delta} \approx \frac{10{,}000 \text{ bp}}{75 \text{ bp/division}} \approx 130 \text{ divisions}.$$

This is only a first approximation — real telomere loss is neither perfectly constant nor the whole story. And the Hayflick limit, in any case, is about fifty, not a hundred and thirty. The discrepancy is itself a clue: a cell does not wait for the buffer to hit zero. It retires when the buffer reaches a *critical* short length — a few thousand base pairs — that it senses as trouble. Model it simply. Let the buffer length after n divisions be

$$L(n) = L_0 - \delta n,$$

and let L_c be the critical length that triggers arrest. The cell retires after

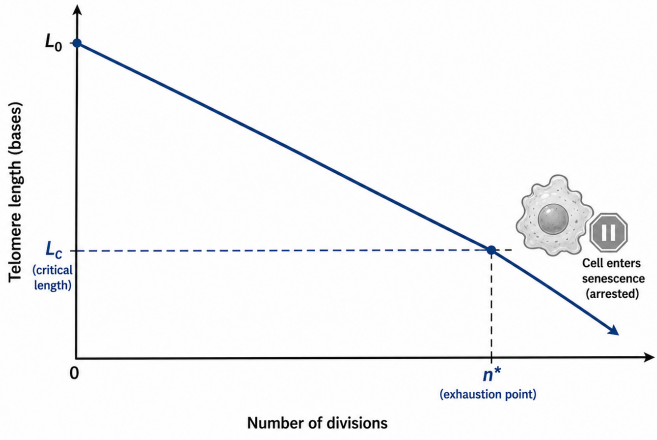
$$n^* = \frac{L_0 - L_c}{\delta}.$$

Put in $L_0 \approx 10{,}000$, $L_c \approx 5{,}000$, $\delta \approx 75$, and you get $n^* \approx 67$ — the right neighborhood for the Hayflick limit. The real dynamics are noisier and not perfectly linear; the evidence suggests it is one or a few critically short telomeres, rather than the average length, that sound the alarm, and there is cell-to-cell scatter. But the linear model captures the essential logic, which is all we need here: **a fixed buffer, eaten at a roughly fixed rate, runs out after a predictable number of divisions.** That is the counter Hayflick was looking for. The tally is written in DNA, at the ends

of the chromosomes, and it is counted down by the very act of copying.

Which raises an obvious worry. If the ends erode every time, why aren't we losing genes with every cell division? Because the eroding ends are not genes. The last stretch of each human chromosome is a **telomere**: the six-letter sequence TTAGGG, repeated over and over — thousands of times — carrying no protein-coding information at all. It is a disposable buffer. And "disposable buffer" here is not a figure of speech dressed up as biology; it is the literal arrangement. The cell feeds meaningless repeats to the end-replication problem so that meaningful sequence sits safely upstream, untouched. Evolution did not solve the problem of un-copyable ends by fixing the copying. It solved it by putting something expendable at the end, to be lost on the cell's behalf.

Figure 1.2 — The counter running down.



Deeper Water — the lagging strand, the overhang, and the cap. The end-replication gap is only half the reason chromosome ends shorten. After replication, nucleases chew the 5′ end back further to leave a single-stranded 3′ overhang — the very overhang the maintenance enzyme we are about to meet needs in order to work. There is also a subtler danger. A chromosome end looks, chemically, exactly like a double-strand break — the most lethal kind of DNA damage there is. If a cell treated its forty-six natural ends as forty-six breaks, it would try to "repair" them by fusing them together, scrambling the genome. It does not, because each telomere is wrapped in a six-protein complex called **shelterin**, which folds the 3′ overhang back and tucks it into the double-stranded telomere to form a lariat — the **T-loop** (worked out largely by Titia de Lange). The T-loop hides the end. So a telomere does two jobs at once: it is the buffer that absorbs the end-replication loss, and it is the disguise that stops the cell from reading its own chromosome ends as catastrophe. When a telomere grows too short to form the loop, the disguise fails, the end is read as a break, the DNA-damage response fires, and

the cell shuts down. That, in molecules, is what "the counter ran out" actually means.

The pond, the yeast, and the Christmas gel

We have been talking about telomeres as if their existence were obvious. It was not. The story of how they were found is worth telling properly, partly because it is a good story and partly because of where the answers came from — none of them from anyone trying to cure aging.

It begins with two people staring at chromosomes through microscopes in the 1930s and 40s. Barbara McClintock, working on maize, noticed that when a chromosome broke, the broken ends were *sticky* — they fused readily with other broken ends. But the natural ends of an intact chromosome never did. So the natural ends had to carry something — some protective identity — that ordinary broken DNA lacked. Hermann Müller, working on fruit flies, saw the same thing and gave the structure its name in 1938, from the Greek *telos* (end) and *meros* (part): the telomere. Both were describing a *function* — end protection — decades before anyone knew what molecule performed it. (McClintock's deeper work, on genes that jump around the genome, was so far ahead of its time that it was largely ignored; she finally won the Nobel for it in 1983, in her eighties.)

To go from "ends are special" to "ends are made of *this*," you need ends — lots of them — and most organisms keep

only a few dozen chromosomes per cell. Elizabeth Blackburn made a shrewd choice of material: a pond ciliate called *Tetrahymena*, which shreds its DNA into thousands of tiny chromosomes, and therefore thousands of telomere ends, per cell. In 1978 she sequenced those ends and found a simple repeat — TTGGGG — written over and over. A pattern at last. (Notice how close it is to the human TTAGGG: similar, not identical, a first hint that this is something ancient and shared across wildly different forms of life.)

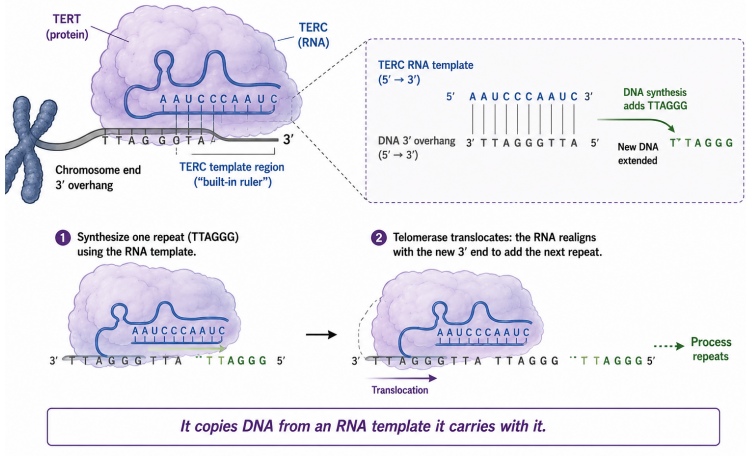
Then came the decisive experiment, in 1982, when Blackburn teamed up with Jack Szostak. Szostak had been working with linear DNA plasmids in yeast, and the yeast kept destroying them — chewing them up from the ends. Blackburn handed him *Tetrahymena* telomere sequences; they stuck those sequences onto the ends of the plasmids; and the yeast stopped destroying them. Telomere repeats from a pond organism protected a piece of DNA inside an entirely different kingdom of life. That is the moment telomeres stopped being a quirk of one ciliate and became a *universal* functional element — the protection was portable across more than a billion years of evolution.

But the same experiment posed a new puzzle. In the yeast, the *Tetrahymena* ends not only survived; they had *yeast*-type repeats added onto them. Something was actively lengthening the ends — adding sequence to fight the erosion. There had to be an enzyme. Blackburn and her graduate student Carol Greider set out to find it.

Greider built an assay: a telomere-like primer, radioactive building blocks, an extract of cell innards, and a gel to separate whatever products came out. On Christmas morning in 1984 she developed the film and saw it — a regular ladder of bands, spaced six bases apart, exactly the signature of an enzyme adding telomere repeats one unit at a time. They had found **telomerase**. Greider was twenty-three.

The enzyme turned out to be stranger and more elegant than anyone expected. How does it know which sequence to add? It carries the answer inside itself. Telomerase is a *ribonucleoprotein* — a protein bound to a short RNA molecule — and that RNA contains a template matching the telomere repeat. The protein half is a *reverse transcriptase*: it reads its own internal RNA template and writes DNA from it, then ratchets forward and writes the same repeat again. The machine that maintains the ends of your chromosomes carries its own little ruler; the sequence to copy is built into the enzyme. (In the modern shorthand: TERT is the protein, TERC is the RNA.)

Figure 1.3 — Telomerase, the enzyme with a built-in ruler.



In 2009 the Nobel Prize in Physiology or Medicine went to Blackburn, Greider, and Szostak. Hold the shape of that story for a moment: a microscopic pond organism, a yeast plasmid, and a graduate student's Christmas gel cracked a problem sitting at the very center of human aging and cancer. It is worth remembering the next time someone asks what possible use there could be in studying pond water.

THE BRAKE AND THE PRICE

Now we can close Hayflick's loop — and then immediately open a much sharper question.

Telomerase exists, and it can keep telomeres long indefinitely. Your germline cells — the lineage that makes eggs and sperm — keep it switched on, which is why your children do not inherit your worn-down telomeres; the counter is reset

each generation. Many stem cells run it at a lower level too. But the great majority of your ordinary body cells switch telomerase *off* after development. With the enzyme gone, their telomeres shorten with every division, and after roughly fifty rounds a critically short telomere trips the DNA-damage response and the cell enters senescence — alive, still working, but permanently retired from dividing. **In normal human fibroblasts, the Hayflick limit largely reflects telomere exhaustion, sensed as damage.** Mystery solved.

But here is the question the rest of this book turns on. Telomerase is the *cure* for telomere shortening. Evolution invented it, kept it, and uses it in the germline. So why switch it off in the body at all? Why deliberately build in an aging clock you already know how to stop?

The answer is one of the most important ideas in all of aging biology. The Hayflick limit is a **tumor suppressor**. A cancer is, at bottom, a cell that divides when it shouldn't, over and over, without limit. A cell that can divide only about fifty times has a built-in ceiling on how large a tumor it can build before it retires. The division counter is a *brake on cancer*. Turning telomerase off in your body cells is the price of that brake: your cells age, in part, precisely so that they cannot easily become immortal.

And we can confirm that this trade-off is real, not just a tidy story — because of what happens when cells *do* become cancer. The overwhelming majority of them switch telomerase back on. Roughly 85 to 90 percent of human can-

cers reactivate telomerase to escape the Hayflick limit and divide without end; most of the rest use a second, recombination-based route called ALT (alternative lengthening of telomeres). Defeating the telomere clock — immortalization — is one of the standard steps a normal cell takes on the road to malignancy.

So telomerase wears two faces at once. The same enzyme that would keep your cells young is the one your tumors must switch on to grow without limit. Aging and cancer are not opposites here; they are two ends of a single trade-off, and the telomere is the fulcrum. This is the first big idea to carry forward, and it generalizes far beyond telomeres: in aging biology, the obvious intervention is almost never free. "Just turn the clock back" usually means "remove a brake that was put there for a reason." Hold that thought. It returns in every chapter.

> **Deeper Water — two barriers, and how cancer gets past both.** Cells face not one telomere barrier to immortalization but two. The first is replicative senescence (sometimes labeled M1), triggered by short telomeres through the p53 and Rb pathways. A cell can be forced past it — for example by viral oncoproteins that disable p53 and Rb — but then it keeps dividing with telomeres that get dangerously short, and slams into the second barrier, *crisis* (M2): widespread

end-to-end chromosome fusions, mangled cell divisions, and mass death. The rare cell that survives crisis has almost always done one specific thing — reactivated telomerase, or switched on ALT, to stabilize its ends. This is why telomerase reactivation is treated as a near-universal hallmark of cancer, and why telomere biology cuts both ways in medicine: too little maintenance and tissues fail early; too much and you have removed a barrier to malignancy. Note the therapeutic inversion this implies. The commercially hyped idea is to *activate* telomerase to fight aging. But the better-validated drug so far runs the opposite way — **imetelstat**, a telomerase *inhibitor*, has been approved for certain bone-marrow disorders, turning the enzyme off to starve abnormal cells. Switching telomerase off as a cancer therapy has reached the clinic. Switching it on as an anti-aging therapy has not.

THE FIRST THING ALMOST EVERYONE GETS WRONG

Here the science meets the marketplace, and a careful reader has to slow down — because what follows is the single most transferable skill in this book. The same trap waits in every chapter to come.

The facts so far are seductive. Telomeres shorten as we age. People with shorter telomeres have, on average, higher rates of age-related disease and earlier death. String those two together and you arrive at a product: a blood test that measures your telomere length and reports your "true biological age," usually paired with a supplement that promises to lengthen it. The pitch writes itself. It is also, at the level of the individual person buying it, close to useless. Understanding *why* is the point.

First, the number is noisy. The cheap and common way to measure telomere length is a method called quantitative PCR, which returns a ratio rather than an absolute length. It is fast and inexpensive, and it carries a large measurement error — split a single blood sample in two and measure both halves, and the two answers can disagree by more than a year's worth of supposed "telomere age." Better methods exist (the Southern blot of terminal restriction fragments; Flow-FISH), but they are slower, costlier, and still not pinpoint. So part of the biological age the test hands you is not your biology at all. It is the scatter of the assay.

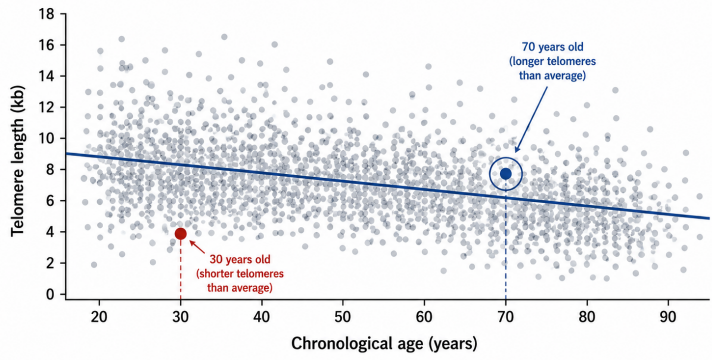
Second — and this is the deep one — a truth about a population is not a truth about a person. Across tens of thousands of people, average telomere length declines with age. That correlation is real and sturdy. But the *spread at any given age is enormous*. Among people who are all seventy, telomere lengths vary so widely that plenty of seventy-year-olds have longer telomeres than plenty of thirty-year-olds.

Put a number on it. The correlation between age and telomere length in blood cells is real but modest — reported values of the correlation coefficient r often fall in the range of roughly -0.3 to -0.5, varying with the age range, the tissue, the assay, and the cohort. Square it:

$$r^2 \approx 0.1 \text{ to } 0.25.$$

That means telomere length accounts for only about ten to twenty-five percent of the variation in chronological age — and even less of the variation in the things you actually care about, like your own risk of heart disease, once age itself is taken into account. A variable can be unmistakably real in a population and still tell you almost nothing about the specific person in the chair. That gap — between a sturdy population average and a weak individual prediction — is the most important and most consistently ignored fact in all of consumer "biological age" testing.

Figure 1.4 — Why the test fails for you.



the trend is real; the spread is enormous.
One person's dot tells you very little.

Third, it is usually the wrong tissue, and a moving target. These tests measure telomere length in your white blood cells, because blood is what you can easily draw. But telomere behavior in immune cells need not match that in your heart muscle, your neurons, or your liver. And immune-cell telomere length is itself a moving target: a serious infection, chronic stress, or inflammation can expand and contract immune-cell populations and shift the reading — so the same person can test differently month to month for reasons that have nothing to do with how fast they are aging overall.

Fourth, cause or symptom? Even the link between short telomeres and disease is not cleanly causal. In many cases short telomeres may be a *record* of accumulated stress and inflammation rather than a *driver* of aging — a smoke detector going off, not the fire itself. Genetic studies that try to separate the two (Mendelian randomization, which uses in-

herited telomere-length variants as a kind of natural experiment) return a humbling result: genetically longer telomeres do lower the risk of some diseases — but they *raise* the risk of several cancers (melanoma, glioma, and certain lung cancers among them). Which is exactly what the brake-and-price trade-off predicted. Longer telomeres are not simply "better." They are a different point on a trade-off — protective against some failures, permissive of others.

So what is telomere length, really? A genuine and important piece of aging biology, and a useful variable for *research* — for asking population questions across thousands of people. As a *personal* diagnostic you can buy and act on, it fails on four counts at once: the measurement is noisy, the individual prediction is weak, the tissue may be the wrong one, and "longer" is not unambiguously "better." Two sentences are worth memorizing, because they are the discipline the rest of this book runs on:

"We can measure it" is not the same as "you should act on it." And "it correlates with aging across a population" is not the same as "it tells you how you *are aging."*

You will need both most exactly where the marketing is loudest.

One nuance is worth holding here, because it sounds like a contradiction and isn't. Longevity clearly *runs in families*. The children and siblings of people who reach a hundred tend to live longer than average, and telomere length itself is substantially heritable. Both are true. But "heritable across

families" and "predictable for one individual from a single measurement" are different claims — the first can be strong while the second stays weak. Your family history is, in fact, one of the better long-range hints you have about your own aging. A one-time telomere readout from a mail-in kit is one of the worse ones. The biology is real; the test is still mostly noise.

When the clock really does break

It would be easy to read that last section as "telomeres don't really matter in people." That would be the opposite error, and the book owes you the other half.

We have a natural experiment that proves telomeres matter in humans: people born with mutations that cripple telomere maintenance — in the genes for telomerase or its partners. These telomere biology disorders (the best-known is dyskeratosis congenita) produce exactly what the theory predicts, and they produce it first in the tissues that divide the most. Bone marrow, which must manufacture blood cells around the clock, fails. The lining of the lungs scars (pulmonary fibrosis). The gut and skin suffer. Affected people show features of premature aging, concentrated in the highest-turnover tissues, sometimes decades early.

So in humans, telomeres are not a decorative correlate. When maintenance genuinely fails, people age prematurely and predictably — fastest where cell division is fastest. The

truth sits in the uncomfortable middle. Telomeres are a real, sometimes rate-limiting part of human aging in specific tissues and specific people, *and at the same time* a weak, noisy yardstick for the gradual aging of a healthy person. Both halves are true. A reader who knows only one half will misread the news for the rest of their life.

What the mouse won't tell you

There is one more result you will see cited — often as proof that aging is basically solved — and it is worth walking through slowly, because it teaches the lesson this entire book is built around.

In 2011, a group reactivated telomerase in mice and reversed signs of aging. Shrunken brains regrew, failing organs recovered, the animals visibly rejuvenated. It is a real, dramatic result. It is also widely misread, and the misreading is instructive.

Two things the headline leaves out. **First, those mice were not normal.** They had been genetically engineered to *lack* telomerase entirely, so that their telomeres collapsed and they aged prematurely; the experiment switched telomerase back on and rescued them. That is not "making a healthy animal younger." It is curing an artificial deficiency the experimenters built in the first place. Restoring something you deliberately removed is a much weaker claim than improving a

normal animal — and only the second would tell you anything about yourself.

Second, and more fundamental: the ordinary laboratory mouse is one of the worst possible models for human telomere aging — because in a mouse, telomeres are not the limiting factor at all. Lab mice carry enormous telomeres, tens to over a hundred thousand base pairs against roughly ten thousand in humans, and they keep telomerase switched on in many adult tissues. By the logic of the Hayflick limit, a mouse ought to be nearly ageless. And yet a mouse grows old and dies in two or three years. Telomere shortening simply is not what paces aging in a normal mouse; other mechanisms get there long before the telomeres do. So if you want to test the hypothesis "telomere shortening drives aging, and restoring telomeres reverses it," the standard mouse is precisely the animal in which telomeres matter least. The 2011 result works *because* the mice were re-engineered to make telomeres limiting. It does not transfer to a normal mouse — let alone to a normal human.

This is the pattern to carry into every chapter that follows. Aging biology is full of astonishing demonstrations in worms, flies, and mice: lifespans doubled, organs regrown, clocks wound back. *Some* of those effects survive the jump to humans. Many quietly do not — because the model organism differs from us in exactly the way that made the experiment work in the first place. A reader who sees "reverses aging in mice" and hears "so it will work in me" has made

the single most common mistake in this field. Telomeres are the right place to learn that lesson, because here the gap between the mouse and the human is not subtle. It is the whole story.

Where the machines come in

This is a book about how artificial intelligence is reshaping the search for medicines, so it is fair to ask where AI touches the telomere story. So far, mostly in measurement and scale. Telomere length can now be estimated computationally straight from genome-sequencing data, and machine-learning methods increasingly help analyze the noisy telomere assays and microscopy images researchers rely on. At the population end, the UK Biobank measured telomere length in the blood cells of roughly half a million people — which turned telomere biology into a large-scale machine-learning and genetics problem, letting researchers map the inherited variants that set telomere length and connect them, disease by disease, to risk.

But here is the honest framing, and it is the bridge to the next chapter. A better algorithm can sharpen the *number*. It cannot make the number mean more than it means. Run the most sophisticated model in the world on one person's blood-cell telomere length, and you still have a noisy, weak, single-tissue predictor of how that person is aging. A precisely measured weak signal is still a weak signal. AI's real

leverage in aging will turn out to live mostly in finding *better signals* and *better drug targets* — not in squeezing certainty out of a measurement that never had much to give.

Which raises the obvious question. If telomere length is a noisy, low-resolution clock, is there a better one — a measurement that tracks biological age tightly enough to be worth acting on? There is a candidate, and it is built entirely out of machine learning. It is called the epigenetic clock, and it is where we go next — along with a trap that looks a great deal like the one in this chapter, only sharper.

✦ ✦ ✦

Carry these three things

1. **A cell counts its divisions, and the counter is a disposable buffer at the end of each chromosome.** Telomeres — TTAGGG repeated thousands of times — are sacrificed to the fact that linear DNA cannot copy its own ends, so that the genes stay safe. When the buffer runs critically short, the cell senses damage and retires. That is the Hayflick limit, in molecules.
2. **The enzyme that keeps us young is the one cancer switches on to live forever.** Telomerase can refill the buffer, but most body cells switch it off — because an unlimited division counter is an unlimited tumor. Aging and cancer protection are entangled, two ends of one trade-off. In this field, "just turn the clock back"

usually means "remove a brake that was there for a reason."

3. **A measurement can be real across a population and nearly useless for one person.** Telomere length tracks aging in a crowd of thousands and tells you very little about yourself: a noisy assay, a weak individual prediction, possibly the wrong tissue, and "longer" isn't simply "better." *We can measure it* is not *you should act on it.* That sentence is the discipline the rest of this book runs on — and the mouse that seems to "prove" telomeres drive aging is, read closely, the clearest warning of all that mouse biology is not human biology.

Chapter 2 · The Clock and the Lever

EPIGENETIC CLOCKS ARE THE BEST MEASURE OF BIOLOGICAL *aging we have — which is exactly why they are the most dangerous place to push.*

> **Companion notebook · Chapter 2.** *Build Your Own Epigenetic Clock (simulated + real Hannum data)* — open it in Google Colab and build it yourself: Open in Colab →

IS THERE A BETTER CLOCK?

We left the last chapter with a complaint and a question. The complaint: telomere length is a low-resolution clock — noisy, single-tissue, weak for any one person. The question: is there a better one? A measurement that tracks how old you really are, tightly enough to be worth acting on?

There is, and the way it was found is a good story about what happens when you let a machine look for a pattern instead of guessing one.

For decades the assumption was that aging would look different in every tissue — that your skin and your liver and your brain were each running their own private clock, so any "biological age" measurement would have to be built tissue by tissue. In 2013 a researcher named Steve Horvath, working largely on his own and for a long time, did something almost stubborn. He gathered DNA methylation data — a kind of chemical annotation on the genome we will get to in a moment — from thousands of human samples spanning dozens of tissue types, all of known age, and asked a computer to find a single mathematical formula that would read someone's age off that chemical pattern, in *any* tissue at once.

It worked far better than it had any right to. One formula, the same few hundred sites with the same fixed weights, predicted chronological age across nearly every human tissue and cell type — blood, brain, breast, kidney, even some tumors — with a median error of only three to four years and correlations approaching 0.95 to 0.97 (performance varies somewhat by tissue and cohort). There was, it turned out, a kind of universal aging signature written in methylation, the same in your kidney and your skin, that nobody had designed and few had anticipated. The result was good enough that it was initially met with skepticism — partly because no one expected a single clock to work across so many different tissues, and on the reasonable grounds that biology is rarely that clean.

That formula is the **Horvath clock**, and the family of measurements it launched — *epigenetic clocks* — are the most accurate biological-age estimators we have. This chapter is about how they work, because they are the clearest example in the whole book of artificial intelligence doing real biology: a clock built entirely by machine learning out of high-dimensional data. And it is about the trap that comes free with them — a trap that is the same one from the last chapter, but sharper, because this time the predictor is not weak and noisy. It is precise. And precision, as we will see, makes the temptation far worse.

What a methylation mark is

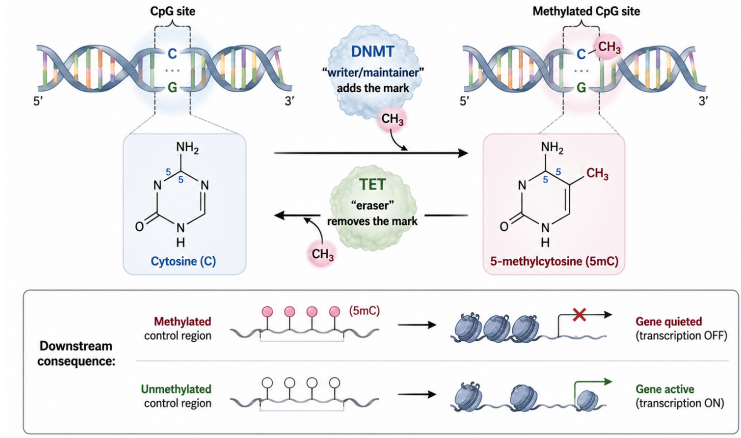
To read the clock you have to know what it reads. The substrate is **DNA methylation**, and the first thing to be precise about is the word *epigenetic*, because it is one of the most abused words in popular science.

An epigenetic change is a change in how a gene is *used* — switched on, switched off, read loudly or quietly — that is passed down when a cell divides, but that does *not* alter the DNA sequence itself. Same letters, different annotation. The sequence is the text; the epigenome is the markup — the highlighting and the margin notes that tell the cell which passages to read and which to skip. This is how one genome builds two hundred different cell types: a neuron and a liver cell carry identical DNA and behave nothing alike, because

different genes are annotated on and off in each. (Be careful with "inherited" here — we mean inherited from a cell to its daughter cells when it divides, which is the kind that matters for a clock. Inheritance across whole organisms, parent to child, is a separate and much more contested topic.)

DNA methylation is the best-understood of these marks. Chemically, it is the addition of a small methyl group to a cytosine — the C of the four DNA letters — turning it into 5-methylcytosine. In mammals this happens overwhelmingly where a C sits directly next to a G along one strand, a spot written **CpG** (the "p" is the phosphate linking them). There are roughly 28 million CpG sites in the human genome, and each one, in a given cell, is either methylated or not. Methylation in a gene's promoter or regulatory region often reduces its expression; loss of it often lets the gene speak more freely — though, as the rest of this chapter will stress, some methylation marks regulate genes directly while others may simply *mark* a broader cellular state without controlling anything themselves. Crucially, the mark is written and erased by dedicated enzymes — the DNMT family writes methylation (and copies the pattern onto new DNA when a cell divides, which is what makes it heritable through division), and the TET family removes it. So methylation is not graffiti the cell can't control. It is a managed, reversible, enzymatically maintained layer of information sitting on top of the sequence.

Figure 2.1 — A methylation mark.



Two features make methylation a near-perfect material for a clock. First, it is close to *digital*: at a given site in a given cell, the mark is essentially present or absent, which makes it cleanly measurable. Second, it is *measurable at enormous scale and low cost*. A methylation array reads hundreds of thousands of CpG sites at once — the common chips cover roughly 450,000 or, in newer versions, about 850,000 — and for each site it reports a **beta value**: the fraction of DNA molecules in the sample that are methylated there, a number between 0 (never methylated) and 1 (always). Run a sample and you get back a vector of several hundred thousand numbers, each between zero and one. Do that for thousands of people of known age, and you have exactly the kind of dataset machine learning handles well.

How you build a clock out of regression

Here is the part that matters most for a book about AI in medicine, so we will do it properly, with the actual mathematics, rather than waving at it.

You have, for each of N people, their true age and a long vector of methylation beta values — call them $x_1, x_2, \ldots, x_p$, where p is the number of CpG sites, around 450,000. You want a formula that turns those beta values into a predicted age. The simplest honest formula is a weighted sum:

$$\widehat{\text{age}} = w_0 + \sum_{i=1}^{p} w_i \, x_i .$$

Each site i gets a weight w_i — positive if that site tends to gain methylation with age, negative if it tends to lose it — and the predicted age is just the baseline w_0 plus every site's beta value multiplied by its weight. "Building the clock" means finding the weights. This is *linear regression*, the oldest tool in the statistics drawer.

But there is a problem that is also the whole point. You have around 450,000 sites and maybe a few thousand people. With that many knobs and that few examples, ordinary regression will fit the training data perfectly and learn nothing — it will memorize noise, contort itself to hit every training point exactly, and then predict garbage on anyone new. With

more dials than data, a model can always draw a curve through every point it has seen; that is not understanding, it is tracing. The fix is to *penalize* the model for using too many sites and for using any of them too strongly. Instead of just minimizing prediction error, you minimize prediction error *plus* a penalty on the weights:

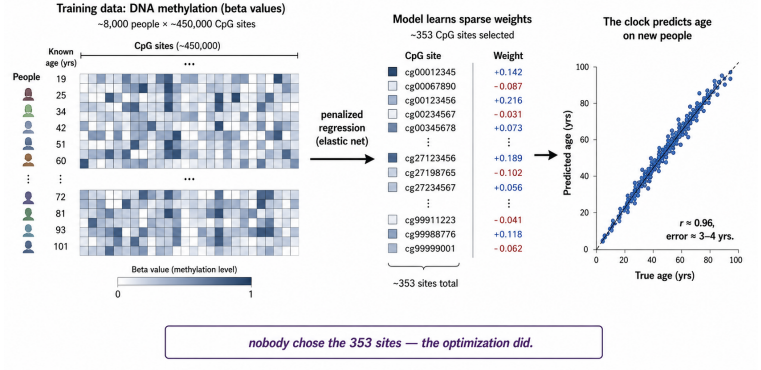
$$\min_{w} \underbrace{\sum_{j=1}^{N} \left(\widehat{\text{age}}_j - \text{age}_j\right)^2}_{\text{fit the data}} + \underbrace{\lambda\Big[\,\alpha \textstyle\sum_i |w_i| \;+\; (1-\alpha)\sum_i w_i^2\,\Big]}_{\text{stay simple}}.$$

The first term is the familiar squared error: make the predictions close to the truth. The second term is the discipline. The piece with $\sum_i |w_i|$ — the sum of the absolute values of the weights, called an L1 penalty — has a remarkable property: as you turn up its strength, it drives most weights to *exactly zero*. Not small. Zero. It performs *feature selection*, silently switching off the vast majority of the 450,000 sites and keeping only the handful that genuinely carry age information. The piece with $\sum_i w_i^2$ — the L2 penalty — instead spreads weight smoothly across correlated sites so the mod-
el doesn't lean its whole bet on one noisy spot. The knob α mixes the two, and λ sets how harsh the overall penalty is. This particular blend is the **elastic net**, and it is exactly what Horvath used.

The result is startling in its economy. Out of roughly 450,000 candidate sites, the elastic net keeps just **353**. Hor-

vath's entire multi-tissue clock — the thing that reads your age across blood and brain and kidney alike — is 353 CpG sites, each with a fixed weight, summed up. That is the whole model. A bright reader could evaluate it by hand with a calculator and the right 353 numbers.

Figure 2.2 — A clock is a trained model.



Notice what kind of object this is, because everything later depends on it. The model was fit to one job and one job only: minimize the gap between predicted and true age on the training people. The 353 sites were chosen *because their methylation correlates with age*. That is all the clock knows how to do — spot the methylation fingerprint that tends to come with a given number of birthdays. Keep that in your pocket.

Three generations, and the trick was the training target

A clock that predicts your age is, on its face, a strange thing to want. You already know your age. The value can't be in the prediction itself — it's in the *error*. When the clock reads older than your driver's license says, the hope is that the gap means something: that your cells are, in some real sense, fur-
ther along than your calendar age. That gap has a name — **epigenetic age acceleration** — and the history of the field is really the history of making that gap mean more.

The cleanest way to define the gap is not the naive subtraction "clock age minus calendar age," but the *residual*: the part of your clock reading left over after accounting for how old you actually are. Formally, you regress clock age on chronological age across everyone, and your acceleration is how far above or below that line you sit:

$$\text{age acceleration} = \text{clock age} - \mathbb{E}\big[\text{clock age} \mid \text{chronological age}\big].$$

A positive residual means the clock reads you as older than expected for your age. And in large studies, that residual does carry signal: people with accelerated first-generation clocks tend, on average, to die somewhat sooner and to carry more age-related disease. Real, but modest — and it is here that the field made its decisive move, and the move was pure machine learning.

The insight: if what you ultimately care about is health and mortality, why train the clock on *age* — a proxy — at all? Why not train it directly on what you actually want to predict? That single change of *training target* produced three generations of clocks:

The **first generation** (Horvath 2013; Hannum 2013) trained on chronological age. They predict age beautifully and health only weakly, for the reason above: they were optimized to recover the calendar, not your condition.

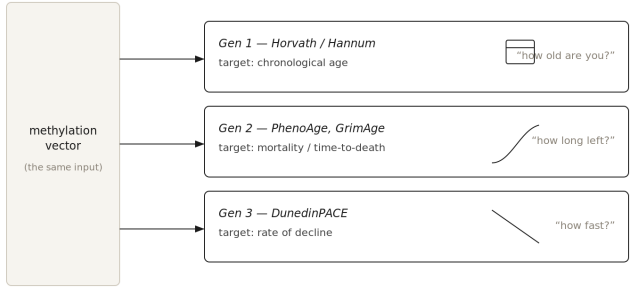
The **second generation** abandoned the calendar. *PhenoAge* (Levine, 2018) first built a "phenotypic age" out of clinical blood markers chosen because they predict death — things like inflammation, glucose, and kidney function — and then trained a methylation clock to predict *that*. *GrimAge* (Lu, 2019) went further still and trained directly on time-to-death, using methylation patterns as stand-ins for proteins in the blood and even for a person's smoking history. These mortality-trained clocks predict who will get sick and when far better than the first generation — because they were aimed at the real target instead of a proxy for it. GrimAge remains among the strongest.

The **third generation** changed the target again, from a *state* to a *rate*. *DunedinPACE* (Belsky, 2022) was built on an unusual and precious resource: a cohort of roughly a thousand people in Dunedin, New Zealand, all born in the same year and followed for decades, with the decline of nineteen different organ-system measures tracked over time. That let the

researchers measure each person's *pace* of aging — how many years of biological wear they accumulate per calendar year — and train a methylation model to estimate that pace from a single blood draw. DunedinPACE doesn't ask "how old are your cells?" It asks "how fast are you aging *right now*?" — which is, arguably, the more useful question, and certainly the more useful one for testing whether an intervention is slowing anything down.

Figure 2.3 — Same data, different question

one methylation vector in → three clocks, three training targets



The question you train on is the question you are really choosing.

This is a genuinely important lesson about machine learning in medicine, and it is bigger than clocks: **a model is only ever as good as the target you train it on.** Train on a proxy and you get a proxy. The art is almost never the architecture; it is choosing what the model should be trying to predict. The clock field's real progress came not from fancier math but from three successive, smarter answers to the question "predict *what*?"

It is worth pausing on an obvious skeptic's question, because the answer matters: are these clocks measuring anything *real*, or just fitting noise that happens to track birthdays? The evidence that they capture genuine biology is strong, and it comes from how they respond to things we already know harm the body. People with rare accelerated-aging syndromes read as biologically older on these clocks; so do people who smoke, who carry significant obesity, who are living with HIV, or who have been through chemotherapy. Confront a clock with a known biological insult, and it consistently registers it as accelerated aging — exactly what you would want from a real measure of aging, and a strong sign the clocks are not arbitrary. But notice that this *sharpens* the chapter's central worry rather than dissolving it. A meaningless number is easy to ignore; a number that genuinely tracks the body's condition is precisely the one you will be tempted to push on directly. Hold both facts at once: the clocks are real — and a real measurement is still not a lever.

> **Deeper Water — why the penalty makes the clock sparse.** It is worth seeing *why* the L1 penalty zeroes out weights while the L2 penalty doesn't, because the reason is geometric and it recurs all over modern machine learning. Picture the weights as a point you're free to move, with two forces acting on it: the data pulls it toward the best-fit values, and the penalty pulls it

toward the origin (all weights zero). The L2 penalty's "budget" of allowed weight is a sphere — smooth, round — and the best-fit solution that just touches that sphere almost never touches it at a point where a coordinate is exactly zero; you get many small nonzero weights. The L1 penalty's budget is instead a diamond (in high dimensions, a cross-polytope), and a diamond has sharp corners that stick out *along the axes* — exactly the points where most coordinates are zero. A surface touching a spiky shape tends to touch it at a spike, so the L1 solution lands on a corner, setting most weights to zero and keeping a sparse few. That is why the elastic net can sift 450,000 sites down to 353: the L1 corner does the sifting, and the L2 roundness keeps the survivors stable when several of them carry overlapping information. Sparsity isn't a bonus feature bolted on. It falls out of the shape of the penalty.

Is the clock cause or consequence?

If methylation reliably changes with age — reliably enough to read age to within a few years — a natural and momentous question follows. Is that changing methylation *driving* aging, or merely *recording* it? Is the epigenome the cause of the breakdown, or the breakdown's logbook?

This is the heart of a hypothesis associated most prominently with David Sinclair, and it deserves to be stated at its strongest, because it is genuinely interesting. The idea is that aging is, at bottom, a loss of *epigenetic information*. Over a lifetime, the cell's careful pattern of marks — the markup that tells it which genes to read — slowly degrades and grows noisy, like scratches accumulating on a disc until the music skips. The DNA text stays intact, but the cell increasingly forgets *which cell it is supposed to be*, and that forgetting is the aging. The provocative corollary is that if the information is merely misfiled rather than erased, it might be *restored* — the cell reminded of its youthful identity — which is the thread we pick up in Chapter 7, on cellular reprogramming.

The hypothesis comes with real experiments, not just analogy. In a provocative and much-debated 2023 study, Sinclair's group built mice in which they could induce breaks in the DNA that the cell repairs *without* changing the sequence — breaks designed to disturb the epigenome while leaving the text untouched. Those mice showed accelerated signs of aging *and* advanced epigenetic clocks, and partial reprogramming pushed some of it back. To the hypothesis's supporters, that is close to a demonstration that epigenetic disruption alone can drive aging.

But this is exactly the kind of claim the book has trained you to hold at arm's length, and two cautions apply. The first is the now-familiar one from Chapter 1: those mice underwent repeated, deliberately induced DNA breaks that do not

obviously resemble ordinary aging — an artificial perturbation, not the gentle drift of the years. Showing that a violent epigenetic insult accelerates aging does not establish that ordinary aging is *caused* by epigenetic drift rather than accompanied by it. The second caution is deeper and is the subject of the rest of this chapter: the clocks themselves cannot settle the cause-or-consequence question, because **a clock is a correlation, and a correlation, however tight, is silent about cause.** The very precision that makes the clock impressive is precisely what makes it unable to tell you whether the methylation it reads is steering the ship or just watching from the deck.

So the honest status is: methylation change is one of the most reliable *correlates* of aging we have ever found, and whether it is a primary *driver* is open and contested. Reasonable scientists land on different answers. The clock does not adjudicate between them — and anyone who tells you it does has skipped the most important step.

The most important sentence in this book

Here it is, and it is worth setting apart, because if you take one idea from these pages it should be this one:

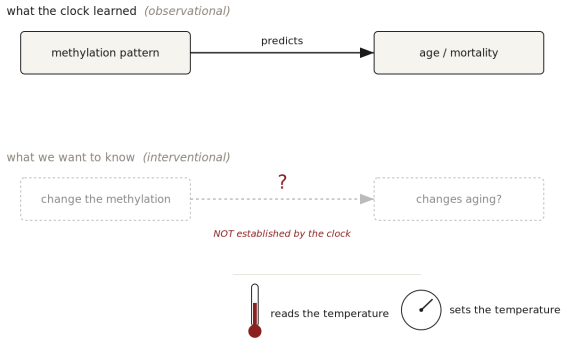
A measurement that predicts an outcome is not the same as a handle that controls it.

The epigenetic clock is the cleanest illustration there will ever be. The clock was *fit* to predict age (or mortality) from

methylation. It selected the sites it uses *because* their methylation correlates with the outcome. That makes it an excellent *thermometer* for aging. It says nothing whatsoever about whether those sites are a *thermostat* — about whether *changing* the methylation, or *moving the clock reading*, would change the aging itself.

Confuse the two, and you get the defining error of the entire longevity-marketing world. The clock reads your "biological age" at, say, fifty-two. You take some intervention — a supplement, a diet, a drug — and the clock now reads forty-nine. The pitch: you reversed three years of aging. The reality: you moved a *predictor* by three years, and whether you moved *aging* by even a day is completely unestablished by that fact. You may simply have nudged the specific methylation sites the clock happens to watch, while the underlying biology rolls on unchanged. In machine-learning terms, you may have *taught to the test* — improved the score on the metric without improving the thing the metric was supposed to stand for. The clock is a model, and any model can be gamed by an intervention that targets its inputs rather than the world.

Figure 2.4 — Thermometer, not thermostat



You can see the trap operating in real studies. The most cited example is a small 2019 trial in which nine men took a combination of growth hormone and two other drugs for about a year, and several epigenetic clocks read them as roughly one to two years "younger" afterward. It made headlines as the first reversal of human aging. But pause on what it actually was: nine people, no control group, a cocktail with its own serious risks, and an outcome measured entirely by clock readings. Moving a clock in a tiny uncontrolled study is a *hypothesis*, not a result — and it is hypothesis-generating in exactly the way that should make you cautious, not excited, because the one thing it cannot rule out is that the drugs nudged the clock's inputs without touching anyone's lifespan.

What would it take to trust a clock as a measure of an intervention's worth? In drug development the concept is a **surrogate endpoint** — a marker you're allowed to use *in*

place of the real outcome (living longer, getting sick less) because moving the marker has been *shown* to reliably move the real outcome. A surrogate has to earn that status: you must demonstrate, across many different interventions, that whenever you move the marker you also move the hard endpoint in step. LDL cholesterol earned much of that status for heart disease the hard way, over decades and many trials — and even there, drugs that lowered cholesterol by the "wrong" mechanism sometimes failed to help, a reminder of how easily a marker and an outcome can come apart. The epigenetic clocks have *not* earned surrogate status for aging. They are candidate surrogates — promising, unvalidated. This is not a technicality; it is why a regulator will not yet accept "we lowered the patient's GrimAge" as evidence that a drug works, and why you shouldn't either. A precise predictor is still only a predictor until something proves it is also a lever.

Deeper Water — predicting versus pushing. The distinction at the core of this chapter has an exact formal form, and it is one of the most important ideas in modern statistics and AI. When you fit a clock, you learn an *observational* relationship — in the notation of causal inference, the distribution of the outcome given what you *see*, $P(\text{age} \mid \text{methylation})$. What you actually want to know before intervening is differ-

ent: the distribution of the outcome given what you *do*, written $P(\text{age} \mid \text{do}(\text{methylation}))$ — the result of *setting* the methylation to a value rather than merely *observing* it. These are not the same object, and no amount of observational data, however large or clean, can convert one into the other on its own; bridging them requires either a real intervention (an experiment) or strong, explicit causal assumptions. There is also a concrete, mundane way clocks get fooled that lives inside this gap: **cell-composition confounding.** Blood is a mixture of cell types, and their proportions shift with age (your naive-T-cell fraction falls, for instance). A clock measured on whole blood can be partly reading the *changing mixture* rather than the aging of any individual cell — so an intervention that alters your blood-cell proportions can move the clock while changing nothing about cellular aging at all. Researchers now try to separate "intrinsic" age acceleration (independent of cell mix) from "extrinsic" (which includes it), precisely because the clock you naively measure is reading two different stories at once. Even the world's best aging predictor doesn't fully know what it's keying on.

Where the machines come in

This whole chapter has been the AI chapter — the clocks *are* machine learning, the elastic net is a learning algorithm, and the generational leap forward was a lesson in choosing training targets. But it is worth asking the two questions a drug-discovery reader would ask next.

First: would a fancier model do better? People have built deep-learning clocks — neural networks trained on the same methylation data instead of a penalized linear regression. The honest and slightly deflating answer is that they perform *comparably* to the humble elastic net, with at best modest gains. That is itself one of the most useful results in the field, and it echoes the last chapter exactly: when the signal is what it is, piling on model complexity does not manufacture more truth. A 353-site linear formula captures most of what a deep network does, because the limit here is the information in the data, not the cleverness of the model. More parameters do not create more reality.

Second, and this is where AI in aging could actually matter: can we use clocks to *find drugs*? The obvious move is to run candidate compounds through cells and screen for whichever ones lower the epigenetic clock the most — and people are doing exactly this. But you can now see the danger written into it. A screen that rewards "lowers the clock" will faithfully hand you the compounds that are best at *gaming the clock*, which is not the same as the compounds

that slow aging — unless and until the clock is a validated surrogate. Used carelessly, a clock-based screen is an efficient machine for finding clever ways to teach to the test.

The genuinely promising frontier runs the other way: using these tools to move *from correlation toward cause*. If methylation at a specific site is suspected of actually driving an aging phenotype, you can now go and *edit just that site* — using engineered proteins (a disabled Cas9 fused to a methylation-writing or methylation-erasing enzyme) to add or remove the mark at one location, then watch what the cell does. That converts the clock's correlations into testable causal claims, one site at a time, and it is precisely the experiment that could eventually tell us which of the clock's 353 sites are thermostats and which are merely thermometers. That is where machine learning and the wet lab together could turn a beautiful predictor into an actual drug target — by doing the one thing the clock alone can never do, which is to *intervene* and see.

✦ ✦ ✦

Carry these three things

1. **An epigenetic clock is a machine-learning model, and a strikingly economical one.** Out of roughly 450,000 methylation sites, a penalized regression keeps a few hundred and sums their weighted values to read your age to within three or four years. The progress

from first-generation clocks to GrimAge to Dunedin-PACE came not from fancier math but from smarter training targets — from predicting the calendar, to predicting death, to predicting the *rate* of decline. A model is only as good as the target you train it on.

2. **A measurement that predicts an outcome is not a handle that controls it.** The clock is a thermometer for aging, not a thermostat. Its sites were chosen because they *correlate* with aging, which says nothing about whether *changing* them changes anything. Moving a clock — especially in a tiny, uncontrolled study — is a hypothesis, not a reversal of aging. Until a clock is shown to be a valid *surrogate* (move it, and the real outcome reliably moves with it), "we lowered your biological age" is a claim about a model's output, not about your life.
3. **Correlation is silent about cause — and precision doesn't change that.** The clock cannot tell you whether methylation drives aging or merely records it; that is the open, contested question behind the "information theory of aging," and no predictor, however tight, can settle it. The way forward is to *intervene* — to edit single sites and watch — because only an experiment can turn a correlation into a cause. The clocks tell us, better than anything before them, *that* we are aging and roughly how fast. They do not tell us *what is break-*

ing. For that we have to open the mechanisms — and the next one refuses to die.

Chapter 3 · The Cells That Won't Die

SENESCENT CELLS PROTECT YOU FROM CANCER WHEN YOU'RE *young and poison your tissues when you're old — which is why clearing them is the field's most beautiful idea, and so far its most expensive disappointment.*

> **Companion notebook · Chapter 3.** *When Cleanup Loses: the Senescent-Cell Tipping Point* — open it in Google Colab and build it yourself: Open in Colab →

THE CELL THAT REFUSES TO DIE

When Leonard Hayflick watched his fibroblasts stop dividing after fifty doublings, he assumed they were simply worn out — old cells running down like a clock with a tired spring. For a long time that is how everyone read it: senescence was passive exhaustion, the cellular equivalent of fatigue. It took decades to see that this was almost exactly backwards. A senescent cell is not a cell that has run down. It is a cell that has been deliberately, actively *switched into a new state* — one it

cannot leave, one with a purpose, and one that turns out to be far more sinister for the tissue around it than mere exhaustion would ever be.

A senescent cell stops dividing, permanently. But it does not die. It sits in your tissue, alive and metabolically busy, often swollen to several times its normal size — and it spends its time secreting a steady cocktail of inflammatory and tissue-corroding molecules into its surroundings. It won't divide, it won't die, and it harms the living cells nearby. The popular name for this is the **zombie cell**, and for once a piece of popular-science shorthand is precise rather than misleading — but only if we are exact about what we mean. The cell is "undead" in three specific, mechanistic senses: it has been locked out of the cell cycle so it cannot proliferate; it has actively reinforced the pathways that keep it from dying, so it resists the normal cellular suicide that would clear it; and through what it secretes, it actively damages the healthy tissue it lives in. Not a tired cell. A cell that has chosen retirement over death, and makes its neighbors pay for the choice.

We met the first reason a cell does this in Chapter 1: when a telomere gets critically short and is read as DNA damage, the cell withdraws permanently from division — replicative senescence, the molecular basis of the Hayflick limit. But telomere exhaustion is only one trigger. A cell will also flip into senescence if an oncogene fires and threatens to turn it cancerous (a beautiful piece of self-defense called oncogene-

induced senescence), if it suffers DNA damage it cannot repair, if it is poisoned by oxidative or metabolic stress, or if chemotherapy batters it. The common thread is danger — specifically, the danger of a damaged cell continuing to divide. And that tells you what senescence is *for*.

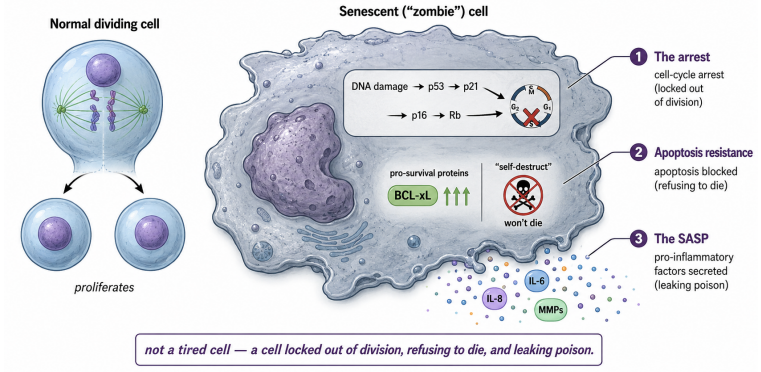
What a senescent cell actually is

Underneath the zombie metaphor is a precise piece of cell-cycle biology, and it is worth getting right, because the same machinery shows up again when we try to kill these cells.

For many proliferative tissues, dividing is the default state; senescence is an enforced, locked exit. The lock has two main bolts. The first runs through the protein **p53** — the cell's master damage sensor from Chapter 1 — which, when persistent damage is detected, switches on a small protein called **p21**. The second runs through a protein with the unlovely name **p16^INK4a**. Both p21 and p16 are *cyclin-dependent kinase inhibitors*: they jam the enzymes (the CDKs) that the cell normally uses to push itself through the cycle from one division to the next. With those enzymes jammed, a gatekeeper protein called Rb stays clamped down on the genes needed to copy DNA, and the cell simply cannot enter the next round of division. Throw both bolts and reinforce them — the cell even repackages the relevant stretches of its genome into dense, silent heterochromatin so the proliferation genes can't be switched back on — and the arrest becomes effectively

permanent. This is why senescence is not reversible the way a cell merely pausing between divisions is. The exit has been welded shut.

Figure 3.1 — Anatomy of a zombie cell.



Now an honest complication that matters for everything downstream: **there is no single, reliable marker that says "this cell is senescent."** The classic stain — for a lysosomal enzyme abbreviated SA-β-gal — is suggestive but not specific; p16 is a workhorse marker but not universal; the enlarged shape, the loss of a nuclear-envelope protein called Lamin B1, the persistent DNA-damage signals, the secretory profile — each is a hint, none is proof. Senescence is a *heterogeneous* state: a senescent liver cell, a senescent immune cell, and a senescent fibroblast can look and behave quite differently, and a cell can be partway in. Researchers identify senescence by a *combination* of features, not a single test. Keep this in your pocket too, because it is one of the central practical reasons the therapies are hard — if you cannot cleanly say

which cells are senescent, you cannot easily confirm that a drug cleared the right ones, and (as we will see) a machine-learning problem is hiding right here.

The poison they spread

The most consequential thing about a senescent cell is not that it stopped dividing. It is what it secretes. The cocktail has a name — the **senescence-associated secretory phenotype**, the SASP — and it is the mechanism by which a handful of retired cells can degrade an entire tissue.

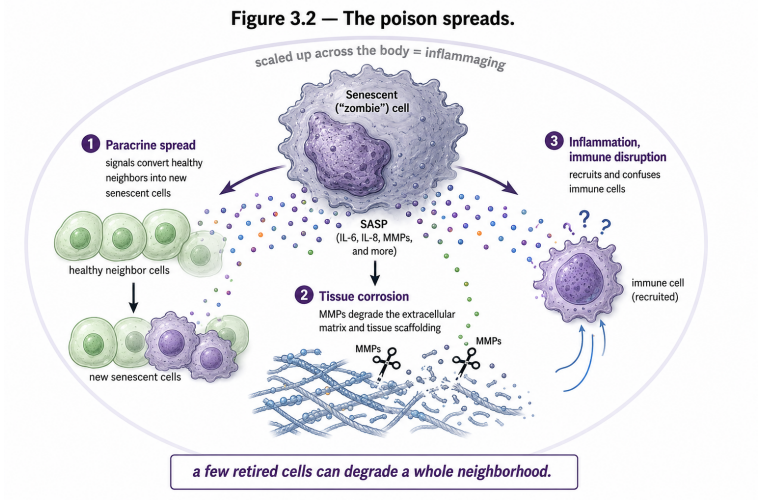
The SASP is a brew of inflammatory signaling molecules (cytokines such as IL-6, IL-1, and IL-8), chemical "come here" signals that summon immune cells, growth factors, and enzymes called proteases (the MMPs) that chew up the structural scaffolding between cells. It is driven from inside the cell largely by a master inflammatory switch called NF-κB, with a contribution from an ancient innate-immune alarm — the cGAS-STING pathway — that normally detects viral DNA but here is tripped by the cell's own damaged DNA fragments leaking where they shouldn't be. The senescent cell, in other words, partially activates the pathways that normally respond to an infection, and calls in the immune system accordingly.

This screaming does three kinds of damage. It inflames the local tissue, chronically. It corrodes the surrounding structure through the proteases. And — most insidiously —

it pushes *neighboring healthy cells into senescence too*, through the very inflammatory signals it is releasing. A senescent cell can recruit its neighbors into the same state, which is called paracrine or bystander senescence. A few zombie cells become a spreading patch.

Zoom out from one tissue to the whole body, and the SASP is one major source of a hallmark of aging we are otherwise not giving its own chapter — **chronic inflammation**, often called *inflammaging*. As you age, the slow, smoldering, body-wide inflammation that quietly raises your risk of nearly every age-related disease, from atherosclerosis to dementia, appears to be driven in important part by accumulated senescent cells, each one a tiny inflammatory siren — though they are not the only contributor: aging of the innate immune system, a leakier gut, inflamed visceral fat, and lingering chronic infections all feed the same fire. Still, senescence is one of the most mechanistically concrete contributors we can point to, which is why we treat inflammaging here, as a major consequence of senescence, rather than as its own thing: in this book, much of the chronic inflammation of aging has a traceable cellular address.

Figure 3.2 — The poison spreads.



WHY THEY PILE UP

If senescent cells are so harmful, why does the body let them accumulate? The answer is a balance between two rates, and it is worth writing down, because the arithmetic explains both why the problem accelerates with age and what the two ways to attack it are.

Let N be the burden of senescent cells in a tissue. They are *produced* at some rate P — new cells flipping into senescence from telomere exhaustion, oncogenic stress, and accumulated damage — and they are *cleared* by the immune system (natural-killer cells and macrophages recognize and remove them) at a rate proportional to how many there are, with a clearance constant c. Then

$$\frac{dN}{dt} = P - cN,$$

which settles at a steady-state burden of

$$N^* = \frac{P}{c}.$$

Here is the trouble. **Aging pushes both terms the wrong way at once.** The production rate P rises — more cells reach critically short telomeres, more accumulate irreparable damage, more oncogenic near-misses occur, as the years go by. And the clearance constant c falls — because the immune system itself ages and grows less effective at removing senescent cells (a decline called immunosenescence). Worse, the decline in clearance is partly *caused by the SASP*, which disrupts immune function: a vicious circle, the zombies sabotaging the very janitors meant to remove them. With the numerator climbing and the denominator shrinking, the steady-state burden $N^* = P/c$ does not creep upward linearly. It accelerates. That is why senescent-cell burden is low and well-managed for most of life and then climbs steeply in old age — not because anything new switches on, but because a balance that was comfortably in your favor tips, from both ends, into a steep climb.

The same equation lays out the two therapeutic strategies cleanly. You can attack the burden by *raising effective clearance*

— drugs that kill senescent cells, which the field calls **senolytics**, effectively boosting c. Or you can leave the cells in place but *muzzle their secretions* — drugs that suppress the SASP, called **senomorphics** or senostatics, which lower the harm per cell rather than the number of cells. Most of the excitement, and the chapter's hard lesson, lives in the first strategy.

Deeper Water — when the janitors saturate, and the burden runs away. The simple model on the main road, $dN/dt = P - cN$, hides a comfort in its structure: no matter how high production P climbs or how far the clearance constant c falls, the burden always settles at *some* finite value, $N^* = P/c$. It can grow large, but it never truly runs away — there is no tipping point, only a smoothly steepening climb (halve c and you double N^*). Real immune clearance, though, has a feature this linear term ignores: it *saturates*. You have only so many natural-killer cells and macrophages, and they can remove senescent cells only so fast; past a certain burden the cleanup crew is working flat out and simply cannot speed up. A standard way to write a clearance rate that rises with burden but plateaus at a ceiling is $dN/dt = P - cN/(K + N)$, where c is

now the *maximum* clearance rate — the most cells the immune system can mop up per unit time, a different quantity (and different units) from the per-cell constant in the linear model — and K is the burden at which clearance runs at half that maximum. At low burden ($N \ll K$) this behaves almost like the linear model, clearing at roughly $(c/K)\,N$; at high burden ($N \gg K$) the clearance term flattens out at c, the crew maxed out. That one change does something the linear model structurally cannot: it creates a genuine threshold. Setting $dN/dt = 0$ gives a steady state *only* when production stays below the clearance ceiling, $P < c$, at which point $N^* = PK/(c - P)$ — an expression that blows up as $P \to c$, the denominator going to zero and the burden to infinity. And once production *exceeds* the ceiling, $P > c$, there is no steady state at all: because the clearance term can never exceed c, the burden grows without bound. The system has *tipped*. This is a sharper and more honest picture of late-life acceleration than the linear model offers — senescent burden may not just be climbing a steepening curve but approaching, and then crossing, a clearance threshold past which accumulation becomes

self-sustaining. And the main road's vicious circle makes it worse from the other side: as the SASP degrades immune function, it drags the ceiling c downward toward the rising production P, hurrying the system toward the tipping point. It also reframes what a therapy must achieve: nudging the burden down a little is futile if you are past the tipping point — you have to push the system back *across* it, lifting effective clearance back above production or cutting production back below clearance, to escape the runaway. (This is still a toy model: real senescence is heterogeneous, spread across tissues and compartments, with feedbacks on both P and c at once. But it captures a qualitative truth the linear form cannot — a system with finite clearance capacity does not merely degrade, it can collapse.)

The brake and the price, again

Before we try to kill these cells, we have to confront a question that should, by now, feel familiar — because it is the same question the telomere chapter asked, and the answer rhymes. If senescent cells are so destructive, why does evolution permit them at all? Why build a cellular state whose whole job in old age is to poison you?

Because in youth, senescence is not a bug. It is one of your best defenses, and it does several genuinely good jobs. It is a powerful *tumor suppressor*: a cell that detects it is becoming cancerous and responds by locking itself permanently out of division is a cell that cannot build a tumor — the same brake-on-cancer logic as the telomere limit in Chapter 1, now generalized. It is essential for *wound healing*: transient senescent cells appear at injury sites and secrete signals that orchestrate repair, then are cleared once the job is done. It even helps *sculpt the embryo*, with programmed senescence shaping certain tissues during development. Acute, temporary, cleared-on-schedule senescence is a tool. It is *chronic, accumulated, uncleared* senescence — the same mechanism, left running too long and never mopped up — that turns toxic.

This is one of the cleanest real examples of an idea at the center of why aging exists at all: **antagonistic pleiotropy.** A trait that helps you when you are young and reproducing can be favored by evolution *even if it harms you when you are old*, because natural selection cares overwhelmingly about early life — about getting you to reproduce — and is nearly blind to what happens to you decades later. Senescence pays its dividend early (no cancer, good healing) and presents its bill late (inflammaging, tissue decay), and by the time the bill arrives, selection has already gotten what it wanted from you. Aging, in this view, is not a failure of evolution. It is partly the *price* of mechanisms that were selected precisely because they served you when it counted.

And from that follows the intervention caution that this book keeps returning to, now in its third chapter and its third form. In Chapter 1 it was "don't just turn on telomerase — you'd be removing a brake on cancer." Here it is: *don't just kill every senescent cell* — because some senescence is doing necessary work, and blanket clearance could cripple wound healing, impair tumor suppression, and have effects we cannot fully predict. The obvious intervention is, once again, not free. The real target is not "destroy all senescent cells" but the much harder "selectively remove the *chronic, harmful* ones while sparing the *acute, useful* ones" — and we do not yet have a clean way to tell them apart.

> **Deeper Water — why evolution allows aging.** Antagonistic pleiotropy, proposed by George Williams in 1957, is one of three classic, complementary explanations for why aging exists, and they are worth holding together. The oldest, from Peter Medawar, is the *mutation-accumulation* theory: because selection weakens with age (anything that only harms you after you've reproduced is barely "seen" by natural selection), harmful mutations whose effects appear only in late life accumulate in the genome, simply because nothing efficiently weeds them out. Williams' antagonistic pleiotropy goes further: late-life harm isn't merely tolerated, it can

be *actively selected for* when it is the flip side of an early-life benefit — the same gene doing good early and ill late. The third, the *disposable soma* theory (Tom Kirkwood), frames it as an energy-budgeting problem: an organism has finite resources to split between reproduction and bodily maintenance, and evolution tunes the split to maximize reproductive success, not longevity — so the body (the "soma") is built to last just long enough, then allowed to decay. Senescence sits squarely inside all three: a late-life-harmful state, the cost of an early-life-beneficial defense, and an example of maintenance that the body could in principle keep up but doesn't. None of these theories is "the" answer; aging is overdetermined, with several true explanations at once. But together they deliver the unsettling and clarifying punchline that aging is not a disease in the ordinary sense, nor an accident — it is, in large part, the designed-in price of being an organism that was optimized to reproduce rather than to last.

Killing the zombies

Now the exciting part — and the reason senescence became, for a while, the hottest target in all of longevity science.

The logic is irresistibly clean. If accumulated senescent cells drive aging, then *removing them* should reverse the damage. And the first real test of that idea is one of the most striking experiments in the field. In 2011, a team led by Jan van Deursen built mice carrying an ingenious genetic switch — a system (nicknamed INK-ATTAC) that let them kill, on command with a drug, cells expressing high levels of the senescence marker p16 — many, though not all, of which are senescent. When they threw that switch and cleared the senescent cells, the mice aged visibly better: less deterioration in muscle, fat, and eye. Five years later, the same group showed that clearing senescent cells from *naturally* aged mice — not an engineered fast-aging model, ordinary old mice — extended their healthy lifespan and pushed out median survival. Remove the zombie cells, and the animals lived better and somewhat longer. It is hard to overstate how electrifying that was: a single, targeted intervention against one cell state, improving aging across multiple organs at once.

The converse experiment seals the case. If removing senescent cells helps, does *adding* them harm? In 2018, a group led by Xu and Kirkland transplanted relatively small numbers of senescent cells into young, healthy mice — and the recipients developed physical dysfunction (slower walking speed, weaker grip, less endurance), persistent ill health, and, strikingly, *spreading* senescence in their own tissues, as the transplanted cells recruited the animals' native cells into the same state. A modest dose of zombie cells was enough;

transplanting the same number of normal, non-senescent cells did nothing. Put the two experiments side by side and the causal logic closes from both directions: clear senescent cells and aged mice improve; introduce senescent cells and young mice deteriorate. Senescent cells are not merely *correlated* with the damage of aging — at least in mice, they can actively *cause* it.

But van Deursen's clearance experiment used a genetic switch built into the mice. To do it in a person, you need a *drug* that selectively kills senescent cells — a senolytic. And here the apoptosis-resistance from earlier becomes the key that unlocks the therapy. Recall that a senescent cell actively reinforces the pathways that keep it from dying. It survives only by *leaning hard* on a specific set of pro-survival proteins — its anti-apoptotic crutches. A normal, healthy cell isn't leaning on those crutches; it has no pressing reason to die, so it isn't propping itself up against death. That difference is the therapeutic window. **Knock the crutch out, and the senescent cell — which depended on it — topples into the suicide it was holding off, while the healthy cell barely notices.** Senolytics are drugs that knock out the crutch.

The first ones were found this way, by hypothesis rather than by accident. In 2015, James Kirkland's group reasoned out which survival pathways senescent cells depend on and tested drugs against them, identifying *dasatinib* (a cancer kinase-inhibitor) and *quercetin* (a plant flavonoid) — dasatinib preferentially killing some senescent cell types and quer-

cetin others, with the combination proving effective across a broader range than either drug alone. Others followed: *fisetin*, another plant flavonoid (though its supporting evidence remains weaker than for dasatinib-and-quercetin); *navitoclax*, a potent blocker of the BCL-xL survival protein. And these drugs carry a genuinely novel and appealing feature. Most medicines must be taken continuously, because they work only while the drug occupies its target. A senolytic doesn't: once you kill a senescent cell, it is *gone*, and it takes time for new ones to accumulate. So senolytics could in principle be given in brief, intermittent courses — a "hit-and-run" approach, dosing for a few days, then stopping for weeks or months. That is a genuinely different therapeutic paradigm, and part of why the idea generated so much hope, and so much money.

Deeper Water — the crutch, and the cost of pulling it. The senescent cell's survival crutches are known collectively as the senescent-cell anti-apoptotic pathways, and they include the BCL-2 family (especially BCL-xL), the PI3K/AKT survival axis, and a p53-regulating interaction involving a protein called FOXO4. Different senolytics pull different crutches — dasatinib and quercetin hit a broad set including kinase- and survival-signaling; navitoclax specifically blocks BCL-xL. And navitoclax illustrates, pain-

fully, why "knock out the crutch" is harder than it sounds: BCL-xL is not used *only* by senescent cells. Platelets — the cell fragments that clot your blood — also depend on it to survive, so navitoclax causes their numbers to crash (thrombocytopenia), a dose-limiting toxicity that has hampered its use. The therapeutic window is real but it is narrow and target-dependent: the more a survival pathway is shared between zombie cells and important healthy cells, the smaller the window. This is the molecular reason senolytics cannot simply be made arbitrarily potent — and a preview of why moving from a clean mouse experiment to a safe, effective human drug is so much harder than the elegant logic suggests.

The most expensive lesson so far

Everything above is true, and the mouse data are real, and the logic is beautiful. And the first major test of that logic in human beings failed — publicly, expensively, and instructively.

The company was **Unity Biotechnology**, and it is important to understand that this was not some fly-by-night operation. Unity was founded by the field's leading lights — including Judith Campisi, who reframed senescence as a driver

of aging, and van Deursen and Kirkland, whose work you have just read — and it was backed by some of the most sophisticated investors in the world, with high-profile names and hundreds of millions of dollars behind it. When it went public in 2018, it carried the field's brightest hopes. Its lead program, UBX0101, was a senolytic designed to be injected directly into an arthritic knee, to clear the senescent cells thought to be driving the inflammation and cartilage breakdown of osteoarthritis. The strategy was local (inject just the joint, sidestep whole-body toxicity), the target was concrete, and the science was first-rate.

In 2020, the Phase 2 trial results came back. UBX0101 did no better than placebo at relieving knee-osteoarthritis pain. The drug failed its primary endpoint. The company's stock collapsed, and Unity reorganized around a different program entirely — a senolytic for retinal eye diseases, which has since produced more encouraging but still early signals.

Sit with what that failure was and was not. It was *not* a failure of chemistry, or of artificial intelligence, or of effort, or of expertise. The people involved understood senescence as well as anyone alive. It was, instead, the same failure this book has now shown you three times in three chapters, in three different costumes: a gorgeous, well-supported result in mice that did not transfer to humans. Maybe the senescent cells weren't actually the dominant driver of osteoarthritis pain in people the way they appeared to be in mouse models. Maybe a single injection didn't clear enough of them, or the

right ones. Maybe the heterogeneity of senescence meant the drug killed some senescent cells but not the ones that mattered — and remember, with no clean marker, it was hard to even confirm what had been cleared. Maybe human joints and mouse joints simply differ in the way that made the experiment work. Any of these could be true; the trial cannot, by itself, tell you which. What it tells you is that the inference "clearing senescent cells fixes the disease" — so clean in the mouse — broke somewhere on the way to the human, exactly as the telomere story and the clock story broke before it.

And now the honest, balanced status of the whole field, because the lesson is "be disciplined," not "give up." Senolytics are *not* dead. Small human trials have shown that drugs like dasatinib-and-quercetin really can reduce the burden of senescent cells in people — the first direct evidence that human senolysis is even possible — in conditions ranging from a lung-scarring disease to diabetic kidney disease, and trials in other diseases including Alzheimer's are underway. The eye program has shown promise. But as of now, the truth is this: no senolytic has been approved for anything; no large trial has clearly demonstrated the anti-aging — or even firm disease-modifying — promise at scale; and the first marquee test of the whole idea failed. Senescence remains one of the most compelling targets in aging biology *and* one of its least proven therapies. Both halves are true at once. A reader who hears only "mice lived longer

when we cleared zombie cells" and not "the first big human trial failed" has, once again, learned the wrong half.

Where the machines come in

Two places in this story are genuinely machine-shaped, and one of them traces directly back to a problem we flagged early in the chapter.

The first is the marker problem. Because there is no single test for senescence — because it is a heterogeneous state recognized only by a *combination* of features — identifying senescent cells reliably is exactly the kind of pattern-recognition task that machine learning is built for. Researchers now train deep-learning image classifiers to spot senescent cells from their morphology and from multiplexed molecular imaging, learning the combination of subtle features that no single stain captures. A major coordinated effort (the NIH's Cellular Senescence Network, SenNet) is building tissue-by-tissue atlases of senescent cells and the classifiers to find them. This matters because better detection is the precondition for everything else: you cannot study, target, or confirm the clearance of cells you cannot reliably identify. Note, though, the now-familiar caution carried over from Chapter 2 — a classifier that *predicts* "this cell is senescent" is still a predictor, and it is only as good as the labeled examples it learned from, in a field where even the experts' labels are uncertain.

The second is drug discovery itself: using computational screening to find new senolytics and senomorphics — predicting which compounds will selectively pull a senescent cell's survival crutch, and, more valuably, designing ones with *wider* therapeutic windows than navitoclax, targeting survival pathways that senescent cells depend on but vital healthy cells don't. This is where artificial intelligence could earn real ground, because the selectivity problem — kill the zombie, spare the platelet — is partly a molecular-design problem, and molecular design is exactly what the newest AI tools are getting good at.

But hold onto the chapter's hardest lesson when you weigh these. The Unity trial did not fail for want of a better algorithm or a cleverer molecule. It failed because a piece of mouse biology did not hold in a human being. AI can sharpen our detection of senescent cells and accelerate the search for drugs that kill them. It cannot, on its own, tell you whether clearing those cells will cure a human disease — that question is answered only in the clinic, and so far the clinic's answer has been "not yet." The binding constraint in this field is not computation. It is the stubborn, expensive gap between a mouse and a person.

✦ ✦ ✦

Carry these three things

1. **A senescent cell is not exhausted — it is locked into an active, harmful state on purpose.** Triggered by danger (a short telomere, an oncogene, irreparable damage), it bolts itself out of the cell cycle with p21 and p16, reinforces its resistance to dying, and pours out the SASP — an inflammatory, tissue-corroding secretion that spreads senescence to its neighbors and, summed across the body, is much of what we call inflammaging. The chronic inflammation of aging has a cellular address.
2. **Senescence is the brake-and-price trade-off again, now as an evolutionary principle.** It suppresses cancer and aids healing when you are young, and poisons your tissues when you are old — a textbook case of antagonistic pleiotropy, a trait selected for its early benefit despite its late cost. So "kill every senescent cell" is, like "turn on telomerase," removing a brake that was there for a reason; the real target is the chronic harmful cells, not the acute useful ones, and we cannot yet cleanly tell them apart. They accumulate because aging raises their production and lowers their clearance at the same time — $N^* = P/c$, with the numerator climbing and the denominator shrinking — which is why the burden doesn't creep but accelerates.

3. **The most beautiful idea in the field met the most expensive disappointment.** Clearing senescent cells rejuvenated mice; the first major human trial — Unity's, run by the field's own founders with serious money behind it — failed for knee osteoarthritis. Not a failure of chemistry or AI, but the same mouse-to-human gap from the last two chapters, in a third costume. Senolytics remain genuinely promising and genuinely unproven: real evidence that we can clear these cells in people, no proof yet that doing so cures anything at scale. So far we have been trying to *remove* the damaged cells, or *measure* their damage. The next chapter asks a more ambitious question — whether we can slow the *rate of damage itself*, by turning the ancient dials a cell uses to sense food. It is where the field's hardest fights live, and where its single most oversold molecule is waiting.

Chapter 4 · The Hunger Switch

SENSING SCARCITY FLIPS A CELL FROM GROWTH TO REPAIR — *the most reliable longevity lever we have, and the corner of the field most drowned in hype.*

> **Companion notebook · Chapter 4.** *The Goldilocks Dose: hormesis and the mTOR dial* — open it in Google Colab and build it yourself: Open in Colab →

THE WORM THAT DOUBLED

For most of the twentieth century, serious biologists did not study aging, because aging was not thought to be a *thing you could study*. It looked like entropy — the slow, unprogrammed wearing-down of a complicated machine, no more controllable than rust. There were no genes "for" aging any more than there were genes for a car rusting. To propose otherwise was to mark yourself as not quite respectable.

Then, in 1993, Cynthia Kenyon broke the field open with a worm.

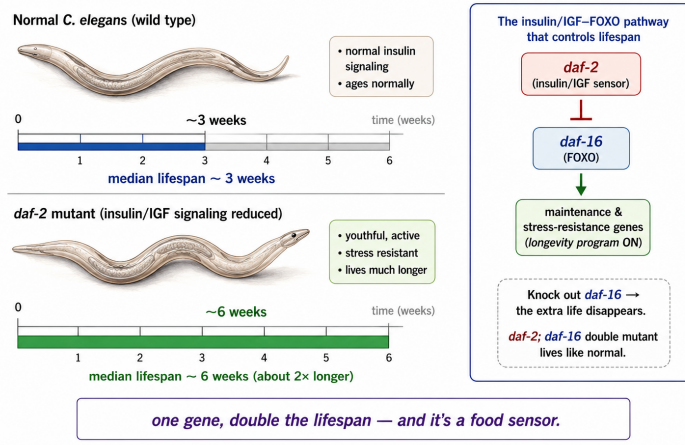
Working with *C. elegans* — a millimeter-long roundworm that lives a few weeks and has exactly 959 cells — Kenyon's lab found that mutating a single gene, called *daf-2*, **doubled the animal's lifespan.** Not by a few percent. Double. And the long-lived worms were not frail invalids dragging out a miserable extra span; they stayed youthful and vigorous, moving and feeding like much younger animals, simply for twice as long. A single gene. One switch.

It got better, and more revealing. The doubling was not automatic — it *required* a second gene, *daf-16*, working downstream. Knock out *daf-16*, and the *daf-2* longevity evaporated entirely; the worms aged and died on schedule. So this was not damage being passively avoided. It was an active program: *daf-2* normally holds *daf-16* in check, and when you release *daf-16*, it switches on a whole suite of protective, maintenance, and stress-resistance genes that keep the animal young. Aging, Kenyon had shown, is *regulated*. There are genetic controls that set how long an animal lives, and flipping a single one can rewrite the lifespan.

This is, without much competition, the most important single experiment in the molecular biology of aging, because it converted aging from an inevitability into a *mechanism* — and a mechanism can be studied, and perhaps adjusted. (Kenyon was not entirely first; in the 1980s Michael Klass and Tom Johnson had found worm mutants that lived longer. But the *magnitude* of the *daf-2* effect, and the clean demonstration that a regulated genetic pathway controlled it, is what

changed the field's mind.) And the moment anyone asked what *daf-2* actually *was*, the whole story snapped into focus around a single theme that organizes this chapter. *daf-2* is the worm's receptor for an insulin-and-IGF-like signal — a **nutrient sensor.** The switch that doubles a worm's life is the switch that tells its cells how much food is around.

Figure 4.1 — Aging is regulated.



THE GROWTH-VERSUS-MAINTENANCE SWITCH

Here is the deep idea the worm handed us, and it runs through every intervention in this chapter.

A cell, and an organism, must constantly make a choice about how to spend limited resources: pour them into **growth and reproduction**, or into **maintenance and repair**. When food is plentiful, the sensible bet is to grow and breed now. When food is scarce, the sensible bet is the opposite —

hunker down, fix what's broken, build stress resistance, and wait for better times to reproduce. Evolution wired cells to read the nutrient supply and shift the dial accordingly. And the discovery at the heart of modern aging biology is that **the maintenance setting is also, as a side effect, the longevity setting.** Tip a cell toward "scarcity — repair and endure," whether by actually restricting food or by drugging the sensors so they *think* food is scarce, and across yeast, worms, flies, and mice, the animal lives longer and healthier. This is the disposable-soma idea from the last chapter made actionable: when the body invests in maintenance over reproduction, it lasts longer.

A small number of conserved sensors do the reading, and they fall cleanly into the two camps. On the *food-is-plentiful, grow* side:

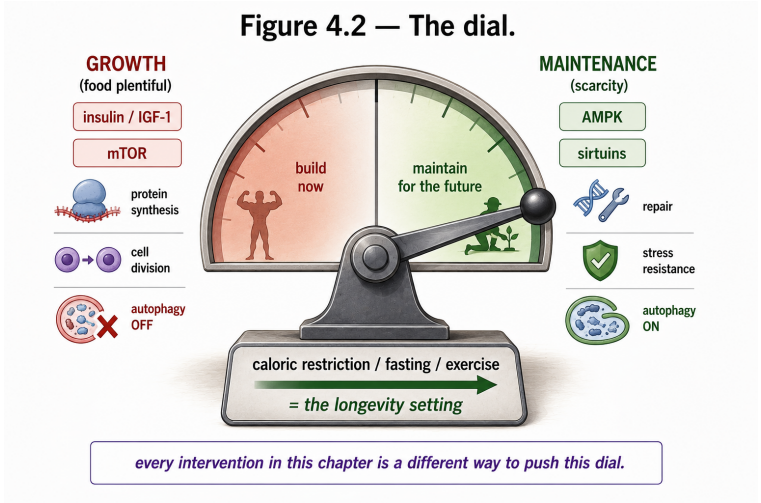
- **Insulin/IGF-1 signaling** — the pathway Kenyon's *daf-2* belongs to. It senses nutrients and growth cues and tells cells to grow and divide. Turn it *down* and lifespan goes *up*, all the way up the tree of life.
- **mTOR** (the "mechanistic target of rapamycin") — the master growth switch. When mTOR senses plentiful amino acids and growth signals, it drives the cell to build proteins and grow, and it *switches off* autophagy, the cell's recycling-and-cleanup system (which we'll meet properly in Chapter 6). Inhibit mTOR, and the

cell stops growing, starts cleaning house, and becomes more stress-resistant.

On the *food-is-scarce, maintain* side:

- **AMPK** — the cell's fuel gauge. It senses when energy is *low* and responds by switching on energy-producing, catabolic, maintenance programs and shutting down wasteful growth (it inhibits mTOR directly). It is the sensor that fires during fasting and exercise.
- **Sirtuins** — a family of enzymes that sense the cell's metabolic *state* and, when energy is scarce, promote stress resistance, DNA repair, and mitochondrial health. They are the protagonists — and the troublemakers — of the second half of this chapter.

These are not four separate stories. They are an interlocking network with a single output: the growth-versus-maintenance dial. Insulin/IGF and mTOR push toward growth; AMPK and sirtuins push toward maintenance and oppose mTOR. **Caloric restriction, fasting, and exercise all tip the entire network toward maintenance** — lowering insulin/IGF and mTOR, raising AMPK and sirtuin activity — and that coordinated tipping is, as far as we can tell, the core of why eating less makes animals live longer. The rest of this chapter is what happens when we try to flip that switch with drugs instead of hunger, and which of those attempts the evidence actually supports.



Rapamycin: the one that travels

If you want to know which longevity intervention to take most seriously, the honest answer is the one that keeps working no matter who runs the experiment — and that is rapamycin.

The molecule has a wonderful origin. In the 1960s, a soil sample from Rapa Nui — Easter Island, the one with the giant stone heads — yielded a compound made by a soil bacterium, an antifungal that was named *rapamycin* after the island. It turned out to be a potent immunosuppressant (it's used to keep transplant patients from rejecting organs), and the reason it works is that it shuts down mTOR — indeed, mTOR is *named after rapamycin*, the drug having been discovered before the protein it hits.

Then came the result that anchors the whole field. The U.S. National Institute on Aging runs an Interventions Testing Program designed to be hype-proof: candidate longevity drugs are tested in genetically diverse mice, simultaneously, at three independent laboratories, so that no single lab's quirks or wishful thinking can produce a false positive. In 2009 it reported that rapamycin extended the lifespan of these mice — and, astonishingly, it worked *even when started late in life*, at an age equivalent to a human in their sixties. A drug, begun in old age, reliably lengthened mammalian lifespan, in both sexes, reproduced across three labs. Nothing else in aging biology has a result that clean. Rapamycin is the closest thing the field has to a sure thing, and everything else gets measured against it.

That said — and this is the discipline the book runs on, applied *symmetrically*, not just to the things we want to dismiss — rapamycin is still not a *proven human* longevity drug. It is the best *animal* result we have, which is not the same thing. At the doses used in transplant patients it suppresses the immune system and can worsen blood-sugar control, which are real liabilities for a drug you'd give to healthy older people for decades. The field's current bet is that *intermittent*, lower-dose regimens, or close chemical cousins, can capture the longevity benefit while shedding the immunosuppression — and there are encouraging early human studies in which brief courses of rapamycin-like drugs actually *improved* immune function in the elderly. We are also,

right now, running the experiment in companion dogs (whose aging resembles ours far more than a lab mouse's does), which is exactly the right move: take the best mouse result and test it in a longer-lived, more human-like animal *before* believing it for people. Rapamycin clears a higher evidentiary bar than anything else in this chapter. It has not yet cleared the highest one.

Eat less, live longer — but compared to what?

Behind rapamycin sits the oldest longevity intervention of all, the one rapamycin partly mimics: simply eating less. **Caloric restriction** — reducing food intake (typically 20–40%) while keeping nutrition adequate — has extended lifespan in essentially every model organism it's been carefully tried in, from yeast to worms to flies to mice, where it can add 30–50% to lifespan. It was first shown in rats in 1935, and it remains the most robust environmental lever on aging known. Mechanistically it is no mystery: CR tips the whole nutrient-sensing network toward maintenance.

But the moment you ask whether it works in *primates* — in animals like us — you walk into the single best lesson in this book about why honest science is hard, and why "a study showed X" is never enough.

Two research teams set out, in the late 1980s, to test caloric restriction in rhesus monkeys — long-lived animals, decades-long studies, enormous expense and patience. One was at the

University of Wisconsin; the other at the National Institute on Aging. And when their results came in, they appeared to **flatly contradict each other.** Wisconsin (reporting in 2009) found that its calorie-restricted monkeys had less age-related disease and better survival — CR worked. The NIA (reporting in 2012) found that its calorie-restricted monkeys did *not* live significantly longer than controls — CR didn't, at least for survival. Same intervention. Same animal. Two rigorous, careful, decades-long studies. Opposite headlines. If you wanted a reason to distrust longevity science, here it was.

Except the contradiction was not a contradiction. When the two teams sat down together and reconciled their studies, the answer turned out to be almost entirely about a single, easily overlooked variable: **what the control monkeys ate.** At the NIA, the "control" monkeys were not gorging freely — they were fed a healthy, portion-controlled diet of natural ingredients, so they were already lean and already mildly restricted. At Wisconsin, the control monkeys ate as much as they wanted of a sugary, purified diet, and got fat. So the two studies were not asking the same question at all. Wisconsin compared *restricted* monkeys to *overfed* monkeys and found a big benefit. The NIA compared *restricted* monkeys to *already-sensibly-fed* monkeys and found a small one. The studies also differed in the age restriction started and in diet composition. Once you align them, the apparent war dissolves into a coherent and rather mundane conclusion: caloric restriction does benefit primate health, but how *much* benefit you see

depends enormously on what you're comparing against — and on a junk-food control diet, even modest restriction looks miraculous.

That is the lesson to carry out of this section and into every health claim you ever read: **"compared to what?"** is the question that breaks or makes a study. The monkey "contradiction" was not fraud, not sloppiness, not a failure of science. It was two excellent studies with different control groups, and the entire field had to learn to read them together. Most of the confusing, whiplash-inducing nutrition and longevity news you encounter has exactly this shape underneath it.

What about humans? We can't ethically run a lifelong severe-CR experiment, but the careful CALERIE trial put non-obese adults on sustained moderate restriction for two years and found real benefits — improved cardiovascular and metabolic markers, less inflammation, and (in a lovely connection to Chapter 2) an actual *slowing* of the Dunedin-PACE pace-of-aging clock. But the honest framing is sober: the effects were meaningful but modest, severe restriction is genuinely hard to sustain and carries its own costs (loss of bone and muscle, cold intolerance, and a quality-of-life toll most people will not accept), and we have no human evidence that CR extends maximum lifespan. CR is the real thing, and it is also not a magic wand — which is precisely why so much effort has gone into finding a *drug* that delivers its benefits without the hunger.

Deeper Water — the sensors read ratios, not amounts. There is an elegant design principle hiding in how these nutrient sensors actually work, and it's worth seeing. AMPK does not measure how much fuel the cell has in absolute terms; it reads a *ratio* — roughly the amount of AMP (the spent, low-energy form of the cell's energy currency) relative to ATP (the charged form). When the cell burns through fuel, ATP falls and AMP rises, so the AMP-to-ATP ratio climbs, and AMPK fires in proportion. Sirtuins, likewise, are tuned to the ratio of NAD^+ to NADH — a readout of how oxidized-versus-reduced, how fed-versus-fasted, the cell's metabolism currently is. Why ratios? Because a ratio is a far more robust signal than an absolute concentration. The total amount of any metabolite swings with cell volume, dilution, and the size of the metabolite pool; a *ratio* of two related molecules cancels all of that out and reports the cell's underlying *state* — its energy charge, its redox balance — independent of how big or dilute the cell happens to be. Evolution built fuel gauges that read a normalized quantity, for the same reason a good instrument reports a ratio rather than a raw reading: it's the quantity that actually means something. This matters for the

rest of the chapter, because it tells you what it would take to fool a sensor. To make a cell *think* it's fasting, you don't necessarily starve it — you shift one of these ratios. That is the entire premise of the molecules we turn to now, and it's also why they're so easy to oversell: moving the ratio is straightforward, but moving the ratio is not the same as making the animal younger. Hold that thought; it's about to become the whole problem. (One more human footnote: the worm's *daf-16* corresponds to the human FOXO genes, and a particular variant of *FOXO3* is among the most reliably replicated genetic associations with reaching extreme old age — the worm's longevity switch is, faintly, still visible in which humans become centenarians.)

The most oversold molecule

Now we arrive at the part of aging biology where the gap between what is claimed and what is known is the widest in the entire field — and the chapter has spent four sections building the tools you need to see it clearly. The science underneath is real and serious. The marketing on top of it is not. Both facts are true at once, and keeping them apart is the whole skill.

Start with the legitimate science. In the late 1990s, Leonard Guarente's lab at MIT found that a gene called *Sir2* extended lifespan in yeast, and that its activity depended on **NAD^+** — the same energy-state molecule from the box above. This was genuinely exciting: here was a longevity-promoting enzyme (a sirtuin) directly wired to the cell's metabolic state through NAD^+. The logic that followed was clean and seductive: sirtuins promote longevity; sirtuins run on NAD^+; NAD^+ levels *decline* with age; therefore, restore NAD^+ (or activate sirtuins directly) and you might slow aging. An entire industry grew out of that syllogism. But notice that it is a chain of "mights," and a chain is only as strong as its weakest link.

The first attempt to exploit it became a cautionary tale that everyone in this field should know by heart. In 2003, a compound was reported to activate the sirtuin SIRT1 and extend lifespan in yeast — and that compound was **resveratrol**, the molecule in red wine. The story wrote itself: the longevity secret of red wine, a natural compound that switches on your anti-aging enzymes. Follow-up work reported that resveratrol improved the health of mice on junk-food diets. A company, Sirtris, was founded to develop sirtuin activators, and in 2008 the pharmaceutical giant GSK bought it for roughly **720 million dollars.** Red-wine longevity was, briefly, one of the hottest ideas in medicine.

Then it came apart. Independent laboratories — including ones at major drug companies — could not reproduce the

central claim that resveratrol directly activates SIRT1. The apparent activation turned out to be, in significant part, **an artifact of the assay**: resveratrol only "activated" SIRT1 when the test used an unnatural, fluorescently-tagged substrate; with the enzyme's real targets, the effect largely vanished. The dramatic lifespan claims in animals did not robustly hold up either. In 2013, GSK shut Sirtris down. Resveratrol today is regarded, at best, as a weak and inconsistent intervention, and the episode stands as a textbook case of how an exciting result, a measurement artifact, and several hundred million commercial dollars can combine to outrun the evidence by years.

And here is why that history matters so much: **the same logic, and several of the same people, drove the next and current wave — NAD^+ precursors.** Since you can't easily get NAD^+ itself into cells, the strategy shifted to feeding the body molecules it converts into NAD^+: **nicotinamide mononucleotide (NMN)** and **nicotinamide riboside (NR).** These are the supplements at the center of today's longevity marketplace, and they are the chapter's titular oversold molecules. To evaluate them honestly, separate cleanly what is established from what is not:

What is solid. NAD^+ really does decline with age in many tissues. And supplementing NMN or NR really does **raise NAD^+ levels** in the blood and tissues of human beings — this part is well-replicated and not in dispute. In mice, NMN

and NR show a range of benefits across many studies — metabolic, vascular, and others. The biology is real.

What is not. There is, as of now, **no convincing evidence that raising NAD^+ in humans produces meaningful functional benefits** — that it slows aging, extends healthy lifespan, or prevents disease. The human trials are mostly small, short, and focused on biomarkers; where they look at actual function — insulin sensitivity, physical performance, strength — the results are mixed and, at best, modest. No human study has shown a lifespan or robust healthspan benefit from any NAD^+ precursor. This is precisely the trap from Chapters 1 and 2, returning in chemical form: NAD^+ is a **biomarker you can reliably move**, and moving a biomarker is not the same as moving aging. "It raises your NAD^+" is a thermometer claim, not a thermostat claim.

There is even a live scientific dispute about whether oral NMN does what the bottle implies. Charles Brenner — the biochemist who identified NR as an NAD^+ precursor in the first place, and a pointed critic of NAD^+ overclaiming — argues that NMN is not efficiently taken up *intact* into cells, and that swallowed NMN likely has to be converted to NR before it can get in, making the choice of NMN over NR partly a marketing distinction. The claim is contested; what matters for you is that even the *mechanism* by which the most popular version is supposed to work is unsettled among the experts.

A word on the conflicts of interest, because they are the most important media-literacy lesson in this chapter: **they run in every direction.** The most prominent NMN champion has co-founded companies built on it. Brenner, the NMN skeptic, has commercial ties to NR. The studies, the supplements, the podcasts, and the patents are tangled together so thoroughly that almost no prominent voice in NAD^+ biology is financially neutral. This does not mean everyone is lying — most of the underlying science is sincere. It means that when you read "new study shows NAD^+ booster reverses aging," your first question should always be *who benefits if I believe this*, and your second should be *what was the actual outcome measured* — and the honest answer to the second is almost always "a biomarker, in a small study, for a short time."

The tangle even reached the regulators. NMN was sold for years as a dietary supplement, until in 2022 the U.S. FDA reversed itself and concluded that NMN could **no longer be marketed as a supplement** — because it was simultaneously being developed as an investigational *drug* (a company co-founded by the field's most famous NMN advocate had begun the formal drug-development process for it). Under U.S. law, a molecule being investigated as a drug can be barred from sale as a supplement. The result has been a regulatory muddle, inconsistently enforced, with NMN still widely sold. But sit with the picture: the *same molecule*, framed in two incompatible ways at once — a harmless natural supplement you can buy by the jar, and a patented investigational phar-

maceutical — depending on which framing serves the seller. There is no sharper illustration in this book of how commercial incentives shape what reaches you and how it is described.

So, the honest bottom line on the NAD^+ story, holding both halves: this is serious biology — NAD^+ matters, it declines with age, precursors reliably restore it, and the mouse data are real and interesting. *And* the human evidence does not remotely support the claims made for these molecules, the mechanism of the most popular one is disputed, the entire field is saturated with financial conflict, and the gap between the marketing and the data is the widest in all of longevity science. NMN may yet prove useful; the point is that **right now, nobody knows that it does, and a great many people are paid to tell you they do.**

> **Deeper Water — what it would actually take to believe a NAD^+ booster.** It is worth spelling out exactly what evidence would move a NAD^+ precursor from "moves a biomarker" to "slows aging," because the same template applies to every supplement in this book. First, you would need a *large, randomized, placebo-controlled* trial — not nine people, not an open-label study where everyone knows what they're taking, but enough subjects, randomized, with a real placebo, to rule out chance and wishful thinking. Second, the

outcome measured would have to be something that *matters* — not NAD^+ levels, not even a clock reading (recall from Chapter 2 that clocks are unvalidated surrogates), but a hard endpoint: fewer diseases, better physical function sustained over years, or longer life. Third, and most demanding, you'd want the result *replicated by independent groups with no financial stake* — because the resveratrol collapse taught us that a single lab's positive result, especially a commercially interested one, can be an artifact. By that standard, rapamycin (multi-lab, placebo-controlled animal data, started late in life, now being tested in dogs and humans) is *years* ahead of NMN, and the difference is not subtle. The reason to insist on this standard is not pedantry. It is that the history of this exact field — antioxidants in the next chapter, resveratrol in this one — is a graveyard of compounds that moved a biomarker, generated excitement and revenue, and then failed the only test that counted.

A MEASURED WORD ON METFORMIN

One more nutrient-sensing drug deserves mention, because it sits in an interesting middle and is the field's most important regulatory test case. **Metformin** is a cheap, dec-

ades-old, widely-used diabetes drug that, among its several effects, activates AMPK — the fasting/maintenance sensor. Some observational studies hint that people taking metformin have lower mortality and cancer rates; it extends lifespan in some animal models (though not all); and it is plausible as a mild CR-mimetic. The veteran researcher Nir Barzilai has long championed a trial called **TAME** (Targeting Aging with Metformin), designed less to prove metformin is a wonder drug than to do something more strategic: to test whether a drug can delay the onset of *multiple* age-related diseases at once, and in doing so to persuade regulators to treat "aging" itself as a condition worth targeting — which would change the economics of the entire field.

The honest status is mixed and modest. TAME has struggled for years to secure funding and has not yet delivered a definitive verdict. The evidence that metformin helps *non-diabetics* live longer is genuinely uncertain — and there is even a finding that metformin may *blunt* some of the benefits of exercise, a reminder that pushing on a nutrient sensor can have costs as well as benefits. Metformin is cheap, safe, and important as a proof-of-concept for regulating aging as a treatable process. It is not, on current evidence, a demonstrated longevity drug for healthy people. File it under "promising and unproven," next to almost everything else here that isn't rapamycin.

Where the machines come in

Nutrient sensing is, in one sense, the most AI-ready of all the hallmarks, because the targets are well-defined molecular nodes — mTOR, AMPK, the sirtuins, the insulin/IGF pathway — and "find or design a molecule that selectively tunes a known target" is exactly the problem the newest computational tools are built for. Companies like Insilico Medicine use AI to hunt for novel longevity-relevant targets in vast biological datasets and then to *generate* candidate molecules to hit them; Insilico's pipeline has carried an AI-discovered drug for a fibrotic lung disease — an aging-related condition — into human trials, a genuine milestone for the field of computational drug discovery (even if the indication is a specific disease rather than "aging"). The same approach could, in principle, yield better CR-mimetics or rapamycin-like drugs with the longevity benefit and fewer of the side effects — the selectivity problem from the last chapter, attacked with design.

But weigh that promise against everything this chapter has shown, because the binding constraints are not the ones AI removes. AI can accelerate the discovery of molecules that hit these nodes; it cannot tell you whether tuning the node *in a human* extends healthy life, and it cannot manufacture the large, long, placebo-controlled, independently-replicated human trials that are the only thing that ever settles the question. There is also a subtler danger specific to this

overhyped corner of biology: AI is very good at finding *patterns and candidates*, and in a field where moving a biomarker is easy and proving a benefit is hard, a clever algorithm trained to optimize a surrogate — lower the clock, raise the NAD^+, shift the ratio — will faithfully hand you compounds that are excellent at *gaming the surrogate*, which is not the same as compounds that make you younger. The lesson of the resveratrol assay artifact was that a measurement can be optimized into meaninglessness; an AI pointed at the wrong target will optimize it faster. The computation is not the constraint. The constraint is what the last four chapters keep insisting on: it has to work in a human, and a biomarker is not a human outcome.

✦ ✦ ✦

Carry these three things

1. **Longevity is, in part, a resource-allocation decision, and you can bias it toward maintenance.** A handful of conserved sensors — insulin/IGF and mTOR on the "grow" side, AMPK and sirtuins on the "maintain" side — read how much food is around and set a growth-versus-repair dial. Caloric restriction, fasting, and exercise tip that dial toward maintenance, and the maintenance setting is also the longevity setting, across the whole tree of life. Kenyon's single-gene

worm proved aging is a *regulated mechanism*, not just entropy — and the mechanism is a food sensor.

2. **Rank interventions by the quality of their evidence, not the volume of their marketing — and the two are nearly inverted.** Rapamycin sits at the strong end: a multi-lab, reproducible, mammalian result that worked even when started in old age, now being tested in dogs and humans — and *even rapamycin* is not yet a proven human longevity drug, which is the bar everything must eventually clear. NAD^+ boosters like NMN sit at the weak end: real biology underneath, but human benefit unproven, mechanism disputed, and the field drowning in financial conflict — the most oversold molecules in aging, a chemical replay of the resveratrol collapse that cost GSK $720 million. "It raises your NAD^+" is a thermometer claim, not a thermostat one.
3. **"Compared to what?" is the question that breaks or makes a study.** The two rhesus-monkey caloric-restriction studies appeared to contradict each other for years — and the difference came down almost entirely to what the *control* monkeys were fed. Not fraud, not sloppiness: two careful studies with different comparison groups, which the whole field had to learn to read together. Most whiplash-inducing health news has this exact shape. So far we've tried to slow aging by removing damaged cells, measuring damage, and dialing

down the rate of damage. The next chapter takes on the most famous theory of *why* the damage happens in the first place — the idea that we rust from the inside — and shows you one of the most expensive, instructive ways a beautiful theory can simply be wrong.

Chapter 5 · The Beautiful Wrong Idea

THE FREE-RADICAL THEORY OF AGING WAS ELEGANT, INTUITIVE, *and supported by everything except the experiment — and the experiment showed that a little of the "poison" is exactly what keeps you healthy.*

> **Companion notebook · Chapter 5.** No separate notebook for this chapter — the inverted-U "Goldilocks dose" you build for Chapter 4 is the same curve at work here: a little ROS signals and protects, while too much harms: Open the Chapter 4 notebook →

WE RUST FROM THE INSIDE

Of all the theories of why we age, one stands above the rest for sheer intuitive force. Iron left in the rain rusts; cut an apple and it browns; a fire consumes its fuel and leaves ash. All of these are *oxidation* — oxygen reacting with matter and degrading it. And we, who breathe oxygen and burn fuel to live, are slowly oxidizing too. We rust from the inside. Aging, on this view, is simply the accumulated chemical wear of a

lifetime of breathing — the unavoidable price of using oxygen to stay alive.

This is the **free-radical theory of aging**, proposed in 1956 by a chemist named Denham Harman, and for half a century it was the dominant theory in the field. It is easy to see why. It is mechanistically precise, it connects aging to the most basic fact of our metabolism, it matches the observation that oxidative damage really does accumulate in old tissues, and — crucially — it makes a clear, hopeful, actionable prediction: if oxidation causes aging, then *antioxidants* — molecules that neutralize the reactive byproducts of metabolism — should slow aging and prevent the diseases that come with it. Eat your antioxidants, mop up the rust, live longer. An entire multi-billion-dollar supplement industry was built on that single inference, and it remains, to this day, the most widely believed idea about how to slow aging.

It is also, in its simple and intuitive form, *wrong* — and the story of how we learned that it was wrong is the most important lesson in this book about how to tell a beautiful idea from a true one. (A small, telling detail: Harman, the father of the theory that aging is unavoidable oxidative decay, lived to 98.) This chapter is about the cell's power plants, the reactive molecules they make, the theory that those molecules age us, and the long, expensive demonstration that the theory had it almost exactly backwards.

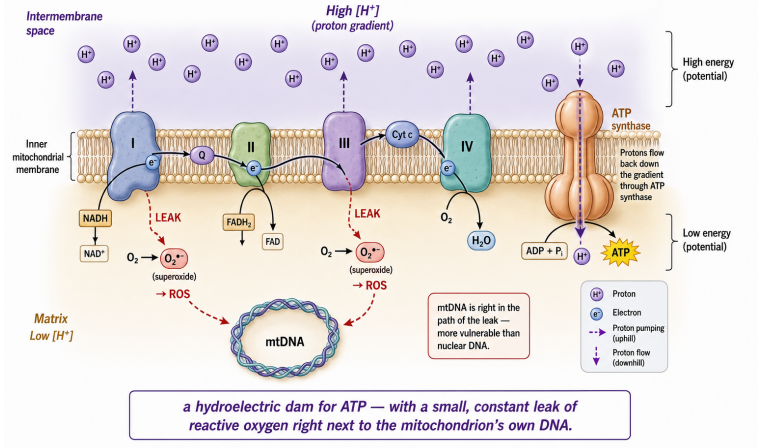
What a mitochondrion actually does

To understand the theory, and its undoing, you need to know where the "rust" comes from, and that means looking inside the mitochondrion — but we have to be precise, because "the cell's power plant" is the kind of phrase that substitutes for understanding rather than building it.

Mitochondria are small compartments inside your cells whose job is to convert the energy in food into ATP, the molecule the cell spends to do everything it does. The conversion works like a tiny hydroelectric dam, and the analogy is exact rather than poetic. Electrons stripped from food are passed down a chain of four protein complexes embedded in the mitochondrion's inner membrane — the **electron transport chain.** As each electron hops down the chain, the energy released is used to pump protons (hydrogen ions) across the membrane, building up a reservoir of protons on one side, like water held behind a dam. That stored-up gradient is then released through a turbine — a remarkable rotating enzyme called ATP synthase — and the flow drives the turbine to manufacture ATP. At the very end of the chain, the spent electrons are handed off to oxygen, which combines with them (and with protons) to form water. This is *why you breathe*: oxygen is the final electron acceptor at the bottom of the chain, and without it the whole electron flow backs up and stops. Most of the ATP that powers your body is made exactly this way, every second.

Now the catch that launched a thousand supplement bottles. The electron chain is fast and busy and not perfectly tidy. A small fraction of the electrons moving down it — far less than a percent — escape prematurely and react directly with oxygen, producing not water but a damaged, reactive form of oxygen: **superoxide**, the first of a family of molecules collectively called **reactive oxygen species**, or ROS. Superoxide and its chemical relatives (hydrogen peroxide, and the especially vicious hydroxyl radical) are reactive precisely because they are unstable — they will grab electrons from whatever is nearby, and "whatever is nearby" includes the cell's proteins, the fatty membranes, and DNA. ROS are, in short, a byproduct of breathing: an unavoidable trickle of chemically corrosive molecules, produced continuously as the price of making energy from oxygen. They are the rust.

Figure 5.1 — The power plant and its leak.



a hydroelectric dam for ATP — with a small, constant leak of reactive oxygen right next to the mitochondrion's own DNA.

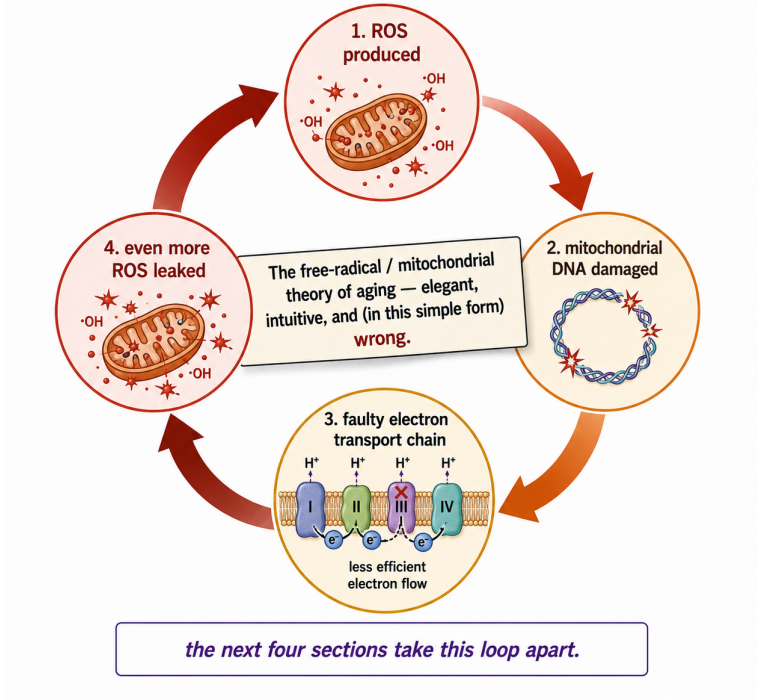
Harman's insight, refined in 1972 into what's called the *mitochondrial* theory of aging, took that leak and built it into a self-reinforcing catastrophe — and the elegance of the construction is exactly why it held the field for fifty years.

The argument runs like this. Mitochondria are the main source of ROS. They also carry their *own* small genome — a circular loop of DNA, separate from the DNA in the cell's nucleus, that encodes several of the essential parts of the electron transport chain. And this mitochondrial DNA sits right next to the leak, in the most ROS-drenched spot in the cell, with what was long thought to be weaker protection and repair than nuclear DNA enjoys. So picture the loop closing: ROS damage the mitochondrial DNA; damaged DNA builds a faulty, sputtering electron transport chain; a faulty chain leaks *even more* electrons as ROS; more ROS do more damage. A vicious cycle — each turn worse than the last — that grinds the mitochondria, and with them the cell, and with them you, slowly into decline. Aging as an error catastrophe, spiraling out from the cell's own power supply.

It explained things, too, which is what made it irresistible. It fit the old "rate of living" observation that animals with faster metabolisms tend to live shorter lives — more metabolism, more ROS, faster aging. It fit the rising tide of oxidative damage markers found in aged tissue. And when researchers discovered, in 1969, an enzyme called superoxide dismutase

whose *entire job* is to find and destroy superoxide, it seemed like confirmation from the cell itself: evolution had built dedicated machinery to defend against ROS, which surely meant ROS were a serious threat worth defending against. The theory had a mechanism, observational support, evolutionary corroboration, and — above all — a clean prediction that anyone could test and anyone could act on. *Take antioxidants and you will age more slowly.*

Figure 5.2 — The theory, drawn as a loop.



The graveyard

A clear prediction is a gift, because it can be tested — and the free-radical theory's prediction was tested, in some of the largest and most expensive clinical trials ever run. The results form one of the most instructive graveyards in all of medicine.

The case for antioxidants did not start with supplements; it started with *food.* People who ate diets rich in antioxidant-laden fruits and vegetables had, observably, lower rates of cancer and heart disease. That correlation was real and robust, and it pointed straight at the supplement aisle: if the antioxidants in vegetables protect you, then antioxidant pills should bottle that protection. So the hypothesis went to the randomized controlled trial — the one method that can tell a real cause from a misleading correlation — and the trials gave back an answer almost nobody expected.

The most shocking result involved **beta-carotene**, the antioxidant that gives carrots their color, tested as a means of preventing lung cancer in people at high risk. Two enormous trials — one in tens of thousands of Finnish male smokers, another in thousands of American smokers and asbestos-exposed workers — gave half the participants beta-carotene and half a placebo, and waited. The antioxidant did not prevent lung cancer. It **caused more of it.** The smokers taking beta-carotene developed lung cancer at *higher* rates, and died at higher rates, than those taking the placebo. The

American trial was halted ahead of schedule, because it had become unethical to keep giving people a supplement that was making them worse. Read that again, because it is the single most useful data point in this book: an antioxidant, given specifically to *prevent* cancer, on the strength of a beautiful and intuitive theory, *increased* cancer and death.

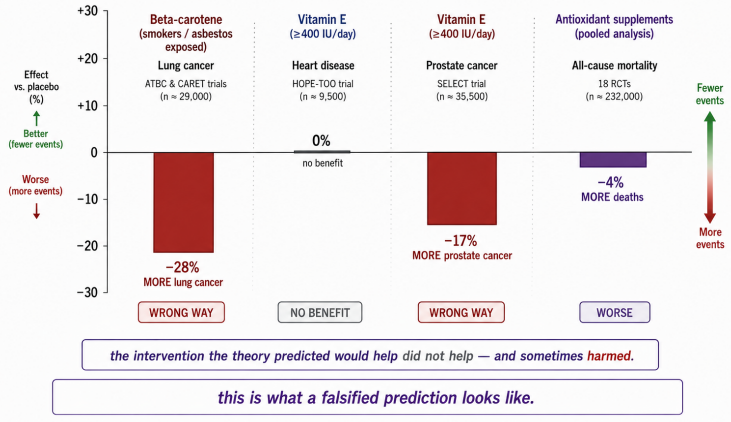
It was not a fluke. **Vitamin E**, the great hope for preventing heart disease, failed trial after trial to prevent it — and a major prostate-cancer-prevention trial found that vitamin E supplementation *increased* the risk of prostate cancer. Pooled analyses combining dozens of antioxidant-supplement trials, covering hundreds of thousands of people, reached a verdict that should be engraved over the supplement aisle: antioxidant supplements do not reduce mortality, and several of them — beta-carotene, vitamin E, high-dose vitamin A — are associated with a small but real *increase* in the risk of dying. The most intuitive theory of aging ever proposed made a clear prediction, the prediction was tested as rigorously as medicine knows how to test anything, and the prediction was not merely unconfirmed. It was inverted.

How could the food correlation have been so misleading? The answer is the lesson of the last chapter — *compared to what?* — wearing a slightly different coat. People who eat lots of vegetables are not a random sample of humanity; they tend, on average, to exercise more, smoke less, be wealthier, see doctors more, and do a hundred other healthful things. When you observe that vegetable-eaters get less cancer, you

are seeing the combined effect of *being the kind of person who eats vegetables*, and the beta-carotene is just one strand tangled up with all the rest. The randomized trial cuts that tangle: by assigning the supplement at random, it isolates the *one* thing — the beta-carotene itself — from everything that usually travels with it. And once isolated, the beta-carotene's true effect turned out to be nothing, or worse than nothing. The vegetables may well be good for you. The beta-carotene pulled out of them and packed into a pill is a different thing entirely.

Figure 5.4 — The graveyard.

Major randomized trials of antioxidant supplements vs. placebo



Deeper Water — why observational studies fooled everyone, and why the trial is king. The antioxidant story is the clearest classroom example of *confounding* and why it makes observational data treacherous for causal questions. A

confounder is a hidden third factor that influences both the thing you're measuring and the outcome, manufacturing a correlation that isn't causal. Vegetable consumption and low cancer risk are correlated, but a cluster of confounders — wealth, education, exercise, not smoking, general health-consciousness, what researchers call *healthy-user bias* — drives both, so the correlation survives even if the vegetables themselves do nothing. You cannot fully fix this by "adjusting" for the confounders you can think of, because there are always confounders you didn't think of, or can't measure well. The randomized controlled trial solves the problem at a stroke and is worth understanding *why* it works: by assigning the intervention with a coin flip, randomization makes the treatment and placebo groups statistically identical in *every* respect — the confounders you know about, the ones you don't, the ones nobody has ever named — *except* the one thing you're testing. Any difference in outcome must then be due to that one thing. This is why a single well-run RCT can overturn decades of consistent observational data, and why it is the highest standard of evidence in medicine. It is also why the trials being stopped *early* matters: there are formal rules (pre-specified stopping

boundaries, monitored by an independent board) that halt a trial when one group is clearly being harmed, because continuing to give people a treatment you now believe is hurting them violates the ethics that make human experiments permissible at all. The beta-carotene trial was stopped not because the science failed but because the science *worked* — it detected harm, and the rules required acting on it.

ROS ARE NOT JUST POISON

So the antioxidants failed, and several of them harmed. *Why?* If ROS are corrosive byproducts, mopping them up should at worst do nothing. How could reducing the rust make things *worse?* The answer is the conceptual correction at the heart of the modern view, and it overturns the theory's founding assumption: **ROS are not only poison. They are also signals.**

The simple theory treated reactive oxygen as pure damage — chemical shrapnel, all cost and no use. But cells, it turns out, *deliberately produce ROS* — especially hydrogen peroxide — as messengers, and depend on them for a startling range of essential jobs. Controlled bursts of ROS regulate gene expression, trigger the cell's own adaptive stress responses, help control how cells grow and specialize, and are the actual *weapon* the immune system uses: when a white blood cell engulfs a bacterium, it kills it by blasting it with a deliberate

burst of reactive oxygen, the so-called respiratory burst. ROS sit at the center of how the cell senses its own condition and responds to it. At low and moderate levels they are not damage at all — they are *information*, a chemical language the cell uses to tell itself what's happening and what to do about it.

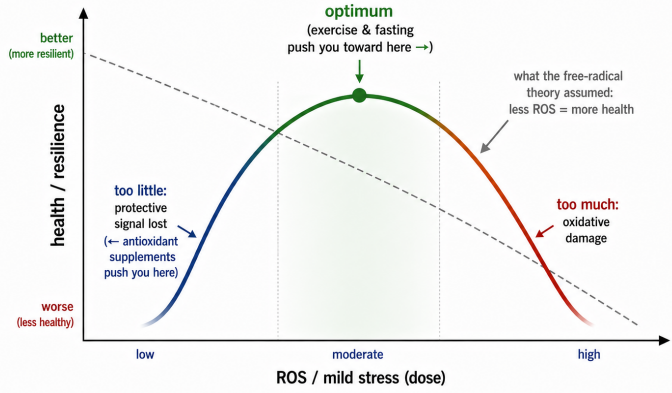
Once you see that, the harm from antioxidants stops being a paradox and becomes almost obvious. When you flood your body with antioxidants, you are not just mopping up corrosive shrapnel. You are also *garbling the cell's signaling* — muffling the very messages it uses to mount its defenses and adapt to stress. You are turning down a signal the cell needs, not just absorbing a poison it fears. An antioxidant pill cannot tell the difference between a damaging ROS and a signaling one; it suppresses both. And suppressing the signal, it turns out, can be considerably worse for you than tolerating the trickle of damage.

A LITTLE POISON, PLEASE

This brings us to the resolution, and it is the same idea that ran quietly through the last chapter, now made central: **hormesis** — the principle that a *small* dose of something stressful or even toxic can be beneficial, because it provokes the body to overcompensate with a protective response, while a large dose is harmful. The dose makes the poison, and at the bottom of the dose scale, the poison is medicine.

Applied to mitochondria, this is called **mitohormesis**, and it works exactly as you'd now expect. A *little* mitochondrial ROS is good for you — not in spite of being a stressor but *because* it is one. A mild pulse of reactive oxygen acts as a warning shot that triggers the cell to strengthen itself: to ramp up its own internal antioxidant defenses, build more and better mitochondria, and raise its general stress resistance. The cell, briefly stressed, overcorrects, and ends up tougher than before. The relationship between ROS and health is therefore not the straight downhill line the free-radical theory assumed — *less ROS, more health* — but an **inverted U**: too little ROS and you lose the protective signal, too much and you suffer real damage, with an optimum somewhere in the middle. The theory's fatal error was assuming a monotonic line where biology had built a curve.

Figure 5.3 — The curve the theory got wrong.



the whole chapter in one picture — biology built a curve where the theory drew a line.

The most beautiful evidence for all of this is a single, clean human experiment that resolves a paradox you may have already noticed. **Exercise increases ROS production** — when you work your muscles hard, their mitochondria burn faster and leak more reactive oxygen. By the free-radical theory, exercise should therefore *accelerate* aging. Instead, exercise is about the most reliably health-promoting thing a human being can do. Mitohormesis explains why: the ROS burst from exercise is the *signal* that triggers all of exercise's adaptive benefits — the improved insulin sensitivity, the strengthened defenses, the metabolic tune-up. And here is the experiment that proved it: researchers gave one group of people exercising regularly a daily dose of antioxidant vitamins (C and E), and another group exercising identically a placebo. The placebo group got the expected healthful adaptations to exercise. The antioxidant group largely **did not** — the supplements *blunted* the key metabolic adaptations to exercise, including the improvements in insulin sensitivity and the cell's own antioxidant defenses. By mopping up the ROS, the antioxidants had muffled the signal, and with the signal damped, the body's adaptation never fully took hold. You could exercise hard and, by taking your vitamins, get much of the benefit canceled out.

Sit with what that means. It is the free-radical theory's prediction falsified at the level of mechanism, in a single elegant study: the thing the theory said was pure damage turned out to be the very signal that makes the single health-

iest behavior we know actually work. (It is also, you'll notice, the exact shape of the metformin-and-exercise finding from the last chapter — a chemical that dampens a stress signal can dampen the adaptation that signal was supposed to trigger.)

When the theory was put to the genetic test

The clinical trials buried the *prediction*; the laboratory then buried the *mechanism*, by testing the vicious cycle directly. If oxidative damage drives aging, several things should follow — and when researchers checked, they mostly didn't.

If more antioxidant defense slows aging, then engineering mice to overproduce antioxidant enzymes should make them live longer. By and large, **it didn't** — most such mice lived perfectly ordinary lifespans, with at best a modest exception or two. If oxidative damage levels set lifespan, then the longest-lived animals should be the least oxidatively damaged. But the **naked mole-rat** — a hairless, subterranean rodent that lives over thirty years, roughly ten times longer than a mouse of similar size, and is strikingly resistant to cancer — carries *high* levels of oxidative damage from a young age. By the free-radical theory it should age in fast-forward; instead it is the longest-lived rodent known. It is a living refutation of the idea that your oxidative-damage burden determines how long you live. (Birds, too, with their

furnace-hot metabolisms and long lifespans, quietly defy the rate-of-living logic.)

And the most direct test of all: researchers built a "mutator" mouse, engineered so that its mitochondrial DNA accumulates mutations far faster than normal. By the theory, this should trigger the ROS vicious cycle and produce accelerated aging — and indeed these mice *do* age prematurely. But when the experimenters looked for the supposed engine of that aging, the runaway ROS, **it wasn't there.** The mice were not drowning in extra reactive oxygen. Their premature aging was driven by the mitochondrial mutations crippling the electron transport chain *directly* — by their power plants simply failing — not by a self-amplifying spiral of oxidative damage. Mitochondrial DNA damage mattered. The specific *mechanism* the theory proposed for how it mattered did not.

> **Deeper Water — what is actually true about mitochondria and aging.** It would be a serious mistake to walk away from this chapter thinking "mitochondria don't matter" or "ROS are harmless." Neither is true, and the honest modern picture is more interesting than the theory it replaced. Mitochondrial dysfunction *is* a genuine hallmark of aging: as we get older, our mitochondria really do work less well — the electron transport chain becomes less efficient, mitochondrial DNA accumulates mutations and deletions,

and, critically, the cell's *quality-control* systems falter. That last part is where the real action turns out to be. Cells continuously police their mitochondria: they tag damaged ones for destruction and digestion (a selective cleanup process called **mitophagy**), and they build fresh ones (**biogenesis**), and they constantly fuse and split them to share components and isolate damage (**dynamics**). With age, this housekeeping degrades — damaged mitochondria are cleared too slowly and replaced too rarely, so dysfunctional ones accumulate. The modern view, then, is not "oxidative damage piles up and ages us" but something subtler: mitochondria decline in *function, signaling, and turnover*, ROS act as both damage *and* essential signal with an optimal intermediate level, and the failure is one of *maintenance and quality control* more than of simple corrosion. The mitochondrion remains central to aging. Harman was right that it matters and right that ROS are real; he was wrong about the shape of the relationship — a curve, not a line — and wrong that mopping up ROS would help. Which is exactly why the therapies that may actually work, as we'll see, target *quality control* — clearing and replacing bad mito-

chondria — rather than blanket-suppressing the chemistry.

Where the machines come in

The collapse of the antioxidant idea reshaped the search for mitochondrial therapies, and the new direction is mechanistically smarter — and a good test of whether we've actually learned the chapter's lesson. Instead of flooding the whole body with antioxidants, the targeted approaches try to be *selective*: antioxidants engineered to home specifically into mitochondria rather than wash through every cell (so they don't garble signaling everywhere), and — more promisingly — drugs that boost the *quality-control* systems rather than the chemistry. The standout example is a compound called urolithin A, a molecule the gut produces from certain fruits, which *induces mitophagy* — prompting cells to clear out their damaged mitochondria — and which, in human trials, has shown modest improvements in muscle function. The whole strategy has shifted from "mop up the rust" to "take out the broken machines and build new ones," which is the right lesson drawn from everything above.

Artificial intelligence enters here in the now-familiar role: screening for and designing molecules that enhance mitophagy or improve mitochondrial function, mining large datasets to work out which mitochondrial-DNA mutations and which mixtures of healthy-and-damaged mitochondria actu-

ally matter for a cell's fate. These are real and useful applications. But this chapter, more than any other, should make you cautious about the *style* of confidence that AI can lend to a hypothesis. The free-radical theory was the most computationally and mechanistically respectable idea imaginable — it had elegant chemistry, decades of supporting data, and an obvious prediction — and it was wrong in a way that no amount of cleverness about the mechanism could have caught. Only the randomized trial caught it. AI can help find a mitophagy-inducing molecule faster than ever before; it cannot tell you whether clearing those mitochondria will make a human being healthier, and it cannot substitute for the large, randomized, placebo-controlled trial that is the only thing that ever could. In a field whose most beautiful and best-supported theory turned out to be inverted, the constraint was never the elegance of the idea or the speed of the search. It was, and remains, the experiment.

✦ ✦ ✦

Carry these three things

1. **A mechanism that sounds right — even with correlations and an industry behind it — is not proof. Only the experiment is.** The free-radical theory of aging had everything: elegant chemistry, observational support from vegetable-rich diets, dedicated cellular defenses confirming ROS matter, fifty years of domin-

ance, and a clear prediction. The prediction was tested in massive randomized trials, and it failed — beta-carotene given to high-risk smokers to *prevent* lung cancer *increased* it, and pooled antioxidant trials showed no benefit and some harm. The vegetables-correlation was confounded (vegetable-eaters differ in a hundred ways); the randomized trial cut the tangle and revealed the truth. This is the single most useful habit this book can give you.

2. **The dose makes the poison — and biology runs on curves, not lines.** ROS are not only damage; they are essential *signals* the cell uses to defend and adapt itself, and the relationship between ROS and health is an inverted U with an optimum in the middle, not a downhill slide. That is why flooding the body with antioxidants can *harm* — it garbles the signal — and why exercise, which *raises* ROS, is profoundly healthy: the ROS burst is the signal that triggers the benefit. Antioxidant supplements were shown to *blunt* key metabolic gains from exercise, by muffling that very signal. "More is not always better" is not a platitude here; it is the literal shape of the data.
3. **Mitochondria still matter for aging — just not the way the theory said.** The simple "oxidative damage accumulates and ages us" vicious cycle failed every direct test: antioxidant-enzyme mice didn't live longer, the naked mole-rat lives thirty years steeped in oxidative

damage, and the mtDNA-mutator mouse aged prematurely *without* the predicted ROS spike. The honest modern picture is about declining mitochondrial *function, signaling, and quality control* — which is why the therapies worth watching aim to clear and replace damaged mitochondria rather than blanket-suppress the chemistry. That idea — that aging is partly a failure to take out the cellular trash — is exactly where the next chapter begins: with the proteins that misfold and clump together, the cleanup crews that are supposed to remove them, and what happens, in the diseases we fear most, when the crews fall behind.

Chapter 6 · Folding and Falling Apart

PROTEINS MUST FOLD INTO EXACT SHAPES TO WORK, AND THE *cell fights a constant, losing war to keep them folded — the war behind the diseases we fear most. The most celebrated AI in biology cracked how proteins fold; misfolding is another matter.*

> **Companion notebook · Chapter 6.** *Why Proteins Don't Fold by Searching (the folding funnel)* — open it in Google Colab and build it yourself: Open in Colab →

THE SHAPE IS THE FUNCTION

Almost everything your body does, it does with proteins. They are the enzymes that run your chemistry, the motors that move your muscles, the receptors that hear signals, the antibodies that defend you, the scaffolds that hold cells in shape. And every one of them is built the same way: as a long chain of amino acids, strung together in an order written by a gene, that then folds up into a precise three-dimensional shape. The shape is not decoration. **The shape *is* the**

function. An enzyme works because its folded form creates a pocket shaped exactly to grip its target; an antibody binds because its folded loops match the thing it's hunting. Unfold the protein — or fold it wrong — and the function is simply gone, the way a key filed to the wrong shape opens nothing.

This poses a problem the cell must solve billions of times a second, and a problem so hard that it took the field of biology fifty years and, eventually, an artificial intelligence to crack even half of it. How does a floppy chain of amino acids find its one correct shape, reliably, fast, in the chaos of a living cell — and what happens, in aging and in the diseases we fear most, when it fails? This chapter is about folding, about the cellular crews whose job is to keep proteins folded and to dispose of the ones that aren't, about what goes wrong when those crews fall behind, and about the most famous AI achievement in the history of biology — which, it turns out, solved precisely the part of this problem that isn't killing us.

The folding problem, and why it's hard

Begin with a puzzle that stopped people cold in 1969, named for the scientist who posed it: **Levinthal's paradox.**

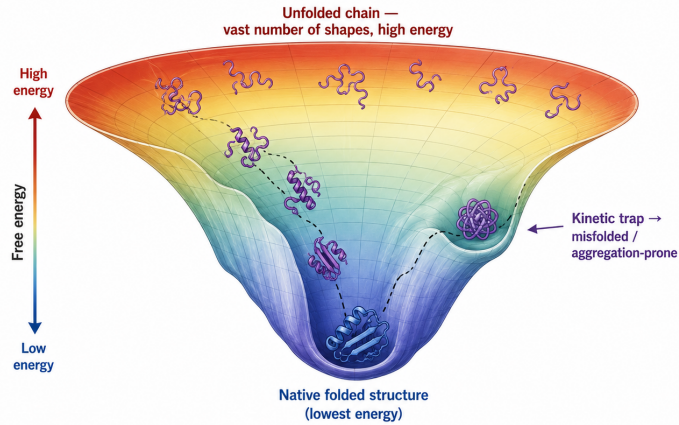
Consider a smallish protein, a chain of a hundred amino acids. At each link, the chain can bend into roughly a few distinct angles. With a hundred links and, say, three choices each, the number of possible shapes the chain could take is about 3^{100} — which is around 10^{48}, a one followed by

forty-eight zeros. Now suppose the protein folded by *trying* shapes, flipping from one to the next as fast as molecules can move, about once every 10^{-13} seconds. To sample all 10^{48} shapes would take roughly 10^{35} seconds. The universe is about 4×10^{17} seconds old. A protein folding by random search would take *vastly* longer than the age of the universe to find its shape — and yet real proteins fold in thousandths of a second to seconds. They cannot possibly be searching at random. Something is guiding them.

The resolution, worked out over the following decades, is one of the loveliest ideas in molecular biophysics: **folding is not a search, it is a fall.** Picture the protein's possible shapes as a landscape of hills and valleys, where height represents energy — and picture that landscape shaped not like a flat plain riddled with pits, but like a **funnel**, sloping down on all sides toward a single deep well at the bottom that is the correct, native shape. The unfolded chain starts at the wide rim, and at every moment the local slope nudges it downhill, toward more stable, more correct partial structures. It doesn't try every shape; it rolls down the funnel, the way water finds the drain without first visiting every point in the sink. This is why folding is fast despite the astronomical number of possibilities: the energy landscape itself does the work of narrowing the search. And — a second deep idea, established by Christian Anfinsen, who showed that an unfolded protein in a test tube can refold to its exact working shape entirely on its own — **the shape of that funnel, and the location of the**

well at the bottom, is encoded in the amino-acid sequence itself. The sequence determines the structure. The instructions for the final shape are written into the chain.

Figure 6.1 — Folding is a fall, not a search.



the protein doesn't try every shape —
the energy landscape funnels it to the right one.
The instructions for the funnel are in the sequence.

So in principle, sequence determines structure and the funnel makes folding fast. In practice, inside a real cell, folding is perilous — and the peril is the seed of everything that follows. The cytoplasm is not a clean test tube; it is one of the most crowded environments in nature, packed so densely with other molecules that a newly made protein chain emerges into a jostling crowd, half-folded and vulnerable, with no room to spare. Heat, mutations, and ordinary wear can knock a protein out of its well. And here is the crucial danger: a misfolded protein usually exposes patches

that are meant to be buried in its core — oily, water-avoiding "sticky" surfaces — and when two misfolded proteins drift together, those exposed sticky patches latch onto each other. They clump. **Aggregation** — misfolded proteins gluing themselves into clumps — is the central catastrophe of this chapter, because a clump is far harder to fix or remove than a single bad protein, and because the clumps are toxic. Folding is a fall, and most of the time it lands true. But it doesn't always, and the cell needs a way to deal with what happens when it doesn't.

The cleanup crews

The cell does not leave folding to chance and luck. It maintains an elaborate, layered system whose entire job is keeping the proteome correctly folded, functional, and clear of junk — a system called **proteostasis**, for protein homeostasis. Three crews do the work, and they form a triage.

The first crew are the **chaperones** — proteins whose job is to help *other* proteins fold. Many of them are called heat-shock proteins, because the cell makes more of them under stress, when folding is most at risk. It's important to be precise about what a chaperone does, because the obvious metaphor is wrong: a chaperone is not an instructor that teaches the protein its shape. The shape, remember, is in the sequence. A chaperone is more like a private folding room — it binds a newly made or wobbling chain, shields its sticky

patches from the crowd, prevents it from clumping with its neighbors, and gives it the space and the time to find its own way down the funnel. Some chaperones can grab a misfolded protein and give it another chance to refold. They are the cell's first response: *try to fix it.*

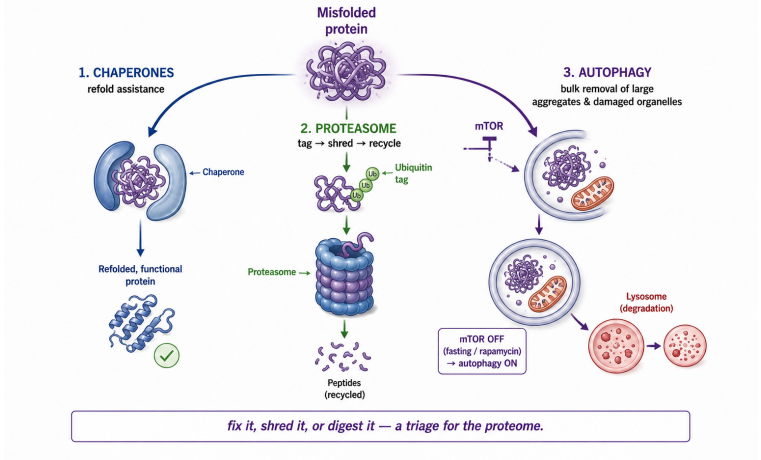
When a protein can't be fixed, the second crew destroys it — selectively. The cell tags doomed proteins with a chain of a small marker protein called **ubiquitin**, a molecular label that means "destroy this." The tagged protein is then fed into the **proteasome**, a barrel-shaped machine that grabs the labeled protein, pulls it apart, and chops the chain into small pieces for recycling. The ubiquitin-proteasome system is precise and fast, and it is the workhorse for disposing of individual, soluble, tagged proteins — one bad apple at a time.

But some things are too big for the proteasome to swallow: large aggregates, clumps that have already formed, and whole worn-out organelles like damaged mitochondria. For those, the cell uses the third crew, and it is the one that gives this chapter its second hallmark. **Autophagy** — from the Greek for "self-eating" — is the cell's bulk recycling system. The cell wraps a portion of itself, clumps and broken organelles and all, inside a newly made membrane sac, then delivers that sac to the **lysosome**, an acid-filled compartment full of digestive enzymes that breaks the cargo down into raw materials to be reused. Autophagy is how the cell takes out its large trash — and, as we saw at the end of the last chapter, *mitophagy*, the clearing of damaged mitochondria, is

just autophagy aimed at a specific target. The machinery of autophagy was worked out, painstakingly, in yeast by Yoshinori Ohsumi, who spent years watching the process under the microscope when almost no one thought it important; it turned out to sit at the center of aging, cancer, immunity, and neurodegeneration, and earned a Nobel Prize.

Together — fix it (chaperones), shred the tagged individuals (proteasome), digest the big stuff (autophagy) — these three crews wage a continuous war to keep the proteome in order. As long as they keep up with the rate at which proteins misfold, the cell stays healthy. The trouble, predictably by now, is that they don't keep up forever.

Figure 6.2 — The three cleanup crews.



Here is the now-familiar shape of aging, in its sixth costume. The misfolding of proteins is a roughly constant hazard — proteins are always at some risk of folding wrong, and stress and damage add to it over time. The cleanup crews hold the line by clearing the misfolded proteins as fast as they appear. Aging is what happens when **the clearing falls behind the misfolding.** And it falls behind because, with age, all three crews weaken: the cell makes fewer chaperones and mounts a feebler stress response, the proteasome grows sluggish, and — most consequentially — autophagy declines. The maintenance side of the balance erodes, exactly as it did with senescent-cell clearance in Chapter 3 and mitochondrial quality control in Chapter 5. The same story keeps recurring because it is, at bottom, *the* story: aging is largely the slow defeat of maintenance by damage, and what differs from chapter to chapter is only which maintenance system we're watching lose.

Once clearance falls behind, misfolded proteins accumulate, find each other's sticky patches, and aggregate — and now we have a vicious feature of the problem. Aggregates are not just inert garbage; they actively *clog the very crews meant to remove them*, jamming the proteasome and overwhelming autophagy, which lets still more aggregates form. The collapse, once it begins, accelerates.

And this is where one of the most satisfying connections in the whole book closes, linking this chapter directly back to Chapter 4. Recall the master growth switch, mTOR — active when food is plentiful, telling the cell to grow. One of the specific things mTOR does when it's active is **shut autophagy off.** Growth mode and cleanup mode are mutually exclusive: when the cell thinks food is abundant, it builds rather than recycles. Which means that the interventions that tip the nutrient-sensing dial toward scarcity — **caloric restriction, fasting, and rapamycin** — do something very concrete and very relevant here: by switching mTOR *off*, they switch autophagy *on*. They turn on the cleanup crew. A substantial part of *why* fasting and rapamycin are geroprotective, it turns out, is that they command the cell to stop growing and start taking out its accumulated trash. The hunger switch of Chapter 4 and the cleanup crews of Chapter 6 are the same circuit seen from two ends — which is also why the most evidence-backed longevity interventions and the proteostasis story are not two topics but one.

> **Deeper Water — the two disposal routes, and why neurons are doomed to lose.** It's worth understanding why the cell maintains two separate destruction systems, and why one kind of cell is especially helpless. The ubiquitin-proteasome system and autophagy are specialized for different cargo. The proteasome handles *soluble,*

individual proteins, tagged one at a time with ubiquitin and threaded into a narrow barrel — it physically cannot accept anything large or already clumped. Autophagy handles exactly what the proteasome can't: bulky aggregates, protein clumps, and whole organelles, swallowed wholesale. There is even a third, hybrid route — chaperone-mediated autophagy — in which chaperones ferry specific proteins one by one directly across the lysosome membrane. The redundancy is the point: different threats need different disposal. Now consider the predicament of a neuron. Most cells, when they accumulate damaged proteins, have an escape hatch unavailable to neurons: they *divide*, and each division halves the concentration of junk between two daughter cells — dilution as a cleanup strategy. But the neurons in your brain are *post-mitotic*: with rare exceptions they never divide again after you're young. The neuron you are using to read this sentence is, very likely, as old as you are. It cannot dilute its aggregates by dividing; it can only clear them in place, with crews that weaken over decades. So a neuron is a long-lived cell that depends *entirely* on its proteostasis crews, for a lifetime, with no fallback — which is precisely why the diseases of protein

aggregation are, overwhelmingly, diseases of the aging brain. The most irreplaceable cells we have are the ones least able to take out their own trash.*

The diseases we fear

This brings us to the most frightening corner of aging biology, and the one where honesty matters most, because these are the diseases that take our parents and that we fear for ourselves. Mechanistically, the major neurodegenerative diseases of aging are, at their core, **diseases of protein misfolding and aggregation** — the cleanup crews losing, in specific cells, to specific proteins.

In **Alzheimer's disease**, two proteins aggregate: fragments called amyloid-beta pile up into plaques *outside* neurons, and a protein called tau twists into tangles *inside* them. In **Parkinson's disease**, a protein called alpha-synuclein aggregates into clumps in the neurons that control movement. In **Huntington's disease** — the clearest case of all, because it is purely genetic — an inherited mutation lengthens a repetitive stretch of the huntingtin protein, and the elongated protein aggregates; the longer the inherited repeat, the earlier the disease strikes, a dose-response written into the genome. In **ALS** and related disorders, still other proteins (TDP-43 among them) misfold and clump. Different protein, different cell type, different symptoms — but the same underlying fail-

ure: a protein adopts the wrong shape, the wrong shapes stick together, the aggregates accumulate faster than the aging crews can clear them, and the affected neurons sicken and die.

Many of these aggregates share a common structural endpoint called **amyloid** — a particular, extremely stable way for misfolded proteins to stack together into fibers. And the most unsettling discovery about aggregation is that it can *spread* in a self-templating way. The principle was established with **prions**, the protein-only infectious agents behind mad cow disease and its human counterparts — an idea so radical when proposed (a protein, with no genes, that could be infectious?) that it was widely dismissed before it won a Nobel Prize. A prion is a misfolded protein that, on contact, induces *correctly* folded copies of the same protein to misfold like it — a bad shape that propagates itself, converting its neighbors. We now know that this seeding behavior is not unique to prions: amyloid-beta, tau, and alpha-synuclein can all seed their own misfolding in this prion-like way and spread the aggregation from neuron to neuron, which helps explain how these diseases march through the brain along predictable routes. The wrong shape is, in a real sense, contagious from molecule to molecule.

> **Deeper Water — why amyloid is so stable, and why a shape can be contagious.** Amyloid is not a random clump; it is a highly ordered struc-

ture, in which the protein chains line up side by side into long, flat, ribbon-like sheets (called cross-β sheets), stacked into fibers held together by a dense web of hydrogen bonds. This architecture is extraordinarily stable — more stable, often, than the protein's own correctly folded shape — which is exactly why it is so hard for the cleanup crews to take apart and why, once formed, amyloid tends to persist. It is also why so many *different* proteins, with no sequence in common, can all end up as amyloid: the cross-β structure is a kind of generic low-energy basin that almost any protein chain can fall into if it misfolds, a deep secondary well in the folding landscape from the start of the chapter. The self-templating — the "contagious shape" — follows from the structure too. An amyloid fiber has an exposed growing end, and a soluble protein of the same kind that bumps into that end is coaxed to add itself on in the same cross-β arrangement, extending the fiber. Each newly recruited molecule becomes a new template surface. This is why aggregation, once seeded, is autocatalytic: the product of misfolding is itself a catalyst for more misfolding. It is the same logic as a crystal growing in a supersaturated solution — and it is why a tiny seed of aggregate, or in

the case of prions a tiny infectious dose, can eventually convert a vast amount of protein.*

Thirty years chasing a plaque

Now the chapter's hard, real-stakes uncertainty — and it is, by design, the book's most important one, because more readers will face this disease in someone they love than any other in these pages.

For roughly thirty years, one idea dominated Alzheimer's research with a force few hypotheses ever achieve: the **amyloid cascade hypothesis**, which held that the accumulation of amyloid-beta is the *primary cause* of Alzheimer's, the first domino, with tau tangles, inflammation, and neuron death all following downstream. It was a powerful and well-motivated idea — and the strongest evidence for it is genuinely compelling, so let me state it fairly before complicating it. The genetics point hard at amyloid: every known mutation that *causes* the rare inherited, early-onset form of Alzheimer's affects the production of amyloid-beta; people with Down syndrome, who carry an extra copy of the gene that makes amyloid-beta's parent protein, develop Alzheimer's almost universally and early; and the main genetic risk factor for ordinary Alzheimer's affects how amyloid is cleared. Amyloid is not an innocent bystander. It is causally entangled with the disease.

So the field made the obvious prediction — *remove the amyloid and you should stop the disease* — and spent decades and tens of billions of dollars building drugs, mostly antibodies, to clear it from the brain. And for most of those decades, the result was a **graveyard** that should sound familiar after the last chapter. Drug after drug succeeded in clearing amyloid from patients' brains — they hit the target, they moved the biomarker — and *failed to improve patients' memory or slow their decline.* One after another, anti-amyloid drugs cleared the plaques and left the patients no better. It was the predictor-versus-lever problem from Chapters 1 and 2, and the biomarker trap from those chapters, playing out at the scale of an entire field and a generation of patients: amyloid was a thing you could *measure* and a thing you could *move*, and for thirty years, moving it did not move the disease. The low point came when a regulator approved an anti-amyloid drug largely on the grounds that it cleared amyloid — the biomarker — over the objections of its own advisors who doubted any real benefit to patients; the drug was mired in controversy and ultimately commercially abandoned. The episode was the surrogate-endpoint cautionary tale of Chapter 4 made horribly concrete: a drug approved for moving a marker, with the thing that actually mattered — whether patients got better — left unproven.

The story grew darker still in 2022, when an investigation raised serious concerns that images in a highly influential 2006 paper — one that had pointed to a specific amyloid spe-

cies as a cause of memory loss and had shaped years of research — had been manipulated. (The paper was later retracted.) It is important to be precise about what this does and doesn't mean, because it is easy to over-read. The alleged manipulation damaged one influential *thread* of the amyloid story and rightly shook confidence; it did *not*, by itself, overturn the amyloid hypothesis, which rests on a much broader base of genetic and biological evidence than any single paper. But it is a sobering reminder, in the spirit of this whole book, that even a field this large and this well-funded can be distorted by a single compelling-but-flawed result, and that prominence is not the same as truth.

And then — the twist that makes this the book's most genuinely nuanced case rather than a simple Chapter 5–style debunking. In the last few years, the *newest* anti-amyloid antibodies finally did what thirty years of predecessors could not: in large trials, they **slowed cognitive decline by a statistically significant amount.** Two such drugs have now been approved. This is real, and it partially vindicates the hypothesis — removing amyloid *does* help, which means amyloid really is causally involved, not merely correlated. But read the result with the discipline this book has tried to build, because the same trials carry the other half of the truth. The benefit is **modest** — a slowing of decline, not a halt and certainly not a reversal, and experts genuinely debate whether the size of the effect is large enough for patients and families to notice. The drugs carry real risks,

including brain swelling and small bleeds visible on MRI that are occasionally serious. They are expensive, burdensome, and work mainly in early disease. So the honest verdict holds both halves at once: after thirty years and a graveyard of failures, clearing amyloid was shown to help — a little — which means the hypothesis was *neither wrong nor right* in the simple way, but **partially correct and badly oversold.** Amyloid is part of the story; the modest size of the benefit from removing it is the clearest possible evidence that it is *not the whole story* — that tau, inflammation, vascular damage, and metabolic failure are all in the mix, and that Alzheimer's is a genuinely multifactorial, still-contested disease.

This is the place in the book where the discipline matters most, because the stakes are highest and the marketing — of drugs, of supplements, of "brain health" everything — is loudest. Anyone who tells you Alzheimer's is *solved*, or that it is *definitely just amyloid*, or that the new drugs *prove amyloid was a dead end*, or that some supplement *prevents* it, is overclaiming in one direction or another. The truth, as of now, is harder and more honest: amyloid matters, removing it helps modestly, the disease is more complex than any single protein, and the field is still genuinely uncertain. A reader who learns only a confident headline — in either direction — will misread the news, and possibly the medical choices, for the rest of their life.

The most celebrated AI in all of biology

We end the chapter where this book's deepest theme — the promise and the limits of artificial intelligence in medicine — reaches its single most spectacular example. Because the folding problem from the start of this chapter, the fifty-year grand challenge of predicting a protein's shape from its sequence, was *solved* — by an AI.

Recall the setup: Anfinsen proved the sequence determines the structure, which meant that, in principle, you should be able to *compute* a protein's shape from its amino-acid sequence alone, without the slow, expensive experiments (crystallizing the protein, bombarding it with X-rays) that had been the only reliable way to see a structure. For half a century, "in principle" was as far as it went; the computation was intractable, and predicting structure from sequence was considered one of the hardest open problems in all of biology, with its own biennial worldwide competition to measure the meager progress. Then, in 2020, a system called **AlphaFold**, built by the AI laboratory DeepMind, entered that competition and effectively ended it — predicting the structures of many proteins to an accuracy rivaling the experimental methods, a leap so large it stunned the field. AlphaFold's predicted structures were soon released, free, for essentially *every protein known to science* — hundreds of millions of them, a resource that would have taken the old experimental methods many lifetimes to produce. Its creators

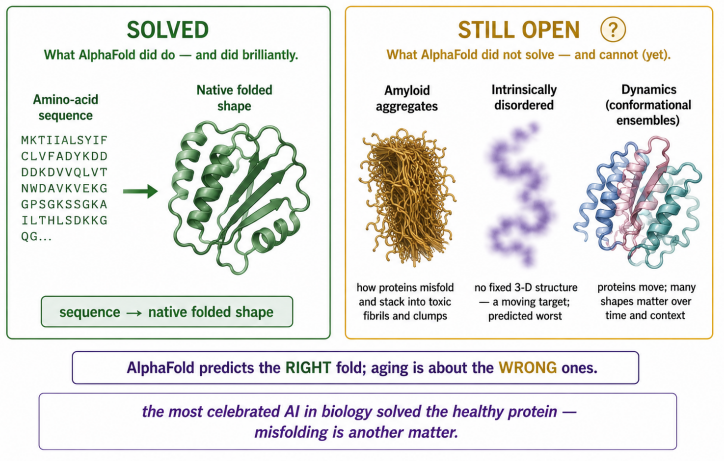
were awarded the Nobel Prize in Chemistry. It is, without serious competition, the most celebrated achievement of artificial intelligence in the history of biology, and it has genuinely accelerated structural biology and the early stages of drug discovery — finding the shape of a target is often the first step to designing a drug against it.

It deserves every bit of that credit. And now the honest footnote, which is not a grudging asterisk but the very heart of this chapter — because of *what specifically* AlphaFold solved, and what it did not. **AlphaFold predicts the native, correctly folded shape of a well-behaved protein.** It tells you, brilliantly, the *right answer* — the bottom of the folding funnel. But everything that has made *this chapter* a chapter about *disease* is about the **wrong answers** — and those AlphaFold largely does not predict. It does not predict **misfolding**: the very thing that causes Alzheimer's and Parkinson's and the rest is a protein going to the wrong shape, and AlphaFold models the right one. It does not predict **aggregation**: how misfolded proteins clump is outside what it computes. It struggles most with **intrinsically disordered proteins** — proteins that don't fold into one fixed shape at all — and here is the cruel irony: several of the most important aggregation-prone proteins in neurodegeneration, including tau and alpha-synuclein and amyloid-beta itself, are themselves intrinsically or partly disordered — the kind AlphaFold handles worst. It largely gives a single *static* snapshot, while proteins flex and breathe and visit multiple

shapes, and function often lives in that motion. And predicting how a small-molecule drug will bind, or how a mutation will change a protein's behavior, remains hard, though newer versions reach toward protein interactions and drug binding.

So AlphaFold is a triumph at predicting the right fold — and aging, this chapter has shown, is largely about the wrong ones. That is not a knock on AlphaFold; it is the precise shape of what it gave us and what it didn't, and it generalizes the lesson the whole book keeps relearning. A spectacular tool that solves one cleanly posed problem — *sequence to native structure* — is not the same as a solution to the messy, downstream, clinical problem you actually care about — *why proteins misfold, how to stop them aggregating, how to keep an aging neuron's cleanup crews from losing.* The most dazzling AI in biology mastered the healthy protein. The diseased one — the misfolded, the aggregated, the disordered — is still, largely, waiting.

Figure 6.4 — What AlphaFold did and didn't do.



✦ ✦ ✦

CARRY THESE THREE THINGS

1. **Aging is the cleanup crews losing — and the same crews are why fasting works.** The shape of a protein is its function, and the cell wages a constant war to keep proteins folded (chaperones), shred the broken ones (the proteasome), and digest the big aggregates and worn-out organelles (autophagy). With age, all three crews weaken, misfolding outpaces clearing, and toxic aggregates accumulate — the same "maintenance loses to damage" story as senescent cells and mitochondria, in a sixth costume. And the circuit ties straight back to Chapter 4: mTOR shuts autophagy off, so fasting and rapamycin — by switching mTOR off — switch the

cleanup crew *on*. The hunger switch and the cleanup crew are one circuit.

2. **The amyloid story is the biomarker trap at the scale of a generation — and the most important place to hold two truths at once.** For thirty years the field chased amyloid, a marker it could measure and move; drug after drug cleared the plaques and left patients no better, and one was even approved on the marker alone with the benefit unproven. Then the newest drugs finally slowed decline — a *modest*, risky, real effect. So the hypothesis was neither simply right nor simply wrong: amyloid is genuinely part of the disease (removing it helps a little), and the smallness of that help is the proof that it is *not the whole story*. Alzheimer's is multifactorial and still contested, and anyone who tells you it's solved — in any direction — is overclaiming where the stakes are highest.
3. **The most celebrated AI in biology solved the right fold; aging is about the wrong ones.** AlphaFold cracked a fifty-year grand challenge — predicting a protein's native structure from its sequence — and earned a Nobel for it, and genuinely transformed structural biology. But misfolding, aggregation, the disordered aggregation-prone proteins of neurodegeneration, and protein motion are largely what it *doesn't* predict — and those are exactly what these diseases are made of. A dazzling solution to one clean problem

is not a solution to the messy one you care about. Every chapter so far has been about damage accumulating and maintenance falling behind, and every intervention about slowing the damage or clearing it. The last chapter asks the most radical question in the field: not whether we can slow the damage, but whether we can *wipe it away* — reset an old cell to a young state, the way each of us was reset at conception. It is the most exciting idea in aging biology, the one just now entering its very first human trial — and, you will not be surprised to learn, the least proven of all.

Chapter 7 · The Reset

INSTEAD OF SLOWING THE DAMAGE, WHAT IF WE ERASED IT — *reset an old cell to a young one? It is the most radical idea in aging biology, the one most entangled with cancer, and, as of this writing, the subject of its very first human trial.*

> **Companion notebook · Chapter 7.** No separate notebook for this chapter — partial reprogramming rides the same hormesis curve as Chapter 4: a narrow safe window between doing nothing and triggering cancer: Open the Chapter 4 notebook →

WHY WE STOP HEALING

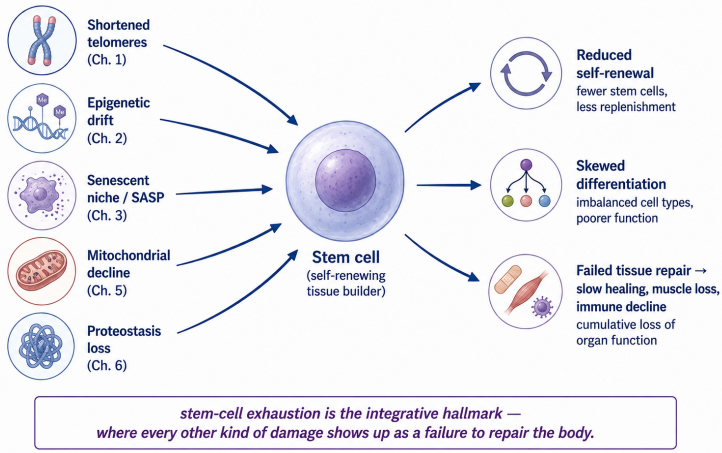
A child's scraped knee heals in days, leaving nothing behind. The same scrape on an eighty-year-old heals slowly, sometimes incompletely, sometimes not at all. Somewhere between those two bodies, the capacity to repair itself quietly drained away — and where it went is the subject of this final chapter, and the answer turns out to gather almost everything in the six chapters before it.

The body repairs and renews itself using **stem cells** — a reserve of unspecialized cells, kept in every tissue, that can do two things no ordinary cell can: copy themselves indefinitely (self-renewal), and turn into the specialized cells a tissue needs (differentiation). Your blood is completely replaced every few months from a small pool of blood-forming stem cells; the lining of your gut is rebuilt every few days; your muscle repairs itself after exercise or injury by calling on dormant muscle stem cells. Stem cells are how a body that is constantly wearing out manages, for most of a lifetime, to keep itself whole. And the hallmark of aging called **stem-cell exhaustion** is the slow failure of that reserve: with age, stem cells decline in function, lose self-renewal capacity, differentiate in skewed and unhelpful ways, and find themselves in a degraded local environment that no longer supports them well. The reserve runs down, and the body's lifelong project of repairing itself falters. This is why old wounds heal slowly, why muscle is lost (sarcopenia), why the immune system renews itself less effectively, why so many tissues thin and weaken.

Here is what makes this the right place to end the tour: stem-cell exhaustion is less a *separate* hallmark than the place where all the others come home to roost. Why do aged stem cells fail? Because their telomeres have shortened (Chapter 1). Because their epigenome has drifted (Chapter 2). Because their niche is poisoned by senescent cells and the SASP (Chapter 3). Because their mitochondria and proteostasis

have declined (Chapters 5 and 6). Every form of damage we have catalogued converges on the stem cell, and the visible result of that convergence is a body that can no longer repair itself. Stem-cell exhaustion is aging made manifest as a failure to heal.

Figure 7.1 — Where all the damage comes home.



There is a famous and slightly eerie experiment that points toward this chapter's radical idea. If you surgically join the circulatory systems of a young mouse and an old mouse, so they share blood, the old mouse's tissues partially *rejuvenate* — its muscle and liver and brain stem cells start behaving more youthfully — while the young mouse ages. The result was electrifying, because it suggested that aging is not solely locked inside each cell but is influenced by *signals circulating in the blood and the cellular environment* — and that an old stem cell is not necessarily broken beyond repair, but

partly *held down* by its surroundings. This launched a wave of excitement about "young blood" as a fountain of youth. But the honest follow-up, very much in the spirit of this whole book, complicated the fairy tale: the rejuvenation appears to come less from magical youthful factors in young blood than from *diluting the pro-aging factors that accumulate in old blood* — with growing evidence that old blood being bad matters more than young blood being good. And the commercial young-blood-transfusion clinics that sprang up to sell the dream were selling something unproven; regulators warned against them, and they faded. The real lesson survived the hype: the aged cell is partly suppressed by its state, which means it might, in principle, be *un*-suppressed — reset. Which raises the most radical question in the field. If an old cell can be coaxed to behave young, how far back can we push it? All the way?

The one-way trip, reversed

To see how astonishing the answer is, you need a picture of cell identity that biologists carried, essentially unquestioned, for most of a century.

The picture is the **Waddington landscape**, named for the biologist who drew it in 1957, and it is the developmental cousin of the folding funnel from the last chapter. Imagine a ball at the top of a hilly landscape, poised to roll down. At the top, the ball is a fertilized egg or an embryonic cell —

pluripotent, able to become anything. As it rolls down, the landscape branches into valleys, and the ball is funneled into one of them: it becomes a neuron, or a skin cell, or a muscle cell, and as it descends further into its valley it becomes more specialized, more committed, more locked in. The valley walls hold it in place; what *is* those walls, physically, is the epigenome from Chapter 2 — the pattern of methylation and chromatin marks that defines and maintains the cell's identity, telling a skin cell to keep being a skin cell. And the deep, century-old assumption was that this is a **one-way trip**. The ball rolls downhill into its valley and stays there. A neuron is a neuron for life. You cannot, the thinking went, push the ball back up the hill.

In 2006, Shinya Yamanaka pushed the ball back up the hill — and won a Nobel Prize for proving the one-way trip was nothing of the kind. He took a fully specialized adult cell — an ordinary skin cell, sitting deep in its valley — and forced it to switch on just **four genes** (master regulators now called the Yamanaka factors). And the skin cell climbed back out of its valley, all the way to the top of the landscape, becoming a pluripotent cell again — able, once more, to become any cell type in the body. These reprogrammed cells are called induced pluripotent stem cells. Cell identity, it turned out, was not a permanent destination but a *reversible state*, held in place by the epigenome and capable of being reset. The landscape was not a one-way slope; with the right push, the ball could roll back up. (Yamanaka's leap built on a stranger,

earlier one: decades before, John Gurdon had shown by transplanting the nucleus of an adult cell into an egg that an adult cell still *contains* all the instructions to build a whole organism — the same insight that, years later, produced Dolly the cloned sheep. The information for youth was never lost. It was only switched off.)

Resetting the clock

Now connect that to aging, and the radical promise comes into focus.

When you reprogram an old cell back toward pluripotency, you do not only reset its *identity*. You reset its *age*. The epigenetic clock from Chapter 2 — the methylation pattern that tracks biological age — runs almost all the way back to **zero**. Telomeres, the counters from Chapter 1, can lengthen again. The accumulated epigenetic damage of decades is wiped, and the cell becomes, by the molecular measures we have, *young* again. This is rejuvenation of a wholly different kind from anything in the previous six chapters. Every intervention we have examined so far tried to *slow* the damage (the nutrient dial, Chapter 4) or *clear* it (senolytics, autophagy, Chapters 3 and 6) — to help the cell's failing maintenance keep up. Reprogramming proposes something far more audacious: not to slow or clear the damage but to *erase* it, resetting the cell's age the way every one of us was reset at conception, when two aged parental cells gave rise to a

brand-new, age-zero embryo. The reset button, it turns out, exists. The cell has always known how to become young, because it does it every generation.

Deeper Water — what reprogramming actually does, and the two things it resets at once. The Yamanaka factors are transcription factors — proteins that bind DNA and switch genes on and off — and what they do, mechanistically, is drive a wholesale remodeling of the epigenome. They recruit the machinery that erases and rewrites methylation marks and reorganizes chromatin, dismantling the epigenetic pattern that defined the old, specialized cell and re-establishing the pattern of a pluripotent one. Crucially, this remodeling erases *two distinct things that happen to be written in the same ink*: the marks that specify the cell's **identity** (skin cell, neuron) and the marks that have accumulated with **age** (the epigenetic drift of Chapter 2). The central scientific bet of this entire field is that these two can be *decoupled* — that you can erase the age marks while leaving the identity marks intact, rolling the ball partway up the hill to shed its years without rolling it so far that it forgets what kind of cell it is. Whether age and identity are cleanly separable, or whether they are too entangled in

the same epigenetic ink to pull apart safely, is one of the deepest open questions in the field. This is also, you'll notice, the most literal possible version of the "information theory of aging" from Chapter 2: if aging is the corruption of epigenetic information, then reprogramming is restoring the information from the cell's own developmental backup — and the question of whether the backup is clean enough to restore from is exactly the question of whether the age marks and identity marks can be told apart.*

The catch: youth and cancer, one last time

By now you can feel the shape of what's coming, because it is the shape this book keeps returning to. The reset button comes with a price, and the price is the same one we have paid in chapter after chapter.

The problem with *full* reprogramming is that erasing a cell's age all the way to zero also erases its identity all the way to zero. A skin cell reprogrammed to pluripotency is no longer a skin cell; it doesn't do the skin's job. And pluripotent cells loose in a living body do something terrible: they form **teratomas** — grotesque tumors containing a chaotic jumble of tissues, teeth and hair and gut lining all tangled together, the unmistakable signature of cells with unlimited developmental potential and no instructions. So you cannot simply

switch the four factors on throughout an old person's body. You would not rejuvenate their organs; you would dissolve them into tumors.

This is the youth-and-cancer entanglement from Chapter 1 and Chapter 3, arriving now in its deepest and most fundamental form. In Chapter 1, the enzyme that keeps cells young (telomerase) was the one cancers switch on to become immortal. In Chapter 3, the cells that protect us from cancer when young (senescent cells) poison us when old. Here, at the bottom of the whole stack, the principle reaches bedrock: **the very plasticity that lets a cell become young again — the loosening of its fixed identity, the reawakening of developmental potential — is exactly the plasticity that cancer exploits.** Cancer, too, is in large part a loss of stable cell identity and a return to a more primitive, unlimited state. Reprogramming and cancer are not neighbors; they are drawing on the same well. The most powerful rejuvenation we know is dangerous for precisely the same reason it is powerful. The brake-and-price theme that opened this book — that in aging biology, the obvious intervention is never free, because the thing that makes a cell young is bound up with the thing that makes it cancerous — is not a coincidence repeated across seven chapters. It is, it turns out, a single deep fact about life, seen from seven angles.

So the whole game, the entire frontier, comes down to a single word: **partial.**

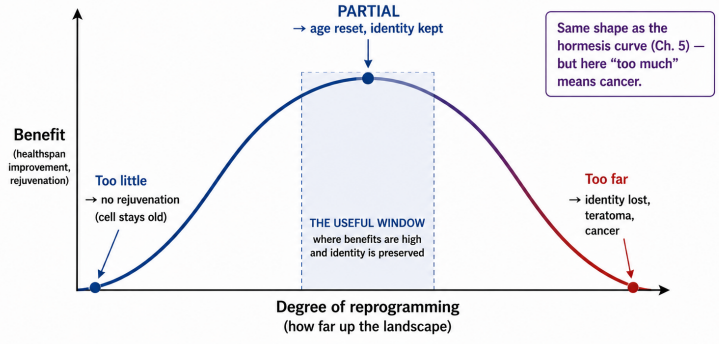
The idea, first demonstrated convincingly in 2016, is to roll the ball only *partway* up the landscape — far enough to shed the age-related epigenetic marks, but not so far that the cell loses its identity or becomes pluripotent. Instead of switching the reprogramming factors on permanently, you apply them in brief *pulses* — on, then off, then on again — or for a limited time, hoping to catch the cell in a rejuvenated-but-still-itself state before it climbs too high. When researchers did this cyclically in fast-aging mice, the animals' aging signs improved and they lived longer, *without* the teratomas that full reprogramming would have caused. Partial reprogramming could rejuvenate a living mammal, at least in that controlled setting. The reset button could, apparently, be tapped rather than held down.

The most striking demonstration of partial reprogramming restoring actual *function* came from David Sinclair's laboratory, and it deserves to be described plainly and credited fully, separate from any of the surrounding hype. His group used three of the four factors (dropping the fourth, c-Myc — the most overtly cancer-causing of them) and delivered them to neurons in the eye. In old mice, and in mice whose optic nerves had been crushed or damaged in a model of glaucoma, this partial reprogramming **restored vis-**

ion — resetting the epigenetic age of the retinal neurons and allowing damaged nerve fibers to regrow, something adult nerve cells normally cannot do. This is real, peer-reviewed, repeatedly examined science, and it is among the most impressive demonstrations in the field that partial reprogramming can not only make cells *measure* younger but make a tissue *work* younger. Later work from the same lab extended the approach toward protecting neurons from the inflammation and degeneration of aging. The science here is genuine, and it is the engine behind everything in the next section.

But notice the catch hiding inside the word "partial," because it is the named uncertainty of this chapter and the reason the frontier is a frontier. *How far is partway?* Reprogram too little, and you reset nothing — the cell stays old. Reprogram too much, and you lose identity and risk teratomas and cancer. The useful window lies between, and it is **narrow and poorly defined** — and it is, you will recognize, the same non-monotonic shape we met as hormesis in Chapter 5: too little does nothing, too much does harm, and the benefit lives at an intermediate dose you have to hit precisely. Except here the failure mode on the "too much" side is not merely losing exercise's benefits — it is cancer. Threading this needle, in the right cells, to the right degree, reversibly, in a living human being, is the entire problem. And we are only just beginning to try.

Figure 7.3 — Threading the needle.



THE GOLD RUSH, AND THE FIRST HUMAN

The promise of reprogramming — erase the damage, reset the age, regrow what's lost — has set off the largest concentration of money and talent in the history of aging science. A single company launched in 2022 with roughly three billion dollars in funding, the most ever raised by a biotechnology startup at birth, devoted to cellular rejuvenation, and it recruited a stunning roster of the field's leaders — including, fittingly, the inventor of the partial-reprogramming approach and the inventor of the epigenetic clock from Chapter 2. Other well-funded ventures followed, backed by some of the most prominent names in technology, several explicitly using machine learning to search for better and safer reprogramming recipes. Reprogramming is, right now, the biggest bet in the field, and the bettors are not naive — they are the people

who know best how hard it is, wagering that the needle can be threaded.

And here, at last, is the moment this entire book has been building toward — stated plainly, and then weighed honestly, because the gap between the two is the lesson of all seven chapters.

At the time of writing, partial reprogramming has entered its very first human clinical trial. Sinclair's company received regulatory clearance to begin a Phase 1 study of partial epigenetic reprogramming — delivering the three-factor combination, as a gene therapy, to treat a form of vision loss in the human eye. The most radical idea in aging biology, the literal resetting of cellular age, is being tested in a human being for the first time. That is a genuine milestone, and it is right to find it remarkable.

Now read it with the discipline this book has spent six chapters trying to build. *What this is:* one gene-therapy approach, entering the earliest and smallest phase of human testing — the phase that asks only "is it safe?", not "does it work?" — for **one narrow condition**, in the eye. The eye was chosen for excellent reasons that reveal exactly how cautious this first step is: it is small, contained, and accessible, you can deliver the therapy locally without exposing the whole body, it is partly shielded from the immune system, and you can watch it closely. This is the most careful possible first toe in the water. *What this is not:* it is not the reversal of human aging. It is not whole-body rejuvenation. It does not yet tell

us whether partial reprogramming is safe over years, whether the age-and-identity decoupling holds up in humans, or whether the cancer risk can be controlled in living tissue over time. Between "a Phase 1 trial of a reprogramming gene therapy for one eye disease" and "we can make people young again" lies the entire distance this book has been measuring — the distance between an astonishing result and a proven human therapy, between the mouse and the person. The first human trial of reprogramming is the most exciting sentence in modern aging biology. It is also a single, narrow, early, careful experiment. Both of those things are completely true, and holding them together — neither dismissing the milestone nor inflating it — is the final exam of everything you have learned.

Deeper Water — the control problem, in full. What stands between the eye trial and real anti-aging reprogramming is a cluster of hard control problems, and they are worth naming concretely. *Dosing and reversibility:* partial reprogramming requires switching the factors on, partway, and off again, at the right level — but delivering this as a gene therapy makes precise, reversible dosing genuinely difficult; once you've put the genes into a cell, controlling exactly how much and how long they act is far harder than with a drug you can stop taking. *Cell-to-cell het-*

erogeneity: different cell types climb the landscape at different rates, so a "pulse" that is gently partial for one cell may overshoot into dangerous dedifferentiation for another — there may be no single dose that is partial for every cell in a tissue at once. *The cancer tail risk:* even three-factor reprogramming reawakens developmental and proliferative programs that overlap with cancer, and a rare cell that goes too far, or stays reprogrammed too long, could seed a tumor — a risk that may be invisible for years and only matters at the scale of a whole body over a lifetime. *Delivery:* getting the therapy into the right cells of an internal organ — a liver, a brain, a heart — without affecting the wrong ones is unsolved for most tissues; the eye is uniquely easy, which is why it is first. *And the deepest unknown:* we do not actually know whether age and identity can be cleanly separated in human cells over the long term, or whether sustained partial reprogramming slowly erodes cell identity in ways that take years to show. None of these is obviously unsolvable. But none is solved, and each is a reason the field is moving — correctly — one cautious, contained step at a time.*

Where the machines come in

The reprogramming story also points to where artificial intelligence may matter most in the future of this field — and, in doing so, closes the book's longest-running theme.

The obstacle above — that gene therapy makes reprogramming hard to dose, control, and reverse — has an appealing potential solution: do it with **small molecules** instead of genes. A drug can be given and stopped, titrated up and down, far more easily than a gene therapy can be switched on and off inside a cell. And it turns out that reprogramming with drugs alone is possible — researchers have reprogrammed cells all the way to pluripotency using nothing but cocktails of small molecules, no added genes at all, and labs (including Sinclair's) have screened for chemical combinations that partially reset cellular age. This is precisely the kind of problem artificial intelligence is suited to attack: searching the vast space of possible molecules and combinations to find the ones that reset age safely and partially, guided by imperfect readouts such as the epigenetic clock, or markers of cell identity. The discovery of safe chemical reprogramming cocktails is one of the most promising places in all of aging research for AI-driven drug discovery to make a real difference, and it is the cleanest possible return to this series' founding premise: using machine learning to find the molecules that do what biology has shown is possible.

And yet — one final time, because it is the discipline the whole book has been built to instill — the same caution applies, sharper than ever. An AI told to search for molecules that "lower the epigenetic clock" will faithfully find molecules that *lower the epigenetic clock* — and you already know, from Chapter 2, that lowering the clock is not the same as making a cell genuinely younger; a model optimizing a surrogate will hand you expert surrogate-gamers. An AI optimizing for "reprogramming" without a precise enough definition of *partial* and *safe* may hand you molecules that reprogram beautifully and cause cancer. The machine can search the chemical space far faster than any human, and that is genuinely valuable. But it cannot tell you whether resetting these cells will make a human being healthier and not sicker, it cannot guarantee that the readout it optimizes is the thing you actually want, and it cannot run the years-long human trials that are the only thing that will ever settle whether the reset button is safe to press. The most radical idea in aging biology will be tested, in the end, the same way every idea in this book had to be: not by the elegance of the mechanism, not by the speed of the search, not by the size of the bet, but in a human being, slowly, carefully, and with the humility to be told no.

✦ ✦ ✦

Carry these three things

1. **Reprogramming is the most radical answer to aging: not slow the damage, not clear it, but erase it.** Stem-cell exhaustion is where every other hallmark converges as a failure to repair the body — and reprogramming proposes to reset the cells themselves. Yamanaka proved cell identity is reversible (the one-way trip up the Waddington landscape *can* be reversed), and reprogramming an old cell resets its epigenetic clock toward zero and can lengthen its telomeres — making many of its molecular markers look young, the way each of us was reset at conception. The reset button exists, because the cell rebuilds a young body from old parents every single generation.
2. **Youth and cancer are entangled at the deepest level — and now you can see it was always one fact.** Full reprogramming erases identity along with age and turns tissues into teratomas, because the very plasticity that makes a cell young is the plasticity cancer exploits. This is the brake-and-price theme from Chapter 1 (telomerase) and Chapter 3 (senescence) reaching bedrock: across seven chapters, the thing that rejuvenates and the thing that causes cancer keep turning out to be the same thing, seen from different angles. The whole frontier is therefore the single word *partial* — threading a narrow window between doing

nothing and causing cancer, the hormesis curve of Chapter 5 with the highest possible stakes.

3. **The most exciting sentence in aging biology is also a single narrow early experiment — and holding both is the final exam.** Partial reprogramming restored vision in mice and extended life in fast-aging mice, and, as of this writing, it has entered its first human trial — a Phase 1 safety study, of a gene therapy, for one eye disease, in the most contained tissue we have. That is a real milestone *and* it is not the reversal of human aging; between the two lies the entire mouse-to-human distance this book has been measuring from its first page. We have now walked all seven hallmarks — counted the telomeres, read the clocks, cleared the zombie cells, turned the hunger switch, buried the beautiful wrong theory, watched the proteins fold and fail, and pressed the reset button. What remains is the reckoning: after all of it, what do we actually know — and can we, in fact, expect to live to a hundred and twenty? That is the last thing this book owes you, and it is where we go to close.

Chapter 8 · The Reckoning

AFTER SEVEN HALLMARKS — CAN WE EXPECT TO LIVE TO A HUN-*dred and twenty? The honest answer, and the tools to keep asking it for the rest of your life.*

> **Companion notebook · Chapter 8.** *Curing One Disease, or Slowing Aging?* — open it in Google Colab and build it yourself: Open in Colab →

We began with a question you may have half-believed and half-doubted: whether you will live to a hundred and twenty. We have now done the work that lets us answer it honestly — walked all seven hallmarks, counted the telomeres and read the clocks, met the cells that won't die and the switch that senses hunger, buried the most beautiful wrong theory in the field, watched proteins fold and fail, and pressed the reset button. So let us gather what we've learned, answer the question we came in with, and — more lasting than any single answer — leave you with the tools to keep asking it for the rest of your life, as the headlines keep coming.

What aging actually is

Start with what aging actually *is*, now that we've seen it from seven angles. The most important thing the tour reveals is that aging is **not one thing.** It is many distinct, interacting failures happening at once: caps wearing off chromosomes, the chemical marks on our genes drifting out of order, damaged cells accumulating and poisoning their neighbors, nutrient-sensing circuits stuck in the wrong setting, mitochondria and protein-folding declining, the stem-cell reserve running down. There is no single switch labeled "aging," which is the first reason there will be no single pill to stop it.

And beneath the diversity, one pattern recurred chapter after chapter until it became unmistakable. Aging is, in large part, **the slow defeat of maintenance by damage.** Your cells are under constant assault — and they carry elaborate systems to repair, replace, and clean up what gets damaged. For most of your life, maintenance keeps up. Aging is what happens, gradually, as the maintenance systems *themselves* weaken and fall behind: the cleanup of worn-out cells slows, the quality control on mitochondria falters, the protein-folding crews lose ground, the stem cells that rebuild tissues exhaust. Damage is roughly constant; maintenance declines; and the widening gap between them *is* growing old.

The tour revealed one more thing — the reason aging is so much harder to fix than the headlines suggest. The levers are **entangled with cancer.** Over and over: the enzyme that

keeps cells young is the one cancers switch on to live forever (Chapter 1); the cells that protect us from cancer when we're young poison us when we're old (Chapter 3); the very plasticity that would let us reset a cell's age is the plasticity cancer exploits (Chapter 7). The thing that would rejuvenate us keeps turning out to be bound up with the thing that would kill us, for the same underlying reason. This is not bad luck repeated seven times. It is a single deep fact: many of the brakes that age us were installed *to protect us from cancer*, and releasing them is dangerous. Aging is, in significant part, the price we pay for not getting cancer at thirty. That is why "just turn the clock back" is almost never free — and why this field is genuinely, structurally hard, not merely underfunded or unlucky, but hard in the bones.

So — a hundred and twenty?

Now the question. Can you expect to live to a hundred and twenty? The honest answer requires both distinctions from the introduction.

Take **maximum lifespan** — the ceiling — first. The longest verified human life ended at 122, in 1997, and in the nearly thirty years since, no one has clearly approached it. One camp of researchers argues there is a real biological ceiling, somewhere around 115 to 125; another argues there is no fixed limit and the ceiling will eventually give way. The debate is genuinely unresolved. But here is what matters for

you: *nothing demonstrated today* — in any human being, in any trial — has moved that ceiling, and the burden of proof sits squarely on those who claim it will. The dramatic life-*doubling* results are all in worms and mice, and you now know to ask whether they survive the jump to a human. For raising the human maximum, the honest answer is that they have not — not yet, and perhaps not for a long time.

Now take **life expectancy** — the average — which is what most people actually mean. Here the most rigorous recent evidence is sobering for the optimists. Careful analysis of the world's longest-lived populations finds that the dramatic gains of the last century have *slowed*, and that, absent some breakthrough that targets aging itself, average life expectancy in even the best-off countries appears to be leveling off in the high eighties — not climbing toward a hundred and twenty. The easy gains, from stopping early death, are mostly spent. Pushing the *average* into the hundreds would require not better treatment of the diseases we already fight, but a fundamental intervention in aging itself — exactly the thing that, as this whole book has shown, remains unproven in humans.

So the honest verdict on the hundred-and-twenty question is this: **as a routine, average outcome, with anything demonstrated today — no.** Not for most people, not soon. Some individuals *do* reach a hundred, even a hundred and ten — but that is the existing high tail of human variation, shaped heavily by inherited luck (longevity runs in families, and a few people are simply built to last), not a sign

that the whole population is shifting upward. Reaching a hundred is real; it is the top of the distribution we already have, not evidence of a new one. The gap between "a few people reach a hundred" and "we can all expect a hundred and twenty" is, one final time, the gap this book has been measuring from its first page.

The prize worth wanting

But here is where the honest answer turns genuinely hopeful — if you've kept the *second* distinction in mind. The most defensible goal of aging science, and arguably the most valuable, may not be lengthening life at all. It is lengthening *healthy* life: compressing the years of frailty and disease at the end into a shorter span, so that you stay vigorous longer and decline quickly at the very end rather than spending a decade or two in slow deterioration.

And there is a powerful argument, from the field itself, for why this is the right target. Because the diseases of aging arrive *together* — driven by the same underlying processes you've now studied — fighting them one at a time yields shockingly little. Eliminate cancer entirely, and you would add only about two to four years to average life expectancy, because heart disease, dementia, and the rest are waiting right behind it. The real prize is to slow the *shared root* — aging itself — and thereby push back *all* of those diseases at once. That goal — more healthy years, fewer sick ones — is

genuinely plausible on the evidence, even where a routine hundred-and-twenty is not. It is the honest hope this book can offer: not that you'll live to a hundred and twenty, but that the science is real enough that you might, plausibly, stay *healthier* for longer than the generation before you.

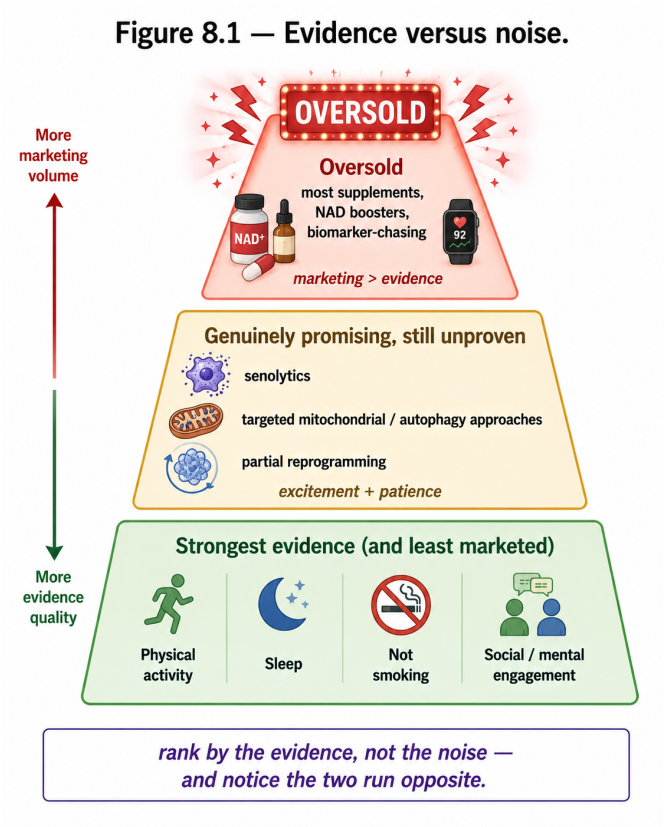
What actually works

Which brings us to the question every reader eventually asks: so what actually *works?* Here the most honest answer is also the least glamorous, and it deserves to be said plainly. The interventions with the *strongest* evidence for a longer, healthier life are not in any bottle sold to you as a longevity breakthrough. They are the boring basics — regular physical activity above all, adequate sleep, not smoking, staying socially and mentally engaged, a reasonable diet, and maintaining a healthy weight and metabolic health. These are unsexy, unpatentable, and unprofitable, which is precisely why they are drowned out by the marketing for things that work far less well, or not at all. (And recall a sharp lesson from Chapter 5: don't sabotage the things that work — the finding that high-dose antioxidant supplements can *erase* the benefits of exercise is a vivid warning that "more intervention" is not always better.)

Beyond the basics sits a tier of genuinely *promising but unproven* science — the senolytics of Chapter 3, the targeted mitochondrial and proteostasis approaches, the reprogramming

of Chapter 7 — worth real excitement and real patience, none of it yet a proven human therapy. And there is a large, loud tier of the *oversold* — most supplements, the NAD boosters of Chapter 4, the endless biomarker-chasing — where the marketing has vastly outrun the evidence. Rank by the quality of the evidence, not the volume of the marketing, and the two, as you've seen all book long, are very nearly inverted. (This is a map of where the evidence stands, which is a different and more durable thing.)

Figure 8.1 — Evidence versus noise.



That ranking is possible only because you now hold the tools to do it — and those tools, not any single fact, are the real thing this book set out to give you. They will outlast every specific result in these pages and serve you for every aging headline you ever read. Gather them:

Does it survive the jump from the mouse to the human? Most astonishing results are in animals, and the gap to humans is where most claims die.

Is this a real lever, or just a marker someone moved? A measurement that predicts aging is not a handle that controls it. "We can measure it" is not "you should act on it." A thermometer is not a thermostat.

Compared to what? The control group makes or breaks a study; most whiplash-inducing health news turns out to be two honest studies with different comparisons.

Is more actually better — or is there a dose? Biology runs on curves, not lines; a little of a stressor can help while a lot harms, and the obvious "more is better" is often exactly wrong.

What would it take to actually believe this? A large, randomized, placebo-controlled trial, measuring an outcome that matters, replicated by people with no stake in the result — and short of that, hold your enthusiasm in reserve.

Who benefits if I believe this? In a field this saturated with money, follow the incentives — and note that the conflicts of interest run in every direction, including the skeptics'.

Carry those six questions and you can navigate this entire field — and the next one, and the one after that — without being easily fooled in either direction. That is a more valuable possession than any answer about telomeres or clocks, because the answers will keep changing and the questions will not.

Amazed and careful

A final word about the strange moment in which this book is written, because it asks you to hold two true things at once that are usually presented as opposites.

The first is that we are living through a *genuine scientific revolution.* In a single human lifetime, aging went from unalterable fate to a regulated, studyable, in-some-ways-adjustable biological process. We can now read a cell's age, clear its damaged components, reset its clock in a dish, and run the earliest human trials of ideas that would have been science fiction a generation ago. That is real, and it is extraordinary, and it is easy to lose under the noise.

The second true thing is that this same field is the most *hype-saturated* in all of medicine — a place where genuine breakthroughs and outright snake oil are sold side by side, often by the same people, in the same language, and where

the gap between what is claimed and what is known is wider than almost anywhere else in science.

Both of these are completely true at once. The science is real *and* the claims are oversold. The revolution is genuine *and* the timeline is a fantasy. Living wisely in this moment — as a patient, a citizen, a parent, perhaps one day a scientist — means holding both, refusing neither the wonder nor the skepticism. That, in the end, is the whole discipline of this book, and of good science itself: to be amazed and to be careful at the same time.

✦ ✦ ✦

So — will you live to a hundred and twenty? Almost certainly not; not as the world is, not with what we can do today. But you might, plausibly, live *healthier* for longer than your grandparents did, because some of this science is real and some of it is arriving. And you can now do something better than either hope or despair about it: you can read this field for the rest of your life with clear eyes — telling the lever from the marker, the mouse from the human, the breakthrough from the bottle — and be far harder to fool in either direction.

And if you are one of the readers this book was secretly written for — the ones who finish it not satisfied but *hungry*, who want to go and do this work themselves — then the most important thing to know is that the field needs exactly that. Not more hype, and not more cynicism, but more

people who can tell the difference. The questions in these pages are not closed. Some of them are waiting for you.

Sources and Further Reading

WHERE TO GO NEXT — THE KEY PAPERS BEHIND EACH *chapter.*

This book is written as a continuous story rather than a string of footnotes, which is the right form for its purpose but leaves a debt: the claims in these pages rest on real, specific, often beautiful pieces of published science, and a reader who wants to check them, or simply to go deeper, deserves to know where to look. What follows is that map — not an exhaustive bibliography, but a curated guide to the landmark papers and primary sources each chapter is built on, the doors into the literature. Most are findable by author and year; many of the older ones are freely available, and the foundational ones are worth reading in the original at least once, because the way a discovery is first reported tells you things a textbook summary never will.

A word of method, in the spirit of the book: when you read one of these papers, ask it the same questions you've learned to ask the headlines. *Was this in a mouse or a human? Is the result a moved marker or a changed outcome? Compared to*

what? What would it take to believe it? The primary literature is not sacred — it is where the arguments actually happen, including the ones that later turned out to be wrong. Reading it critically is the whole point.

Introduction and framework

López-Otín, C., Blasco, M.A., Partridge, L., Serrano, M., & Kroemer, G. (2013). *The hallmarks of aging.* Cell. — The paper that organized the field into the framework this book follows.

López-Otín, C., Blasco, M.A., Partridge, L., Serrano, M., & Kroemer, G. (2023). *Hallmarks of aging: An expanding universe.* Cell. — The updated, expanded framework.

Olshansky, S.J., Willcox, B.J., Demetrius, L., & Beltrán-Sánchez, H. (2024). *Implausibility of radical life extension in humans in the twenty-first century.* Nature Aging. — The central evidence behind this book's answer to the hundred-and-twenty question: that average life expectancy in the longest-lived countries is leveling off, not climbing toward the hundreds.

Oeppen, J., & Vaupel, J.W. (2002). *Broken limits to life expectancy.* Science. — The optimists' classic: record life expectancy has risen steadily for over a century, with every predicted ceiling eventually broken.

Dong, X., Milholland, B., & Vijg, J. (2016). *Evidence for a limit to human lifespan.* Nature. — The other side of the maximum-lifespan debate: an argued ceiling near 115–125. Read

alongside the published rebuttals to see a genuine scientific disagreement in motion.

Chapter 1 — Telomeres

Hayflick, L., & Moorhead, P.S. (1961). *The serial cultivation of human diploid cell strains.* Experimental Cell Research. — The paper that overturned Carrel and established that normal cells count their divisions.

Olovnikov, A.M. (1973). *A theory of marginotomy.* Journal of Theoretical Biology. — and Watson, J.D. (1972). *Origin of concatemeric T7 DNA.* Nature New Biology. — The end-replication problem, identified independently.

Blackburn, E.H., & Gall, J.G. (1978). *A tandemly repeated sequence at the termini of the extrachromosomal rRNA genes in Tetrahymena.* Journal of Molecular Biology. — The pond ciliate gives up the telomere repeat.

Szostak, J.W., & Blackburn, E.H. (1982). *Cloning yeast telomeres on linear plasmid vectors.* Cell. — The experiment that proved telomeres are a universal, transferable element.

Greider, C.W., & Blackburn, E.H. (1985). *Identification of a specific telomere terminal transferase activity in Tetrahymena extracts.* Cell. — The discovery of telomerase (the "Christmas gel"). Blackburn, Greider, and Szostak shared the 2009 Nobel Prize for this line of work.

Jaskelioff, M., et al. (2011). *Telomerase reactivation reverses tissue degeneration in aged telomerase-deficient mice.* Nature. —

The dramatic mouse "rejuvenation" — and a case study in why the engineered-deficiency caveat matters.

Codd, V., et al. (2021). *Polygenic basis and biomedical consequences of telomere length variation.* Nature Genetics. — Large-scale human genetics of telomere length, including the trade-off in which longer telomeres lower some risks and raise cancer risk.

Armanios, M., & Blackburn, E.H. (2012). *The telomere syndromes.* Nature Reviews Genetics. — How telomere-maintenance mutations cause premature aging in high-turnover tissues — the proof that telomeres matter in humans.

Chapter 2 — Epigenetic clocks

Horvath, S. (2013). *DNA methylation age of human tissues and cell types.* Genome Biology. — The multi-tissue clock: 353 sites, an elastic-net regression, near-universal across tissues.

Hannum, G., et al. (2013). *Genome-wide methylation profiles reveal quantitative views of human aging rates.* Molecular Cell. — The contemporaneous blood-based clock.

Marioni, R.E., et al. (2015). *DNA methylation age of blood predicts all-cause mortality in later life.* Genome Biology. — Age acceleration as a mortality signal.

Levine, M.E., et al. (2018). *An epigenetic biomarker of aging for lifespan and healthspan.* Aging. — PhenoAge: the second-generation move to training on health rather than the calendar.

Lu, A.T., et al. (2019). *DNA methylation GrimAge strongly predicts lifespan and healthspan.* Aging. — GrimAge: training directly on mortality.

Belsky, D.W., et al. (2022). *DunedinPACE, a DNA methylation biomarker of the pace of aging.* eLife. — The third-generation shift from a state to a rate.

Yang, J.-H., et al. (2023). *Loss of epigenetic information as a cause of mammalian aging.* Cell. — The ICE-mouse experiments behind the "information theory of aging."

Fahy, G.W., et al. (2019). *Reversal of epigenetic aging and immunosenescent trends in humans.* Aging Cell. — The small, uncontrolled, much-cited TRIIM trial — the canonical "moved a clock ≠ reversed aging" cautionary case.

Zou, H., & Hastie, T. (2005). *Regularization and variable selection via the elastic net.* Journal of the Royal Statistical Society B. — For the reader who wants the math of the clock's engine.

Chapter 3 — Cellular senescence

Campisi, J. (2013). *Aging, cellular senescence, and cancer.* Annual Review of Physiology. — The reframing of senescence from a curiosity into a driver of aging.

Coppé, J.-P., et al. (2008). *Senescence-associated secretory phenotypes…* PLoS Biology. — The SASP, characterized.

Baker, D.J., et al. (2011). *Clearance of p16^Ink4a-positive senescent cells delays ageing-associated disorders.* Nature. — and

Baker, D.J., et al. (2016). *Naturally occurring p16^Ink4a-positive cells shorten healthy lifespan.* Nature. — The INK-ATTAC mice: clearing senescent cells rejuvenates, first in a progeroid model, then in naturally aged mice.

Xu, M., et al. (2018). *Senolytics improve physical function and increase lifespan in old age.* Nature Medicine. — The causal complement to the clearance experiments: transplanting small numbers of senescent cells into young mice induces physical dysfunction and spreads senescence — and senolytics reverse it.

Zhu, Y., et al. (2015). *The Achilles' heel of senescent cells: from transcriptome to senolytic drugs.* Aging Cell. — The first senolytics (dasatinib + quercetin), derived from the anti-apoptotic-pathway logic.

Demaria, M., et al. (2014). *An essential role for senescent cells in optimal wound healing.* Developmental Cell. — Why senescence is beneficial in its acute form — the basis for the "don't clear all of them" caution.

Hickson, L.J., et al. (2019). *Senolytics decrease senescent cells in humans…* EBioMedicine. — First evidence that senolysis is achievable in people (in diabetic kidney disease).

Williams, G.C. (1957). *Pleiotropy, natural selection, and the evolution of senescence.* Evolution. — Medawar, P.B. (1952). *An Unsolved Problem of Biology.* — and Kirkwood, T.B.L. (1977). *Evolution of ageing.* Nature. — The three classic evolutionary theories of why aging exists (antagonistic pleiotropy, mutation accumulation, disposable soma).

(The Unity Biotechnology UBX0101 osteoarthritis Phase 2 result is documented in the company's 2020 disclosures and the trial record rather than a primary paper.)

Chapter 4 — Nutrient sensing

Kenyon, C., Chang, J., Gensch, E., Rudner, A., & Tabtiang, R. (1993). *A C. elegans mutant that lives twice as long as wild type.* Nature. — The worm that converted aging from fate into a regulated mechanism.

Harrison, D.E., et al. (2009). *Rapamycin fed late in life extends lifespan in genetically heterogeneous mice.* Nature. — The Interventions Testing Program result: the field's most reproducible drug, working even when started in old age.

McCay, C.M., Crowell, M.F., & Maynard, L.A. (1935). *The effect of retarded growth upon the length of life span and upon the ultimate body size.* Journal of Nutrition. — The origin of caloric restriction.

Colman, R.J., et al. (2009). *Caloric restriction delays disease onset and mortality in rhesus monkeys.* Science. — Mattison, J.A., et al. (2012). *Impact of caloric restriction on health and survival in rhesus monkeys.* Nature. — and Mattison, J.A., et al. (2017). *Caloric restriction improves health and survival of rhesus monkeys.* Nature Communications. — The two famous monkey studies that appeared to contradict each other, and the joint reanalysis that resolved them — the chapter's "compared to what?" set-piece. Read all three together.

Howitz, K.T., et al. (2003). *Small molecule activators of sirtuins extend Saccharomyces cerevisiae lifespan.* Nature. — Baur, J.A., et al. (2006). *Resveratrol improves health and survival of mice on a high-calorie diet.* Nature. — and Pacholec, M., et al. (2010). *SRT1720, SRT2183, SRT1460, and resveratrol are not direct activators of SIRT1.* Journal of Biological Chemistry. — The rise and fall of resveratrol: the claims, and the assay-artifact debunk.

Rajman, L., Chwalek, K., & Sinclair, D.A. (2018). *Therapeutic potential of NAD-boosting molecules.* Cell Metabolism. — The case for NAD^+ precursors, from a leading proponent — read critically.

Holzenberger, M., et al. (2003). *IGF-1 receptor regulates lifespan and resistance to oxidative stress in mice.* Nature. — and Willcox, B.J., et al. (2008). *FOXO3A genotype is strongly associated with human longevity.* PNAS. — The insulin/IGF and FOXO pathway, from worm to mouse to human centenarian.

Waziry, R., et al. (2023). *Effect of long-term caloric restriction on DNA methylation measures of biological aging in healthy adults (CALERIE).* Nature Aging. — Human caloric restriction slowing a pace-of-aging clock.

Barzilai, N., Crandall, J.P., Kritchevsky, S.B., & Espeland, M.A. (2016). *Metformin as a tool to target aging.* Cell Metabolism. — The rationale for the TAME trial.

Harman, D. (1956). *Aging: a theory based on free radical and radiation chemistry.* Journal of Gerontology. — and Harman, D. (1972). *The biologic clock: the mitochondria?* Journal of the American Geriatrics Society. — The free-radical theory and its mitochondrial refinement, in the originator's own words.

McCord, J.M., & Fridovich, I. (1969). *Superoxide dismutase: an enzymic function for erythrocuprein.* Journal of Biological Chemistry. — The discovery that cells have dedicated antioxidant machinery.

The Alpha-Tocopherol, Beta Carotene Cancer Prevention Study Group (1994). *The effect of vitamin E and beta carotene on the incidence of lung cancer…* New England Journal of Medicine. — and Omenn, G.S., et al. (1996). *Effects of a combination of beta carotene and vitamin A on lung cancer and cardiovascular disease (CARET).* New England Journal of Medicine. — The two trials in which an antioxidant given to *prevent* lung cancer *increased* it. The most important data in the chapter.

Klein, E.A., et al. (2011). *Vitamin E and the risk of prostate cancer (SELECT).* JAMA. — and Bjelakovic, G., et al. (2007). *Mortality in randomized trials of antioxidant supplements for primary and secondary prevention.* JAMA. — More of the antioxidant graveyard, and the meta-analysis showing no benefit and some harm.

Ristow, M., et al. (2009). *Antioxidants prevent health-promoting effects of physical exercise in humans.* PNAS. — The elegant

experiment showing antioxidants abolish the benefits of exercise — mitohormesis in action.

Trifunovic, A., et al. (2004). *Premature ageing in mice expressing defective mitochondrial DNA polymerase.* Nature. — and Kujoth, G.C., et al. (2005). *Mitochondrial DNA mutations, oxidative stress, and apoptosis in mammalian aging.* Science. — The mutator mouse: premature aging from mtDNA mutations, but *without* the runaway ROS the theory predicted.

Andziak, B., et al. (2006). *High oxidative damage levels in the longest-living rodent, the naked mole-rat.* Aging Cell. — The wrinkled refutation of "oxidative damage sets lifespan."

Schriner, S.E., et al. (2005). *Extension of murine life span by overexpression of catalase targeted to mitochondria.* Science. — The modest exception that proves how disappointing the antioxidant-overexpression experiments mostly were.

Andreux, P.A., et al. (2019). *The mitophagy activator urolithin A is safe and induces a molecular signature of improved mitochondrial and cellular health in humans.* Nature Metabolism. — The shift from "mop up ROS" to "improve quality control."

Chapter 6 — Proteostasis

Anfinsen, C.B. (1973). *Principles that govern the folding of protein chains.* Science. — The Nobel lecture: structure is encoded in sequence.

Levinthal, C. (1969). *How to fold graciously.* (In *Mössbauer Spectroscopy in Biological Systems*, proceedings.) — The paradox of folding speed.

Hartl, F.U., Bracher, A., & Hayer-Hartl, M. (2011). *Molecular chaperones in protein folding and proteostasis.* Nature. — The chaperone machinery, reviewed.

Ohsumi, Y. (2014). *Historical landmarks of autophagy research.* Cell Research. — Autophagy from its discoverer (2016 Nobel laureate).

Prusiner, S.B. (1982). *Novel proteinaceous infectious particles cause scrapie.* Science. — The radical, Nobel-winning prion idea — a self-templating misfolded shape.

Hardy, J.A., & Higgins, G.A. (1992). *Alzheimer's disease: the amyloid cascade hypothesis.* Science. — The hypothesis that organized — and arguably distorted — three decades of Alzheimer's research.

van Dyck, C.H., et al. (2023). *Lecanemab in early Alzheimer's disease.* New England Journal of Medicine. — and Sims, J.R., et al. (2023). *Donanemab in early symptomatic Alzheimer's disease (TRAILBLAZER-ALZ 2).* JAMA. — The first anti-amyloid drugs to slow decline — a modest, risky, genuinely real effect that both partly vindicates and partly indicts the amyloid hypothesis.

Jumper, J., et al. (2021). *Highly accurate protein structure prediction with AlphaFold.* Nature. — Varadi, M., et al. (2022). *AlphaFold Protein Structure Database…* Nucleic Acids Research.

— and Abramson, J., et al. (2024). *Accurate structure prediction of biomolecular interactions with AlphaFold3.* Nature. — The solving of protein-structure prediction (Hassabis and Jumper, 2024 Nobel Prize in Chemistry), the database, and the extension toward interactions — read with the chapter's footnote about what it does *not* predict.

Rubinsztein, D.C., Mariño, G., & Kroemer, G. (2011). *Autophagy and aging.* Cell. — Why the cleanup crews falling behind is central to aging.

Chapter 7 — Reprogramming

Gurdon, J.B. (1962). *The developmental capacity of nuclei taken from intestinal epithelium cells of feeding tadpoles.* Journal of Embryology and Experimental Morphology. — The nuclear-transfer experiments showing an adult cell retains all the instructions for a whole organism.

Takahashi, K., & Yamanaka, S. (2006). *Induction of pluripotent stem cells from mouse embryonic and adult fibroblast cultures by defined factors.* Cell. — The four-factor reset (2012 Nobel Prize, shared with Gurdon): proof that cell identity is reversible.

Lapasset, L., et al. (2011). *Rejuvenating senescent and centenarian human cells by reprogramming through pluripotency.* Genes & Development. — Reprogramming resets the markers of cellular age toward zero.

Ocampo, A., et al. (2016). *In vivo amelioration of age-associated hallmarks by partial reprogramming.* Cell. — The proof of concept that *partial*, cyclic reprogramming can rejuvenate a living mammal without teratomas.

Lu, Y., et al. (2020). *Reprogramming to recover youthful epigenetic information and restore vision.* Nature. — Partial reprogramming (OSK) restoring vision in mice — the strongest demonstration of restored function, and the science behind the first human trial.

Conboy, I.M., et al. (2005). *Rejuvenation of aged progenitor cells by exposure to a young systemic environment.* Nature. — and Mehdipour, M., et al. (2020). *Rejuvenation of three germ layers' tissues by exchanging old blood plasma with saline-albumin.* Aging. — Heterochronic parabiosis, and the later evidence that diluting old blood matters more than adding young.

Hou, P., et al. (2013). *Pluripotent stem cells induced from mouse somatic cells by small-molecule compounds.* Science. — and Guan, J., et al. (2022). *Chemical reprogramming of human somatic cells to pluripotent stem cells.* Nature. — Reprogramming with drugs instead of genes — the route most amenable to AI-driven discovery.

(The first human trial of partial epigenetic reprogramming — a Phase 1 study for an eye disease — was cleared by the U.S. FDA in early 2026; see the sponsor's announcements and the public trial registry for current status.)

Olshansky, S.J., et al. (2024). *Implausibility of radical life extension in humans in the twenty-first century.* Nature Aging. — The core of the chapter's answer (and listed under the Introduction above).

Goldman, D.P., et al. (2013). *Substantial health and economic returns from delayed aging may warrant a new focus for medical research.* Health Affairs. — The geroscience case: why targeting aging itself beats fighting one disease at a time, and why curing any single disease buys surprisingly little.

López-Otín, C., et al. (2013; 2023). *The hallmarks of aging / Hallmarks of aging: An expanding universe.* Cell. — The synthesis to return to.

✦ ✦ ✦

A closing note. This list is a starting point, not a boundary. Every paper here is embedded in a web of others — the studies it built on, the ones that challenged it, the ones that came after. Follow those threads. The most important thing these references can teach you is not any single result but the texture of how the work is actually done: provisional, argued-over, self-correcting, and far more honest with itself than the headlines that summarize it. That texture is the real subject of this book.