

The First Second of Infinity

A Collection of Essays

Rodrigo Reyes Feregrino

“But it might take me a little while, so do you want me to tell you a story? The Brothers Grimm, lovely fellas. They’re on my darts team. According to them, there’s this emperor and he asks this shepherd’s boy, how many seconds in eternity? And the shepherd’s boy says, there’s this mountain of pure diamond. It takes an hour to climb it, and an hour to go around it. Every hundred years, a little bird comes and sharpens its beak on the diamond mountain. And when the entire mountain is chiselled away, the first second of eternity will have passed. You must think that’s a hell of a long time. Personally, I think that’s a hell of a bird.”

— AFTER THE BROTHERS GRIMM

(AND ONE OF THEIR DARTS TEAMMATES)

Contents

Foreword	vii
I Of Diamonds and Dust	1
“Lifespan” Is Misleading	3
Why Humans Are Not Immortal	4
Evolution Has No Direction	5
The Grandmother Hypothesis	7
Where Now?	9
The Twelve Horsemen	11
A New Perspective on Aging	11
Impact and Evolution of the Framework	12
How Does the End Come?	13
Genomic Instability	13
Telomere Attrition	14
Cellular Senescence	14
Epigenetic Alterations	15
The Cascade	15
The Path Forward	17
Sources and Further Reading	17
Tic Toc Goes... Which Clock?	19
DNA Methylation and How Scientists Measure It	20
The Horvath Clock	22
Different Clocks, Different Questions	22
The Machine Learning Behind the Clocks	23
The Commercial Clock	25

CONTENTS

The Conquest Begins 26

II Of Bird Beaks and Bargains 28

The Fluor de Lys Affair 30

Act I: The Scene of the Crime 30

Interlude: What Are Sirtuins, and Why Should Anyone Care? 31

Act II: The First Cracks in the Mirror 33

Act III: The Pfizer Papers 33

Act IV: Whodunit? 35

Act V: The Accomplices 37

Coda: The \$720 Million Footnote 38

The Second Foundation 42

What Cellular Reprogramming Actually Is: the Yamanaka
Factors 43

The 2020 Paper That Made the Cover 45

ER-100: The First Human Trial of Partial Reprogramming . 49

Interlude: Five Days in July 50

March 2024: The Cascade 53

What ER-100 Actually Promises, and Doesn't 54

The Merchant's Century 58

The Genome Racer 58

The Engineer of Negligible Senescence 59

The Billionaire's Dilemma 60

The Quantified Selves 61

The Boring Truth 63

A Brisk Trade in Beaks, and the Conquest Deferred 63

A More Surgical Scalpel 67

What Cellular Senescence Is, and Why You Want It (Some-
times) 67

PROTACs 67

UNITY Biotechnology and its mess	67
Why Any of This Has to Do with Aging	68
Why Computers, Why Now	69
All Theories are Wrong, but Quantum Improvement is Real	73
What Quantum Computation Actually Is	73
The Pen That Writes on Every Line	75
The Input Problem and the Output Problem	78
The Quantum Decoherence Problem	79
Feynman to Shor: A Forty-Five-Year Argument	80
What's Real	80
What's Noise	81
What's Probably Wrong	82
The Shor Gap	83
Every Hundred Years	84
Quantum Everything? No. Quantum Biochemistry? Yes.	86
The Hype	86
The Sharpening Tools	87
AlphaFold: Structure From Sequence	87
How AlphaFold Works	88
ESM: Learning Protein Grammar	89
Designing What Evolution Never Tried	90
Aging: A Plausible Frontier	91
Targeting Zombie Cells	91
The Loop, Not the Leap	92
What Changed	92
The Administrative Revolution	95
Coda: That's a Hell of a Bird	96
Author's Note	97

CONTENTS

Acknowledgments

98

Foreword

THERE are many challenges sciences takes on, but arguably the greatest and most important one is the conquest of death. If not the end goal of science, then at least it is, I think, a motivation upstream of most of it. We are apes who are scared of dying, who prefer pleasure to pain, and who figured out somewhere along the way that controlling the external world is a reliable way to get more of the former and less of the latter. Science claims to seek understanding, but understanding is useful mainly because it permits prediction, and prediction is useful because it permits control: control over disease; control over suffering; control over how long a frightened ape gets to remain alive.

"But..." I hear you saying "what about scientific curiosity? What about the *love* for elegant theories? for mathematical proofs?" I would say these are also downstream effects of behaviours that gave some apes in Africa three hundred thousand years ago a slightly better shot at outrunning the thing that was going to kill them anyway; they are adaptations, and adaptations do not have to feel like adaptations from the inside. The ape who finds a proof beautiful is still the ape who once found a ripe fig beautiful. It is the same machinery, pointed at a different layer of the world.

The understanding part of science, in this view, is mostly a side effect. We build models—abductive scaffolding bolted to a restrictive set of assumptions—and the models hold until they don't, and then we build new ones that hold until they don't either. The point is that the models are stepping stones. What they buy us, in the end, is control. Comfort. Healthspan. A longer afternoon on the savannah.

Thus, that is our diamond mountain. Our route out of the savannah. The first second of eternity. Our attempt to access infinity with ape hardware. Whatever sits on the other side of it—call it longevity, transcendence, the end of disease, whatever the keynote speaker is calling it this year—is what the whole enterprise is, in some sense, working toward. After that first second has passed, who knows.

But the collection is named after the bird, not the mountain. Every paper published, every undergrad learning how to pipette, every new drug, every faster algorithm, every new instrument: we are the bird - and that's one hell of a bird. Scientists, the human spirit, the *Zeitgeist*, whatever you want to call it.

The bird returns every century and makes another mark. Most are invisible. Every so often the mountain cracks audibly. Penicillin. DNA. PCR. AlphaFold. We hear the crack, we cheer, and briefly forget that the mountain is made of diamond. Forgetting is what produces the press release, the supplement bottle, the eleven-figure valuation, the keynote speaker announcing that eternity has entered beta testing. Sagan's Baloney Detection Kit exists for exactly this reason: skepticism is easiest to apply to ideas you dislike, and hardest to apply to the ones you desperately want to be true.

Half of what is currently being promised about longevity science—and about the new computational tools now orbiting it, especially artificial intelligence and quantum computing—is hype. However, the chiselling underneath is real, and slow, and it is worth watching even though most of us will not live to see the second tick over.

The collection does not argue that the buzzwords are empty. It argues that beneath the press releases, supplement companies, valuations, and conference keynotes there usually exists a real scientific core. Longevity science is the central problem. Artificial intelligence

and quantum computing enter the story as possible instruments: new ways of modelling biology, searching chemical space, simulating physical systems, or compressing decades of trial-and-error into something faster. Whether they will actually do this remains unclear. The work of separating the real chiselling from the noise surrounding it has turned out to be more interesting to me than either cynicism or evangelism.

Part I treats the biology of aging on its own terms: what evolution does and does not select for, what the twelve hallmarks of aging are, and how epigenetic clocks measure something we did not know how to measure twenty years ago. Part II treats the marketplace that has formed around that biology—the Sirtris arc, the Sinclair arc, the senolytics arc, the venture capital that has so far failed to convert into approved drugs. Part III treats the two emerging technologies whose promoters currently claim, with varying degrees of evidence, to be about to change everything: artificial intelligence and quantum computing.

A note on what I am bringing to this. I spent two years helping out with quantum computing research. I have a chemistry major, a computer science major, and a deeply skeptical mind. I am also, like everyone else, a scared ape who would prefer not to die. There is a real temptation, when you are working on one of these problems, to believe that the buzzword in the next building over is going to solve yours. I have felt it. The essays that follow are partly an attempt to talk myself out of it.

FOREWORD

PART I

Of Diamonds and Dust

Mortality as a Scientific Problem

“Lifespan” Is Misleading

An Evolutionary Perspective on Aging

Imagine being born into an aristocratic family in ancient Rome, over 2000 years ago. If you survived childhood, you could generally expect to live well into your seventies. Indeed, a demographic analysis by historian Walter Scheidel at Stanford University shows that aristocratic Romans lived as long as modern humans - once past the hurdle of early mortality.

We can even go further back in time; pop science media love claiming that ancient humans (think late neolithic/ early bronze age) would all die by age 30. This is not true. While paleodemography is less gracious than modern historical estimates when it comes to providing sharp quantitative results for lifespan distribution amongst ancient humans, one thing is clear: adults who survived childhood regularly lived into their fifties, sixties, and beyond. J. Lawrence Angel of the Smithsonian Institution, analyzing skeletal remains across prehistoric Eastern Mediterranean populations, and later Boldsen and Paine, surveying eight thousand years of European skeletal data from the Mesolithic to the Middle Ages, confirm the pattern: the upper bound of human lifespan simply did not change.

Their maximum lifespan was identical to ours. Yet, it is quite common to hear claims that “science has increased human lifespan” or “being 40 was considered quite old in ancient times.” Journalist Christopher Wanjek pinpointed the issue: “Has the human race increased its life span? Not at all.” The confusion stems from conflating life expectancy with maximum lifespan.

Charles Brenner, Chair of the Department of Diabetes Cancer Metabolism at the Beckman Research Institute, puts it bluntly: "... all vertebrate animal species have a distribution of natural lifespans that are limited by their gene sets—human longevity appears top out at about 120 years".

Why Humans Are Not Immortal

This upper bound on "natural" longevity, even ignoring premature death, makes sense when seen from an evolutionary perspective. Brenner explains it through the concept of selective pressure: "animals in the wild are under little to no direct genetic selection for longevity beyond that to produce reproductive success."

Imagine a predator that hunts humans with perfect lethality, and this predator lives everywhere, like mosquitoes or ticks. But here is the catch: it only kills people over 120 years old. If some people carry genes making them invisible to this predator, would these genes spread through the population? Would bearers of this genetic immunity have more children, more grandchildren?

No.

By age 120, reproduction has long since concluded. Any direct or indirect contribution you could make to the replication of your genes has become negligible by this point (more on this later). Natural selection cannot see threats that appear only after the evolutionary game has ended.

This is why animals rarely live much longer than the time required to reproduce and raise offspring to independence. Maintaining a body demands enormous resources; DNA repair, cellular cleanup, protein quality control—these biological maintenance systems are expensive.

The human body in particular seems to be quite hard to maintain. The human brain, for instance, represents 2% of body weight but consumes 20–25% of resting metabolic energy—roughly 260 to 500 calories daily. That’s double what chimpanzee brains require, and up to ten times more, relative to body size, than other land mammals.

A 2024 study in the *Proceedings of the National Academy of Sciences* by Yegian, Lieberman, and colleagues used a new metabolic-scaling method—comparing observed energy expenditures to those predicted for the average mammal corrected for size, body composition, and environmental temperature—and showed that humans evolved *both* resting and active metabolic rates that substantially exceed expectations for our body mass. Among nonindustrial human populations the elevation runs to 40–85% above expected total energy expenditure, while chimpanzees sit at the mammalian average. Most strikingly, humans appear to have escaped the tradeoff between resting and active metabolism that constrains other primates: we are uniquely expensive to maintain *and* expensive to move.

Why would evolution invest in indefinitely maintaining such an expensive machine when most individuals die from external causes long before cellular maintenance failures matter?

Evolution Has No Direction

The example above requires some extra nuance. Evolution doesn’t “want” anything. It has no goals, no direction, no inherent, *a priori* preference for longer or shorter lifespans. It just happens to be the case that immortality tends to be a bad bet.

Richard Dawkins, the evolutionary biologist at Oxford – and the person who first coined the term “meme” – crystallized this

perspective in his book "The Selfish Gene." From the gene's perspective, organisms are temporary vehicles: survival machines built to transport DNA into the next generation. Once that job is done, the vehicle becomes expendable. Natural selection doesn't care about the individual organism's longevity *per se*; it cares only about genetic replication. As Dawkins puts it, we are "lumbering robots" built by genes to ensure *their* survival, not ours.

The chicken is only an egg's way for making another egg
- Richard Dawkins

This highlights another way we should *not* think about aging and lifespans: our bodies aren't programmed to die "for the good of the species." Rather, genes that maintain bodies indefinitely simply don't – *usually* (more on this later) - out-compete genes that invest those same resources in reproduction.

Thus, animal lifespans span an astonishingly large range. The shortest living animals, the female American sand-burrowing mayflies, live just five minutes as adults, whilst an ocean quahog clam nicknamed "Ming" lived 507 years. Between these extremes lies remarkable diversity: gastrotrichs, microscopic aquatic worms, complete their entire lives in 3-21 days; Labord's chameleon in Madagascar lives only 4-5 months after hatching; bowhead whales can exceed 200 years; Greenland sharks cruise the Arctic for 400 years.

Perhaps unsurprisingly, immortality has indeed emerged as an evolutionarily stable strategy a few times. The most famous example is *Turritopsis dohrnii*, the so-called immortal jellyfish—a creature smaller than a pinky nail that can reverse its life cycle indefinitely. When stressed by injury, starvation, or age, this Mediterranean jellyfish sinks to the seafloor, retracts its tentacles, and transforms back into a juvenile polyp.

The exact mechanism behind *Turritopsis dohrnii*'s immortality was understood recently, and while it lies beyond the scope of this article, the curious reader should venture to read the fantastic "Genome assembly and transcriptomic analyses of the repeatedly rejuvenating jellyfish *Turritopsis dohrnii*" paper.

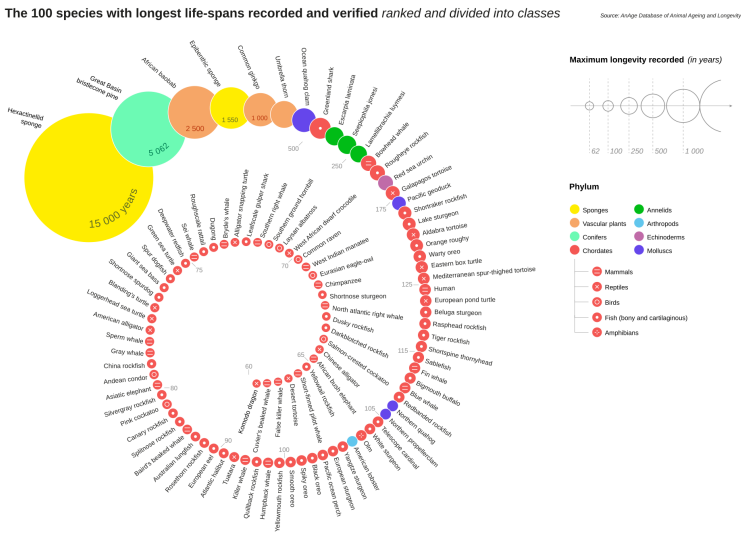


FIGURE 1. Nature’s longevity champions: Glass sponges live 15,000 years. Greenland sharks, 500. Humans barely crack 120. Evolution didn’t optimize for immortality—it optimized for reproduction. These wildly different lifespans reveal the arbitrary ceiling that natural selection placed on each species

The Grandmother Hypothesis

We established earlier that natural selection "stops caring" about organisms after reproduction, but that rule has an exception.

Consider killer whales. Female orcas stop reproducing in their 30s but live into their 90s. When a post-reproductive grandmother orca dies, her grandoffspring become 4.5 times more likely to die within two years. These grandmother whales lead their pods to

salmon runs using decades of accumulated knowledge. During food scarcity, they take charge—guiding their families to feeding grounds younger whales don’t know exist.

A similar pattern appears in humans. Analysis of 18th and 19th century Finnish church records revealed that children with living maternal grandmothers had 30% better survival odds. Among the Hadza people of Tanzania, grandmothers forage roughly 3,000 calories daily—and children’s growth correlates directly with grandmother’s work once mothers have new babies to nurse.

This is the grandmother hypothesis: post-reproductive survival can provide evolutionary advantages through helping descendants survive. This, however, is rare.

Only six mammal species out of more than 5,000 evolved true menopause with extended post-reproductive lifespans: humans, killer whales, short-finned pilot whales, false killer whales, belugas, and narwhals. This extreme rarity reveals how specific the conditions must be.

These species share key features.

- First, low extrinsic mortality: the rate at which predators and accidents kill regardless of biological condition.
- Second, social structures where female relatedness to group members increases with age. In killer whales, pods gradually fill with a matriarch’s descendants. In ancestral humans, young women typically left their natal groups, so older women became increasingly related to everyone around them.
- Third, accumulated knowledge must provide survival advantages that outweigh the metabolic costs. An orca grandmother leading her pod to salmon during lean years converts decades

of ecological memory into grandoffspring survival at rates no reproductive female could match.

Where Now?

Human lifespan is hard-capped by our genes at around 120 years. This ceiling exists not because evolution programmed us to die, but because evolution stopped maintaining us after reproduction. Genes that extend post-reproductive survival face little to no selective pressure, so they never spread.

Yet biological immortality is not impossible. Jellyfish achieve it; orcas and humans evolved loopholes through grandmother effects. The current human lifespan is not a biological inevitability—it's the arbitrary result of a blind watchmaker optimizing for reproduction, not longevity.

This realization transforms the problem. If aging isn't programmed but is instead accumulated damage from insufficient maintenance, then extending human lifespan becomes an engineering problem, not a fantasy. The molecular machinery exists. Evolution just never invested in maintaining it past age 50. Understanding exactly what breaks down—and why—is the first step toward fixing it.

Sources and Further Reading

- Angel, J. L. (1969). The bases of paleodemography. *American Journal of Physical Anthropology*, 30(3), 427–438. <https://doi.org/10.1002/ajpa.1330300314>
- Boldsen, J. L., & Paine, R. R. (1995). The evolution of human longevity from the Mesolithic to the Middle Ages: An analysis based on skeletal data. In B. Jeune & J. W. Vaupel (Eds.), *Exceptional longevity: From prehistory to the present* (Monographs on Population Aging 2, pp. 25–36). Odense University Press.
- Brenner, C. (2022). Sirtuins are not conserved longevity genes. *Life Metabolism*, 1(2), 122–133. <https://doi.org/10.1093/lifemeta/loac025>
- Croft, D. P., Brent, L. J. N., Franks, D. W., & Cant, M. A. (2015). The

evolution of prolonged life after reproduction. *Trends in Ecology & Evolution*, 30(7), 407–416. <https://doi.org/10.1016/j.tree.2015.04.011>

Dawkins, R. (1976). *The selfish gene*. Oxford University Press.

Hasegawa, Y., Watanabe, T., Otsuka, R., Toné, S., Kubota, S., & Hirakawa, H. (2023). Genome assembly and transcriptomic analyses of the repeatedly rejuvenating jellyfish *Turritopsis dohrnii*. *DNA Research*, 30(1), dsac047. <https://doi.org/10.1093/dnares/dsac047>

Karsten, K. B., Andriamandimbiarisoa, L. N., Fox, S. F., & Raxworthy, C. J. (2008). A unique life history among tetrapods: An annual chameleon living mostly as an egg. *Proceedings of the National Academy of Sciences*, 105(26), 8980–8984. <https://doi.org/10.1073/pnas.0802468105>

Scheidel, W. (1999). Emperors, aristocrats, and the Grim Reaper: Towards a demographic profile of the Roman élite. *Classical Quarterly*, 49(1), 254–281. <https://doi.org/10.1093/cq/49.1.254>

Wanjek, C. (2002). *Bad medicine: Misconceptions and misuses revealed, from distance healing to vitamin O*. Wiley.

Yegian, A. K., Heymsfield, S. B., Castillo, E. R., Müller, M. J., Redman, L. M., & Lieberman, D. E. (2024). Metabolic scaling, energy allocation tradeoffs, and the evolution of humans’ unique metabolism. *Proceedings of the National Academy of Sciences*, 121(48), e2409674121. <https://doi.org/10.1073/pnas.2409674121>

The Twelve Horsemen

A Diagnostic Manual for Mortality

Albert Camus, philosopher and Nobel Laureate, famously wrote that the only important philosophical question is whether one should commit suicide. Camus asked whether one should choose to die sooner. Scientists have inverted the question: can we choose to die later—or never?

Many attempts have been made throughout history. Medieval alchemists sought the philosopher's stone. Qin Shi Huang, China's first emperor, consumed mercury compounds in pursuit of immortality—and died at 49 from mercury poisoning. Ponce de León searched Florida for the Fountain of Youth. Countess Elizabeth Báthory allegedly bathed in the blood of virgins.

Now, modern biology offers something those quests lacked: a mechanistic blueprint. The trouble is that the blueprint has twelve separate failure modes, and no single intervention closes more than one or two of them. The reader who came here looking for “the cause of aging” is going to leave disappointed. The reader who came here looking for the *causes* (plural, interacting, mutually reinforcing) is in the right essay.

A New Perspective on Aging

In 2013, Carlos López-Otín and his colleagues at the University of Oviedo published what would become the most cited framework in aging research: the hallmarks of aging. This was something rarer and much more ambitious than another laboratory study generating new data, or even than a meta-study about the effects of some specific

treatment on some specific condition; rather, it was a comprehensive review synthesizing thousands of studies across molecular biology, genetics, and gerontology into a unified map of human deterioration.

The researchers asked a deceptively simple question: What processes, verified across multiple species and experimental systems, consistently drive aging? They demanded rigorous criteria:

- Each hallmark had to manifest during normal aging.
- Experimental intensification had to accelerate aging.
- Therapeutic intervention had to offer the possibility of slowing, stopping, or reversing the aging process.

These twelve hallmarks do not merely catalog symptoms, but rather map the causal architecture connecting molecular damage to organismal collapse—the cascading network of failures that transforms a robust twenty-year-old body into a fragile ninety-year-old one.

Impact and Evolution of the Framework

The framework detonated across the scientific landscape. Research programs reorganized around it. Biotechnology companies structured drug development pipelines according to its logic. By 2023, López-Otín and colleagues published an updated analysis titled "Hallmarks of Aging: An Expanding Universe." Their assessment: "Recent research has confirmed and extended the importance of all these hallmarks. They have withstood scrutiny by tens of thousands of aging researchers, but they require an update to deal with the discoveries of the last decade".

The 2023 update raised the count from nine to twelve, adding *disabled macroautophagy*, *chronic inflammation*, and *dysbiosis* (disruption of the gut microbiome) to the original list. None of these is a small addendum. Each is a systems-level phenomenon that interacts with

most of the other hallmarks—more musicians falling out of tempo with the orchestra, not new instruments joining it.

The twelve hallmarks had become the field's working blueprint for understanding—and potentially defeating—biological aging.

How Does the End Come?

Everyone faces their own personal apocalypse eventually; everyone faces the end of days. But how exactly does the end of days come (at the molecular level, of course)?

The human body at twenty operates like a classical orchestra. Every section knows its part. By fifty, one can hear a few missed notes, wrong tempos, and odd improvisations. By eighty? Monkeys banging on pipes. A comprehensive discussion of all twelve hallmarks is well beyond the scope of this article. However, here are some of the most noteworthy:

Genomic Instability

Your genome accumulates errors. Not occasionally. Constantly. Every cell division can and does produce transcription mistakes. Every. Single. One. Every metabolic breath generates reactive oxygen species—unstable molecules that damage DNA like microscopic bullets. According to López-Otín and colleagues, "DNA damage is the most important driver of aging."

The body deploys sophisticated repair mechanisms— with very fancy names such as base excision repair, nucleotide excision repair, mismatch repair, and double-strand break repair. These molecular maintenance crews patrol your genome constantly, identifying and correcting errors. But with age, these repair systems accumulate their own damage. Key repair proteins become less abundant. The machinery slows down. Error correction becomes less accurate. Mutations accumulate. The instruction manual for cellular function

becomes increasingly corrupted.

Telomere Attrition

Imagine the plastic tips on shoelaces—the aglets that prevent fraying. Telomeres serve this function for chromosomes. These repetitive DNA sequences cap chromosome ends, protecting genetic information during cell division. But there's a problem: whenever a cell divides, it must untangle all its DNA, copy it, and then re-tangle it. This process of DNA replication comes with a hidden cost at the chromosome ends—with each division, telomeres shorten.

Elizabeth Blackburn won the Nobel Prize for discovering telomerase, the enzyme that rebuilds telomeres (Greider Blackburn, 1985). Most adult cells lack sufficient telomerase activity. The result? A molecular countdown. After roughly fifty divisions—what is currently known as the Hayflick limit—telomeres become critically short. The chromosome frays. The cell either dies or enters senescence, a zombie state we'll examine shortly.

Cellular Senescence

Senescent cells stop dividing but remain metabolically active in a way that is actively harmful to surrounding cells.

Judith Campisi at the Buck Institute discovered that senescent cells secrete a cocktail of inflammatory cytokines, growth factors, and proteases—the senescence-associated secretory phenotype, or SASP (Coppé et al., 2008). This secretion has a purpose: in young, healthy individuals with robust immune systems, SASP acts as a distress signal, calling immune cells to destroy the damaged cell before it can cause problems.

The system works beautifully—when the immune response is fast and efficient. Think of it like a building with a burst pipe: SASP is the alarm that evacuates the building and summons emergency

crews. If the repair team arrives quickly, the damage is contained and everyone returns to normal function. But in older individuals with declining immune function, the cavalry never arrives. The senescent cells persist, continuing to pump out inflammatory signals. The metaphor is zombies: neither alive nor dead but spreading their condition to everything nearby.

As Campisi puts it, "Senescent cells are like bad neighbors in a good neighborhood." When the plumber never shows up, the entire building slowly degrades while displaced residents remain inefficient and unproductive. Neighboring cells become inflamed, dysfunctional, or senescent themselves.

Epigenetic Alterations

Your genome contains approximately 20,000 genes. At any moment, each cell type expresses only a fraction of these—liver cells activate liver genes, neurons activate neural genes. This control system operates through epigenetic marks: chemical modifications to DNA and histones that determine which genes activate when.

Aging scrambles this system catastrophically. According to Steve Horvath at UCLA, epigenetic patterns change so predictably with age that they function as a biological clock. The metaphor López-Otín uses: a library where books randomly jump between fiction and reference sections. The information remains intact, but retrieval becomes chaotic. Cellular identity blurs. Function deteriorates.

The Cascade

These mechanisms interconnect viciously. Genomic instability triggers senescence. Senescent cells pump out inflammation. Inflammation disrupts communication. Disrupted communication affects stem cells. Dysfunctional stem cells fail to maintain tissues. Mitochondria accumulate damage. Damaged mitochondria generate reactive oxygen species. Reactive oxygen species cause more ge-

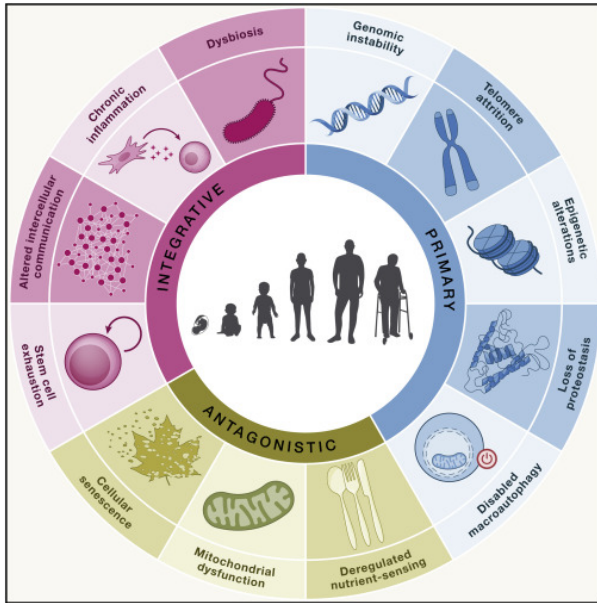


FIGURE 2. schematic of the 12 hallmarks of aging, from Dr. López-Otín’s latest paper.

Genomic instability.

Professor João Pedro de Magalhães, Chair of Molecular Biogerontology at the University of Birmingham, mapped these interactions computationally (Pun et al., 2022). Targeting one hallmark affects multiple others. The orchestra collapse is not twelve independent failures but a cascading network where each musician’s breakdown throws off everyone else’s timing.

López-Otín threads one implication through all twelve: "The opportunity to decelerate, stop, or reverse aging by therapeutic interventions on them." The diagnostic manual of mortality doubles as a strategic blueprint for its defeat. First, identify which section of the orchestra is failing. Then determine whether you can retrain the musicians, replace them, or at minimum stop them from disrupting everyone else’s performance.

The Path Forward

Aging might be the deepest expression of the absurdity in existence Camus describes: the gap between our desperate need for meaning and biology's complete indifference to whether we exist past our fertile years.

Camus argued for revolt as the response to absurdity. Perhaps the most profound revolt is understanding the mechanisms of our deterioration well enough to override them. The hallmarks of aging aren't existential truths requiring acceptance, they're mechanical failures inviting intervention. Camus might have found molecular biology the ultimate act of rebellion: not raging against mortality in the abstract, but dismantling it mechanism by mechanism, refusing to let evolutionary indifference have the final word.

Sources and Further Reading

Camus, A. (1991). *The myth of Sisyphus* (J. O'Brien, Trans.). Vintage Books. (Original work published 1942)

Coppé, J.-P., Patil, C. K., Rodier, F., Sun, Y., Muñoz, D. P., Goldstein, J., Nelson, P. S., Desprez, P.-Y., & Campisi, J. (2008). Senescence-associated secretory phenotypes reveal cell-nonautonomous functions of oncogenic RAS and the p53 tumor suppressor. *PLoS Biology*, *6*(12), 2853–2868. <https://doi.org/10.1371/journal.pbio.0060301>

Greider, C. W., & Blackburn, E. H. (1985). Identification of a specific telomere terminal transferase activity in *Tetrahymena* extracts. *Cell*, *43*(2 Pt 1), 405–413. [https://doi.org/10.1016/0092-8674\(85\)90170-9](https://doi.org/10.1016/0092-8674(85)90170-9)

Horvath, S. (2013). DNA methylation age of human tissues and cell types. *Genome Biology*, *14*(10), R115. <https://doi.org/10.1186/gb-2013-14-10-r115>

López-Otín, C., Blasco, M. A., Partridge, L., Serrano, M., & Kroemer, G. (2013). The hallmarks of aging. *Cell*, *153*(6), 1194–1217. <https://doi.org/10.1016/j.cell.2013.05.047>

López-Otín, C., Blasco, M. A., Partridge, L., Serrano, M., & Kroemer, G. (2023). Hallmarks of aging: An expanding universe. *Cell*, *186*(2), 243–278. <https://doi.org/10.1016/j.cell.2022.11.001>

Pun, F. W., Leung, G. H. D., Leung, H. W., Liu, B. H. M., Long, X., Ozerov, I. V., Wang, J., Ren, F., Aliper, A., Izumchenko, E., Moskalev, A., de Magalhães, J. P., & Zhavoronkov, A. (2022). Hallmarks of aging-based dual-purpose disease and age-associated targets predicted using PandaOmics AI-powered discovery engine. *Aging*, *14*(6), 2475–2506. <https://doi.org/10.18632/aging.203960>

Tic Toc Goes... Which Clock?

To Reverse Aging, One Must First Learn to Measure It

When does someone reach “old age”?

Fauja Singh ran the Toronto Waterfront Marathon in 5 hours and 40 minutes—beating several runners in their twenties—at what was reportedly age 92.¹ John B. Goodenough was awarded the Nobel Prize in Chemistry at age 97 and was still publishing research well into his nineties.

Clearly we mean something different than chronological age when we talk about being “old.” It is precisely this fuzzy measurement of biological age that current scientific research aims to stop—and even reverse.

But how exactly should we define biological age? What makes a reliable, replicable way to predict observable traits like frailty, disease risk, and cellular failure?

That number is the business end of any claim about “reversing aging.” You can’t fix what you can’t measure. Epigenetic clocks provide exactly that—a quantifiable, molecular measurement that translates invisible biochemical changes into a single, testable prediction: your biological age.

¹His date of birth has never been verified by an official document; British authorities declined to certify a centenary record because the supporting paperwork did not exist. The story is true in spirit, asterisked in detail. It is, in this respect, an appropriate opener for an essay about measurement.

DNA Methylation and How Scientists Measure It

The story begins with a simple chemical mark on DNA called methylation. Remember, DNA is a long chain of molecules (each of which can only be either A, T, C, or G) holding hands. Sometimes some molecules get a methyl group (CH₃) attached. Specifically, methylation at CpG sites—locations where a cytosine nucleotide sits next to a guanine nucleotide in the DNA sequence.

Methylation doesn't change the DNA sequence itself. Instead, it controls which genes get expressed when. Scientists discovered that certain CpG sites change their methylation patterns predictably with age. Some sites gain methylation. Others lose it. These changes happen so consistently across people that they function as a molecular clock ticking away inside every cell.

Imagine you're baking cookies. Most are plain pecan flavor, but sometimes your little sister sneaks chocolate chips into the dough while you're not looking. In the raw dough, you can't tell which cookies have chocolate chips. They all look identical.

After baking, though, something revealing happens. The cookies with chocolate chips develop visible puddles of melted chocolate on their surface. Now you can easily tell which cookies have chocolate chips just by looking at the melted chocolate puddles.

In real molecular biology, this is achieved through bisulfite treatment. Scientists treat DNA with bisulfite, which chemically converts unmethylated cytosines into uracil (which then reads as thymine during sequencing). But methylated cytosines—like those chocolate chips—remain unchanged as cytosines.

The melted chocolate puddle is the key. After sequencing, unmethylated sites appear as thymine (the conversion product), while

methylated sites still appear as cytosine. The hidden chemical mark becomes encoded directly in the DNA sequence itself, as readable as any other genetic letter.




FIGURE 3. Methylation marks DNA. Enzymes add methyl groups (CH₃, red circles; chocolate chips) to cytosines at CpG sites.

Now imagine you don't just have a few dozen cookies, but millions—and you bake millions of batches. If you stack cookies from each batch together, and then align the different stacks, you can look at each position and calculate the ratio of chocolate chip to plain cookies.

This ratio is just a number. A number between 0 and 1. Biologists often call this number β (beta). A single epigenetic clock is just a recipe that turns millions of those beta values into one age estimate.

The Horvath Clock

In 2013, Steve Horvath at UCLA published a breakthrough. He analyzed methylation patterns from 8,000 samples spanning 51 different tissue types and identified 353 CpG sites that changed predictably with age.

According to Horvath, this clock “can be used to measure age acceleration in aging disorders” and “offers a universally applicable multi-tissue estimator of biological age.” His clock worked across almost every human tissue—blood, brain, liver, lung.

The accuracy was stunning.

The clock predicted chronological age within 3.6 years on average. For a biological measurement spanning decades, this precision was unprecedented. Nevertheless, Horvath’s clock essentially tells you only how your methylome compares to the population average for your age. It measures correlation with chronological age, but—as previously mentioned—this is not necessarily what matters most for health.

Different Clocks, Different Questions

The field quickly realized that different clocks could answer different questions. Gregory Hannum and colleagues at the University of California, San Diego developed a blood-specific clock in 2013.

Morgan Levine and Steve Horvath created PhenoAge in 2018, trained on health outcomes rather than just chronological age. According to Levine and Horvath, PhenoAge “captures differences in risk of all-cause mortality and a variety of age-related diseases better than chronological age.”

Ake Lu, Steve Horvath, and colleagues pushed further with GrimAge in 2019, combining methylation with plasma proteins. As Lu and Horvath reported, GrimAge “strongly predicts lifespan and healthspan.” Someone whose GrimAge exceeds their chronological age faces elevated mortality risk.

The most interesting recent entry is conceptually different. In 2022, Daniel Belsky and colleagues at Columbia and Duke published DunedinPACE, trained on the Dunedin cohort—a New Zealand birth-year group followed since 1972 with serial measurements of organ function. DunedinPACE does not estimate *how old you are*. It estimates *how fast you are aging*. The unit is years of biological aging per chronological year. A reading of 1.0 means you are aging at the expected rate; 1.2 means twenty percent faster; 0.85 means fifteen percent slower. The CALERIE caloric-restriction trial, the only randomized human trial of an anti-aging intervention with a clean endpoint, used DunedinPACE and detected a measurable slowdown. That a single longitudinal cohort produced the only clock currently sensitive to a real intervention is a useful reminder of how much the clocks depend on the data underneath them.

The Machine Learning Behind the Clocks

But what makes these clocks “tick”? Machine learning and statistics.

The current state-of-the-art clocks rely on a framework called elastic-net regression. Elastic-net regression solves a feature-selection problem through a clever mathematical trick. Remember high school algebra, when you had to find the line between two points? This is the same idea, but instead of a two-dimensional space, we have an n -dimensional space—where n might be 450,000 CpG sites. And we don’t want the line that exactly fits all the points (that would just memorize noise), but rather the line that best fits the underlying

pattern.

The algorithm minimizes prediction error while simultaneously penalizing model complexity. Mathematically, it optimizes:

$$\text{minimize: } \sum_i (\text{observed age}_i - \text{predicted age}_i)^2 + \lambda_1 \sum_j |\beta_j| + \lambda_2 \sum_j \beta_j^2.$$

The first term is just prediction error—how far off are our estimates? The second and third terms are regularizers: they restrict the kind of model we can pick, either by forcing most coefficients to exactly zero—thereby eliminating irrelevant CpG sites entirely—or by forcing the few nonzero coefficients to be still very small. Think of it as saying “unless you really help predict age, you’re out; otherwise you will only add noise to the model.”

According to Horvath, “The appeal of the elastic net is that it automatically performs variable selection and shrinkage.” Out of 450,000 measured CpG sites, the algorithm identifies the 50 to 700 sites that genuinely matter, and ignores everything else.

Some researchers have recently turned to less traditional approaches—deep neural networks and other methods commonly referred to as deep learning. These models can theoretically capture nonlinear interactions between methylation sites that linear models miss. But as Lara Srour and colleagues noted in their 2025 review, “limited sample sizes, confounding factors, and overfitting risk mean that elastic-net pipelines remain the pragmatic baseline in most applications.” Translation: the fancy stuff requires data that aging research simply doesn’t have yet. The mathematically elegant solution from 2013 still works best.

The Commercial Clock

By the early 2020s, the methodology had escaped the laboratory. Elysium Health licensed Horvath's group's technology and launched Index, a consumer methylation test marketed in glossy boxes. Tally Health, co-founded by David Sinclair, sells TallyAge based on its own proprietary clock. MyDNAge, Zymo's earlier consumer entry, advertises a Horvath-based result. Each test costs between \$200 and \$500. Each returns a single number, often accompanied by a curated set of lifestyle recommendations.

The clocks themselves are real science. The retail packaging is something else. The published clocks have prediction intervals on the order of 3–5 years for an individual. The companies selling them generally do not advertise this. The clocks were validated against chronological age in trained populations; their behavior on out-of-distribution individuals is less well-characterized than the glossy boxes suggest. And no clock—not Horvath's, not GrimAge, not DunedinPACE—has been validated as a clinical decision tool. They are, at best, research instruments dressed for retail.

It would be a mistake, however, to read this as a story about decay or fraud. It is the same story as the rest of the collection: scientists publish, companies wrap, journalists amplify, customers buy, the cycle repeats with a sharper instrument. Each iteration chips away a little more of the mountain. The methylome is better understood now than it was in 2013, the clocks are better than they were, and a future generation of clocks will be better still. The direction is forward even when the storefront is somewhat sappy, silly and unrigorous. Keep the storefront in mind for the next part of the collection. The clocks are an early instance of a pattern; they will not be the last.

The Conquest Begins

On the night of November 10, 1619, René Descartes was having fever dreams. In one of them, an angel told him:

The conquest of nature is to be achieved through measure and number.

As we now know, Descartes got to work right away, and many generations of scientists followed. Yet, the invisible machinery of aging—the molecular choreography, the chemical marks accumulating in every cell—remained beyond his reach.

Epigenetic clocks transform biological time into a number. They turn the conquest of nature inward, toward the conquest of our own deterioration. You can't reverse what you can't measure. Now we can measure. The fixing comes next.

Sources and Further Reading

- Belsky, D. W., Caspi, A., Houts, R., Cohen, H. J., Corcoran, D. L., Sugden, K., Baccarelli, A. A., Poulton, R., & Moffitt, T. E. (2022). DunedinPACE, a DNA methylation biomarker of the pace of aging. *eLife*, *11*, e73420. <https://doi.org/10.7554/eLife.73420>
- Hannum, G., Guinney, J., Zhao, L., Zhang, L., Hughes, G., Sada, S. V., Klotzle, B., Bibikova, M., Fan, J.-B., Gao, Y., Deconde, R., Chen, M., Rajapakse, I., Friend, S., Ideker, T., & Zhang, K. (2013). Genome-wide methylation profiles reveal quantitative views of human aging rates. *Molecular Cell*, *49*(2), 359–367. <https://doi.org/10.1016/j.molcel.2012.10.016>
- Horvath, S. (2013). DNA methylation age of human tissues and cell types. *Genome Biology*, *14*(10), R115. <https://doi.org/10.1186/gb-2013-14-10-r115>
- Levine, M. E., Lu, A. T., Quach, A., Chen, B. H., Assimes, T. L., Bandinelli, S., Hou, L., Baccarelli, A. A., Stewart, J. D., Li, Y., Whitsel, E. A., Wilson, J. G., Reiner, A. P., Aviv, A., Lohman, K., Liu, Y., Ferrucci, L., & Horvath, S. (2018). An epigenetic biomarker of aging for lifespan and healthspan.

- Aging (Albany NY)*, 10(4), 573–591. <https://doi.org/10.18632/aging.101414>
- Lu, A. T., Quach, A., Wilson, J. G., Reiner, A. P., Aviv, A., Raj, K., Hou, L., Baccarelli, A. A., Li, Y., Stewart, J. D., Whitsel, E. A., Assimes, T. L., Ferrucci, L., & Horvath, S. (2019). DNA methylation Grim-Age strongly predicts lifespan and healthspan. *Aging*, 11(2), 303–327. <https://doi.org/10.18632/aging.101684>
- Prosz, A., Zrimec, J., Yordanov, P., Börlin, C. S., Neumann-Staubitz, P., Alili, A., Kretschmer, F., von Hagen, J., Jouhten, P., & Neumann-Schaal, M. (2024). Biologically informed deep learning for explainable epigenetic clocks. *Scientific Reports*, 14, 287. <https://doi.org/10.1038/s41598-023-50495-5>
- Srouf, L., Heydari Seradj, S. M., & Leung, A. K. L. (2025). Deep aging clocks: AI-powered strategies for biological age estimation. *Trends in Molecular Medicine*, 31(1), 54–68. <https://doi.org/10.1016/j.molmed.2024.09.007>
- Waziry, R., Ryan, C. P., Corcoran, D. L., Huffman, K. M., Kobor, M. S., Kothari, M., Graf, G. H., Kraus, V. B., Kraus, W. E., Lin, D. T. S., Pieper, C. F., Ramaker, M. E., Bhapkar, M., Das, S. K., Ferrucci, L., Hastings, W. J., Kebbe, M., Parker, D. C., Racette, S. B., ... Belsky, D. W. (2023). Effect of long-term caloric restriction on DNA methylation measures of biological aging in healthy adults from the CALERIE trial. *Nature Aging*, 3(3), 248–257. <https://doi.org/10.1038/s43587-022-00357-y>

PART II

Of Bird Beaks and Bargains

Current Research and Venture-Capital Black Holes

The Fluor de Lys Affair

A Five-Act Play About Resveratrol, Sirtris, and a \$720 Million Mistake

Act I: The Scene of the Crime

Harvard Medical School, 2003. In a laboratory tucked into the warren of research spaces on Longwood Avenue, David Sinclair, an assistant professor of pathology who had recently escaped from Leonard Guarente's shadow at MIT, was examining the results of his experiment with a huge smile on his face; every single experiment, every technical replicate, every independent test showed the same stunning result: resveratrol—a simple polyphenol found in red wine, the kind of molecule that wouldn't look twice at you if you passed it on the street—was activating SIRT1 with extraordinary potency.

The numbers were beautiful. Clean. Reproducible.

According to Sinclair and his collaborators, resveratrol was “as close to a miraculous molecule as you can find.” The compound extended lifespan in yeast. The mechanism seemed clear: activation of Sir2, the founding member of the sirtuin family—enzymes that, when boosted, appeared to slow the molecular clock of aging itself.

The results were published in *Nature* in September 2003. This was not some minor observation tucked into a specialty journal, but a front-page discovery that promised to crack the code of human longevity. *Science* ran an admiring perspective piece by Sinclair himself a few weeks later, on a related theme. Media coverage exploded; red wine sales spiked; and by 2004, Sinclair had co-founded Sirtris Pharmaceuticals with a dream team of entrepreneurs and scientists.

INTERLUDE: WHAT ARE SIRTUINS, AND WHY SHOULD ANYONE CARE?

By 2008, GlaxoSmithKline (GSK) acquired Sirtris for \$720 million.

The story should have ended there: brilliant young scientist makes breakthrough, builds company, exits rich, continues making breakthroughs. A perfect arc.

But science often fails to be a Hollywood screenplay, and something was wrong.

Interlude: What Are Sirtuins, and Why Should Anyone Care?

Before proceeding, we need to understand what was supposedly being activated.

Sirtuins are enzymes. There are seven flavours of them in humans, designated SIRT1 through SIRT7, and they remove acetyl groups from proteins (see Figure 4). Think of proteins as machines with switches; adding an acetyl group flips the switch one way, removing it flips it the other way. Sirtuins are the switch-flippers.

The sirtuin story begins with yeast in the 1990s. Leonard Guarente's laboratory at MIT discovered that a gene called *Sir2* extended lifespan when overexpressed. Amplify *Sir2*, and yeast cells lived longer. Delete it, and they died younger.

Then came the connection to caloric restriction. Guarente's group showed that *Sir2* activity increases during caloric restriction, and that deleting *Sir2* blocks the lifespan extension entirely. Now, caloric restriction is *the* most reliable way to extend lifespan across essentially every organism tested. Starve (a bit, not to death) yeast, worms, flies, or mice, and they live longer.

The clear implication was that sirtuins might be the molecular




FIGURE 4. The seven human sirtuins (SIRT1–7) occupy distinct cellular locations. SIRT1, SIRT6, and SIRT7 reside in the nucleus; SIRT2 in the cytoplasm; SIRT3, SIRT4, and SIRT5 in the mitochondria. Resveratrol’s proposed target, SIRT1, operates in the nucleus where it regulates gene expression.

mediators of why eating less makes organisms live longer. Activate SIRT1, the thinking went, and you could mimic caloric restriction pharmacologically. All the benefits of eating less without actually starving.

In mammals, SIRT1 became the star. It regulates p53 (the tumor suppressor), FOXO proteins (which control stress resistance), and PGC-1 α (which builds new mitochondria—the cell’s power plants). Activate SIRT1, the thinking went, and you could mimic caloric restriction pharmacologically. All the benefits of eating less without actually starving.

And resveratrol appeared to be the key that unlocked it.

Act II: The First Cracks in the Mirror

Matt Borra, working at the University of Wisconsin–Madison, noticed something peculiar in 2005 while trying to replicate the resveratrol activation. He was using different versions of the substrate—the molecular target that SIRT1 acts upon; the thing it’s supposed to deacetylate—and the activation kept changing: sometimes resveratrol worked brilliantly, sometimes it barely worked at all, and the only variable that seemed to matter was which fluorescent tag he had attached to the peptide.

This was strange. This was *very* strange.

The fluorescent tag wasn’t supposed to matter; it was just there to help measure the reaction, like a tracer dye in a plumbing system that lets you see where the water goes but doesn’t change the flow itself. Yet, Borra’s data suggested otherwise. In fact, the results showed that SIRT1 activation only took place when the fluorescent tag was present. Remove the tracer dye, and suddenly the water stopped flowing.

Other laboratories began noticing the same pattern. Matt Kaeberlein at the University of Washington couldn’t replicate the lifespan extension in yeast unless the fluorescent assay was used; with native substrates, resveratrol had no effect. Whispers started circulating at conferences. Emails were exchanged. Eyebrows were raised.

Something was wrong with the assay.

Act III: The Pfizer Papers

In 2009, Dave Beher and colleagues published a systematic take-down: resveratrol did not activate SIRT1 when tested against p53-derived peptides lacking fluorophores; it did not activate SIRT1

when tested against PGC-1 α isolated from cells. It did not, in short, activate the enzyme under conditions that remotely resembled what the actual SIRT1 protein encounters inside living cells—as opposed to artificial lab-created substrates with fluorescent tags attached.

According to Beher's team, "the Fluor de Lys-SIRT1 peptide is an artificial SIRT1 substrate because in the absence of the covalently linked fluorophore the peptide itself is not a substrate of the enzyme." Covalently linked means chemically bonded—the fluorophore wasn't just sitting nearby, it was molecularly glued to the peptide. The assay substrate was a lie.

Then came the kill shot. In 2010, Michael Pacholec and a team of scientists at Pfizer—yes, Pfizer, the pharmaceutical giant, which had every reason to want the sirtuin story to be true—published the most comprehensive demolition yet. They tested not just resveratrol but the entire suite of compounds that Sirtris had developed: SRT1720, SRT2183, SRT1460. They used multiple independent methods that approach the problem from completely different angles: high-performance liquid chromatography, mass spectrometry, nuclear magnetic resonance spectroscopy, surface plasmon resonance, isothermal titration calorimetry.

Every method told the same story. None of the compounds activated SIRT1 when tested against native substrates lacking fluorophores; worse, the compounds actually bound directly to the fluorophore-tagged peptides even in the complete absence of SIRT1. The drugs weren't interacting with the enzyme at all. They were interacting with the measuring tool.

The activation was an artifact—a false signal created by the measurement process itself. Smoke and mirrors.

According to Pacholec's team, "SIRT1 activation by resveratrol

was shown to be completely dependent on the presence of a covalently attached fluorophore in the fluorescent peptide substrate.” Later structural biology—using techniques that let scientists see the three-dimensional shape of molecules—revealed the mechanism: resveratrol binds in a pocket formed partly by the fluorophore itself. Remove the bulky fluorescent tag and that binding pocket disappears, collapsing like a house of cards.

The killer had been hiding in plain sight the entire time.

Act IV: Whodunit?

Was this fraud? Deliberate deception? Scientific misconduct?

The evidence says no. But it’s complicated.

Here’s what actually happened: in 2003, scientists at Biomol Research Laboratories developed a biochemical assay for measuring sirtuin activity. An assay is just a test—a standardized procedure that lets you measure whether something is happening. In this case: is SIRT1 being activated?

The assay used fluorescently labeled peptides. A peptide is a small chain of amino acids—think of it as a tiny piece of protein. The fluorescent label is like a glow stick attached to the peptide; it lets you see what’s happening because fluorescence is easy to measure at scale using standard plate readers, which makes high-throughput screening commercially viable. You can run thousands of tests quickly.

Here’s the crucial bit: the fluorescent tag is supposed to be an inert observer. It’s like a camera filming an event; it should record what’s happening without changing the event itself.

Imagine you ask your sister whether there are cookies left in the cookie jar. She runs to the kitchen, opens the jar, sees there are indeed cookies left—and then eats them all. When she comes back, she reports: “No cookies left.” Technically true. But misleading. The cookies existed until she interfered.

That’s essentially what the fluorophore was doing. It wasn’t just observing the reaction between resveratrol and SIRT1; it was participating in it, creating a binding interaction that wouldn’t exist without its presence.

When Sinclair’s laboratory screened small molecules using this assay, resveratrol showed robust activation. The result was reproducible. The biochemistry was real.

Nobody initially questioned whether the fluorophore might be interfering. Why would they? After all, fluorescent tags are ubiquitous in drug discovery, and they behave as inert observers.

Later crystallography—essentially taking molecular-scale photographs of the proteins—revealed what was happening. The AMC fluorophore wasn’t just a passive observer; it was creating a binding pocket that doesn’t exist with natural substrates. Picture it like this: resveratrol acts as molecular glue, but it only sticks when the bulky fluorescent tag is there to create the right surface. Remove the tag, and suddenly there’s nothing for the glue to grip onto.

The activation was real. But it was also artifactual. The measurement system was creating the phenomenon it claimed to observe.

This explains why the original results seemed so robust and reproducible; in the fluorescent assay, the activation genuinely occurred—it wasn’t fabricated, it wasn’t cherry-picked, it wasn’t statistically

manipulated. The chemistry was real. But it also wasn't... real real. It was a reaction between a drug, an enzyme, and a synthetic fluorescent tag that doesn't exist anywhere in human cells.

Act V: The Accomplices

What amplified a laboratory artifact into a \$720 million acquisition wasn't malice; it was the perfect storm of legitimate scientific excitement colliding with venture-capital incentives and media amplification.

Sinclair's early work on sirtuins and aging in yeast was genuinely groundbreaking; he had discovered that Sir2 extends lifespan by reducing accumulation of extrachromosomal DNA circles, a finding published in *Cell* when he was still a postdoc, and he wrote a thoughtful *Science* perspective on the resulting picture of yeast genome instability the same year resveratrol made him famous. The idea that small molecules could pharmacologically activate longevity pathways wasn't absurd at all.

However, nuance doesn't play well in boardrooms; uncertainty doesn't close investment rounds. "Resveratrol might activate SIRT1 in specific contexts depending on substrate identity" is accurate but unmarketable. "Red wine compound could extend human lifespan" sells.

The result was a feedback loop. Media hype generated investor interest; investor interest generated pressure to move compounds into clinical trials quickly; speed meant less time for the careful biochemical validation that might have caught the assay problem earlier. By the time independent laboratories started publishing contradictory results, Sirtris was already public and GSK had already written the check.

Nobody deliberately deceived anyone. But the system—the publish-or-perish academic incentive structure combined with the move-fast-and-break-things ethos of biotech—created conditions where preliminary findings could be transformed into commercial certainty before the science was ready.

Coda: The \$720 Million Footnote

In March 2013, GSK quietly closed Sirtris's Cambridge, Massachusetts office and absorbed the program into its main research operations in the United Kingdom. The clinical compounds—SRT1720, SRT2183, SRT1460, and the resveratrol formulation that had launched the whole enterprise—were either shelved or transferred to Sinclair's academic laboratory at Harvard. By 2014, the Sirtris brand no longer existed inside GSK. The \$720 million had been written down in the same quiet, cumulative way that pharmaceutical write-downs always happen: not as an announcement but as a series of footnotes in quarterly reports that nobody reads cover to cover.

That is one way to read the story: as a cautionary tale, a warning sign, a closed file. It is not the only way to read it.

Look at what the \$720 million also bought, for the field, while it was going down the drain. The Pfizer team's 2010 demolition was a master class in modern biochemistry: HPLC, mass spectrometry, NMR, surface plasmon resonance, isothermal titration calorimetry, all trained on a single artifact, all converging on the same answer. Borra's 2005 paper, Kaeberlein's 2005 yeast work, Beher's 2009 peptide series, the structural biology that finally caught the fluorophore in the act—these papers raised the methodological floor of an entire discipline. Sinclair's own laboratory, after the write-down, went on to publish the 2020 *Nature* paper on optic-nerve regeneration via

partial reprogramming that we will meet in the next essay. Crystallographers a decade later identified a narrower window in which sirtuin activators may, in fact, work on specific substrates. The supplement market kept selling resveratrol the whole time - take that as you will.

Each of these is a beak-stroke against the mountain. The hundred years pass. The bird returns. The shepherd boy's answer is not that the bird is futile—it is that nobody, individually, gets to see the second tick over. Scientists publish and move on. Investors unwind their positions. Journalists file the next story. Customers buy the next bottle. And the methylome is better understood now than it was, the sirtuin biology is better understood than it was, the assays are cleaner than they were. The Zeitgeist has a clear direction.

Capital allocation is the field's honest report card, but it is not the field's honest job. The peer-reviewed literature is a slow, contested, ambiguous record. The press is a fast, credulous, breathless record. The acquisitions are neither: they are bets placed by people who have read the literature and the press and made a decision with their employers' money. When those bets unwind, they unwind quietly. They settle into footnotes. The science the bets paid for, however, persists.

Keep this in mind for the rest of Part II. The numbers will recur: billions raised by Calico, billions by Altos, the ASPIRE failure at UNITY, the BioAge wind-down. Each one is a footnote in a quarterly report. Each one is the field telling you, in the only language the market understands, that the science was not ready when the money arrived. Each one is also another beak-stroke. The mountain is large.

Sources and Further Reading

Beher, D., Wu, J., Cumine, S., Kim, K. W., Lu, S.-C., Atangan, L., & Wang,

- M. (2009). Resveratrol is not a direct activator of SIRT1 enzyme activity. *Chemical Biology & Drug Design*, 74(6), 619–624. <https://doi.org/10.1111/j.1747-0285.2009.00901.x>
- Borra, M. T., Smith, B. C., & Denu, J. M. (2005). Mechanism of human SIRT1 activation by resveratrol. *Journal of Biological Chemistry*, 280(17), 17187–17195. <https://doi.org/10.1074/jbc.M501250200>
- Cao, D., Wang, M., Qiu, X., Liu, D., Jiang, H., Yang, N., & Xu, R. M. (2015). Structural basis for allosteric, substrate-dependent stimulation of SIRT1 activity by resveratrol. *Genes & Development*, 29(12), 1316–1325. <https://doi.org/10.1101/gad.265462.115>
- Dai, H., Sinclair, D. A., Ellis, J. L., & Steegborn, C. (2018). Sirtuin activators and inhibitors: Promises, achievements, and challenges. *Pharmacology & Therapeutics*, 188, 140–154. <https://doi.org/10.1016/j.pharmthera.2018.03.004>
- Howitz, K. T., Bitterman, K. J., Cohen, H. Y., Lamming, D. W., Lavu, S., Wood, J. G., Zipkin, R. E., Chung, P., Kisielewski, A., Zhang, L. L., Scherer, B., & Sinclair, D. A. (2003). Small molecule activators of sirtuins extend *Saccharomyces cerevisiae* lifespan. *Nature*, 425(6954), 191–196. <https://doi.org/10.1038/nature01960>
- Hubbard, B. P., Gomes, A. P., Dai, H., Li, J., Case, A. W., Considine, T., Riera, T. V., Lee, J. E., e, S. Y., Lamming, D. W., Pentelute, B. L., Schuman, E. R., Stevens, L. A., Ling, A. J., Armour, S. M., Michan, S., Zhao, H., Jiang, Y., Sweitzer, S. M., ... Sinclair, D. A. (2013). Evidence for a common mechanism of SIRT1 regulation by allosteric activators. *Science*, 339(6124), 1216–1219. <https://doi.org/10.1126/science.1231097>
- Kaeberlein, M., McDonagh, T., Heltweg, B., Hixon, J., Westman, E. A., Caldwell, S. D., Napper, A., Curtis, R., DiStefano, P. S., Fields, S., Bedalov, A., & Kennedy, B. K. (2005). Substrate-specific activation of sirtuins by resveratrol. *Journal of Biological Chemistry*, 280(17), 17038–17045. <https://doi.org/10.1074/jbc.M500655200>
- Pacholec, M., Bleasdale, J. E., Chrunnyk, B., Cunningham, D., Flynn, D., Garofalo, R. S., Griffith, D., Griffor, M., Loulakis, P., Pabst, B., Qiu, X., Stockman, B., Thanabal, V., Varghese, A., Ward, J., Withka, J., & Ahn, K. (2010). SRT1720, SRT2183, SRT1460, and resveratrol are not direct activators of SIRT1. *Journal of Biological Chemistry*, 285(11), 8340–8351. <https://doi.org/10.1074/jbc.M109.088682>

- Schmidt, C. (2010). GSK/Sirtris compounds dogged by assay artifacts. *Nature Biotechnology*, 28(3), 185–186. <https://doi.org/10.1038/nbt0310-185>
- Sinclair, D. A. (2003). An age of instability. *Science*, 301(5641), 1859–1860. <https://doi.org/10.1126/science.1090465>

The Second Foundation

Sinclair's Shot at Redemption: Life Biosciences's FDA Investigative New Drug (IND) approval

The Sirtris arc—the whole eight-year drama of resveratrol and the Fluor de Lys assay and GSK's \$720 million—was about exactly one of the twelve hallmarks of aging. Sirtuins are deacetylases; they clip acetyl groups off histones and other proteins; histone acetylation is one of the chemical languages the cell uses to manage which genes are open and which are closed. Sirtuins, in other words, were a play to fix *epigenetic alterations*, hallmark number four on the list. One enzyme family, one type of chemical mark, one slice of one hallmark.

The Yamanaka factors, the main subject - other than David Sinclair himself - of this article, are a different kind of object. It is harder to categorize their effect under a single hallmark because, in some sense, it is a tool that operates upstream of most of them. It is more powerful, in a way; more general.

That same reach is what makes it so hard to deploy responsibly, and it is the technical heart of the next chapter of the Sinclair arc. The science is real. The scandals are real. They are running on the same researcher's CV simultaneously. This essay walks through both: the Phase 1 trial that began in 2026, and the journal paper and dog chew that should not have happened in the months on either side of it. Whether the science survives the scandals is a question the field is still in the middle of answering.

What Cellular Reprogramming Actually Is: the Yamanaka Factors

There are two types of cells in eukaryotic organisms: somatic cells and germline cells. Somatic cells, named after the Greek (*sôma*), meaning body, are every cell in the organism that carries out a biological function other than reproduction — neurons firing in your brain, cardiomyocytes contracting in your heart, fibroblasts knitting together your connective tissue. Germline cells are the cells that pass genetic information to the next generation, giving rise to gametes — sperm and eggs.

Both somatic and germline cells arise from the differentiation of Embryonic Stem Cells (ESCs). Imagine a batch of first-year university students; they could study mathematics, engineering, chemistry, history, or some other specialized field (somatic cells), but they could also study education and become high school teachers (germline cells) who are then able to eventually help produce new first-year students (ESCs).

Most students don't become teachers. Most of our cells are somatic cells. However, what if we could turn a fourth-year student — filled with specialized knowledge, but also biases and trauma and with less neuroplasticity — back into a first-year student, a much cleaner slate? While that is impossible for us students, fortunately the situation is quite different for somatic cells.

M.D., Ph.D., Professor at both Kyoto University and the University of California, San Francisco (UCSF), and rock star of stem cell research, Shinya Yamanaka found a way to turn somatic cells back into pluripotent stem cells — ESCs being one specific variety. He received the 2012 Nobel Prize in Physiology or Medicine "for the discovery that mature cells can be reprogrammed to become pluripotent." In a landmark 2006 paper, Yamanaka and his colleague Kazutoshi Takahashi wrote: "By combining four selected factors,

we were able to generate pluripotent cells, which we call induced pluripotent stem (iPS) cells, directly from mouse embryonic or adult fibroblast cultures."

By "factors" he means transcription factors. Transcription factors are proteins that bind directly to DNA and regulate the process of gene transcription, which eventually impacts the concentration in a cell of the protein produced by that gene. Think of DNA as flat-packed IKEA furniture: the instructions and parts are all there, but nothing gets built without someone actually reading the manual and picking up the tools. Transcription factors are those tools. A cell flooded with the right transcription factors builds proteins faster and more precisely than one with few — just as an IKEA assembly crew armed with powered screwdrivers and torque wrenches will outbuild one working with a single bent Allen key (which in practice is usually what IKEA includes).

Yamanaka showed that introducing extra copies of just four genes — *OCT4*, *SOX2*, *KLF4*, and *c-MYC* — delivered into somatic cells via retroviral vectors, was sufficient to flood the cell with the corresponding transcription factors Oct4, Sox2, Klf4, and c-Myc (yes, the genes and the proteins they code for have the same name, but often that name is written in *italics* when it refers to the gene). Those four proteins then systematically dismantled the cell's specialized identity and rebooted it toward pluripotency. Four keys, and an entirely different lock opened. *OCT4*, *SOX2*, *KLF4*, and *c-MYC* came to be known as the *Yamanaka* factors, and abbreviated as *OSKM* in the literature

Even though Yamanaka's work was centered on stem cells, it has had a profound impact on longevity research. Epigenetic age — the accumulated chemical damage to the genome's control system, the scrambled methylation patterns that the Horvath clock measures — is not permanent. It can, at least in principle, be erased and reset.

The immediate problem, however, is that full reprogramming to iPSCs also erases cellular identity. A neuron reset all the way to pluripotency is no longer a neuron; it has forgotten everything it knew. For longevity purposes, that is trading one catastrophe for another.

The field's solution came in 2016, when Alejandro Ocampo and colleagues published a landmark paper in *Cell* demonstrating partial reprogramming in living mice. Rather than running the Yamanaka factors long enough to fully erase cellular identity, they applied them transiently—briefly, in pulses—using a clever delivery system, which we will dissect in detail in the next section. The result was a cell with its epigenetic clock partially reset, shedding markers of aging without losing its identity in the process. The fourth-year student forgot some of their accumulated stress and rigidity, but remained, crucially, a fourth-year student.

“Collectively,” Ocampo wrote, “these results demonstrate that short-term *in vitro* induction of OSKM in cells derived from a premature aging mouse model ameliorates multiple age-associated hallmarks observed during physiological aging.” It was the first demonstration that the rejuvenating power of the Yamanaka factors could be decoupled from the chaos of full reprogramming, and it lit a fuse in longevity research that is still burning.

The 2020 Paper That Made the Cover

A small problem, before the science: how do you get four genes *into* a living cell?

Yamanaka used **retroviruses**. The name is doing real work, so it is worth slowing down. A normal virus, like the rhinoviruses or coronaviruses you and I have spent the past few years dodging,

carries its genetic material as RNA, dumps it into the cytoplasm of your cell, hijacks your ribosomes to make more virus, and leaves. The cell's nucleus—the vault where your DNA lives—never gets touched. (No, COVID did not alter your genome. SARS-CoV-2 is a positive-sense single-stranded RNA virus that replicates in the cytoplasm; it lacks the molecular machinery to enter your nucleus, let alone integrate into your chromosomes. Sleep well.)

Retroviruses are different. The “retro” refers to the fact that they go *backwards* relative to the central dogma of biology—DNA makes RNA makes protein. Retroviruses carry their genome as RNA, but they also bring along an enzyme called **reverse transcriptase**, which copies that RNA *into* DNA after entry into the cell. The newly synthesized viral DNA then gets shuttled into the nucleus and welded into the host's chromosomes by another viral enzyme called integrase. The cell, from that moment on, contains the virus's genes as part of its own genome. Forever. HIV is the most famous retrovirus, for exactly this reason.

Retroviral integration is fantastic for making iPSCs in a Petri dish, where the integration happens once and you keep the resulting cells indefinitely. It is a terrible idea for a therapy, where you would be permanently rewriting the DNA of specific tissues in a living person, with all four Yamanaka factors running in those tissues forever. *c-MYC* in particular is a well-known oncogene; leaving it constitutively turned on inside an organism is how you make tumors, not how you make young eyes.

By 2020, the field had switched to two better ideas. First, drop *c-MYC*. The remaining three factors—*OCT4*, *SOX2*, *KLF4*, abbreviated *OSK*—turn out to be sufficient for partial reprogramming, with a much friendlier safety profile. Second, deliver them with **adeno-associated viruses (AAVs)** instead of retroviruses. AAVs do not integrate the way retroviruses do; they sit alongside the genome

rather than rewriting it.² Better still, you can put the OSK genes under the control of a chemical switch—a so-called **Tet-On promoter**, which only turns the genes on when the patient is given a tetracycline-family antibiotic such as doxycycline. Stop the doxycycline, stop the reprogramming. The Yamanaka factors become a faucet you can close.




FIGURE 5. Schematic of the Tet-On and Tet-Off dual AAV vectors used by the Sinclair lab to deliver OSK with chemical control. In the Tet-On configuration (top), doxycycline (DOX) binds the transactivator rtTA, which then docks at the TRE promoter and switches the OSK gene on. Tet-Off works in mirror image: doxycycline turns the system *off*. From Lu et al., *Nature* (2020).

This is the architecture Yuancheng Lu, David Sinclair, and colleagues used in their landmark December 2020 *Nature* paper. They packed OSK into AAVs under a Tet-On switch, injected them into the eyes of mice, fed the mice doxycycline, and watched what happened.

What happened was astonishing. The mice in question had glaucoma, or had been subjected to optic-nerve crush injuries, or were

²This is a slight simplification. AAVs *primarily* exist as episomes—loops of DNA floating in the nucleus separately from the chromosomes—but a small fraction can integrate at low frequency. “Sit alongside” is the right first-order picture.

simply old. The retinal ganglion cells—the neurons that send vision from the retina to the brain—were dying or dead. Three weeks of doxycycline-controlled OSK expression, and those neurons grew back. The mice could see again. The methylation patterns on the DNA of the recovered cells looked like the methylation patterns of younger mice. The epigenetic clock, in those specific cells, had been wound back.




FIGURE 6. Confocal images of optic-nerve sections from 12-month-old mice five weeks after optic-nerve crush. Top: control mice expressing only a fluorescent reporter (d2EGFP) show essentially no axonal regeneration past the crush site (asterisks). Bottom: mice receiving the OSK gene therapy show extensive axon regrowth far beyond the injury. From Lu et al., *Nature* (2020).

The mechanism turned out to be elegant. OSK does not erase the methylation marks on DNA directly; it activates a family of enzymes called TETs (ten-eleven translocation methylcytosine dioxygenases), which then go on to remove the age-associated methyl groups themselves. Yamanaka factors as foremen rather than foot soldiers, calling in the cleaning crew rather than mopping the floor themselves. The paper made the cover of *Nature*. It earned the benefit of the doubt that Sinclair would, over the next few years, spend.

ER-100: The First Human Trial of Partial Reprogramming

It took five years, several rounds of nonhuman primate work, and a lot of regulatory paperwork, but on January 28, 2026, Life Biosciences—the company Sinclair co-founded around exactly this technology—announced that the FDA had cleared its Investigational New Drug application for **ER-100**, a partial epigenetic reprogramming therapy delivered as an intravitreal injection. The Phase 1 first-in-human trial (NCT07290244) is now recruiting patients with non-arteritic anterior ischemic optic neuropathy (NAION) and open-angle glaucoma at sites in the United States. Approximately a dozen patients in the first cohort. Primary endpoints: safety, tolerability, immune response. Secondary endpoints: visual function. According to Life Bio, this is the first epigenetic reprogramming therapy of any kind to be cleared for human trials.

It is worth being precise about what that means and what it doesn't. ER-100 is the 2020 *Nature* paper, basically: AAV2 carrying the three OSK genes, doxycycline-controlled, injected directly into the eye. The first quarter of 2026 will see human retinal ganglion cells exposed to Yamanaka factors for the first time, in patients with diseases that destroy those cells faster than the body can replace them. NAION, sometimes called “stroke of the eye,” has no approved treatment. Glaucoma has treatments that lower eye pressure but do nothing to repair already-damaged ganglion cells. The trial exists, in a real medical gap; this is not aging research dressed up as a disease indication, but disease research that happens to use a tool from aging research.

There remains, however, a non-trivial worry. Pushing somatic cells toward pluripotency is the reverse of the safety direction we usually want from a drug, because pluripotent cells unconstrained by their normal biology grow into **teratomas**: benign tumors that, in their efforts to differentiate into multiple tissues at once, sometimes

produce hair, bone, teeth, and the occasional eye in places hair, bone, teeth, and eyes ought not to be. If teratomas did not exist, people would still have nightmares about teratomas. Partial reprogramming is designed—in principle—to stop short of pluripotency, and the doxycycline switch is designed to allow you to halt the process if something looks wrong, but designed and demonstrated are not the same word. The Phase 1 trial is, in part, where we find out.

Paul Knoepfler, the UC Davis stem-cell biologist who has covered Sinclair’s work skeptically for years, summarized the moment with careful enthusiasm: this is a real first-in-human trial of a real cellular reprogramming approach, and that matters; it is also not a pill you will be able to take, not aging reversal, not even necessarily an effective glaucoma treatment until the data come in. A dozen patients is a dozen patients. The biology is plausible. The trial is small. The hype, if it comes, will be larger than either.

Interlude: Five Days in July

* * *

If this was a theater play this is the part where the theatre lights would dim and the orchestra would re-enter with a darker chord. The science we just described is the upstage spotlight. We now turn to the downstage shadow—to the same researcher’s other ledger, opened in the same period, audited by the same field, and producing a very different kind of paper trail. Roll the curtain back.

On June 30, 2023, Sinclair’s lab submitted a paper to the journal *Aging*. Five days later, on July 4—American Independence Day, a federal holiday—the paper was accepted. Twelve days after that it was online. The paper, “Chemically induced reprogramming to reverse cellular aging,” described six small-molecule cocktails that, in cultured human fibroblasts, appeared to reproduce some of the

rejuvenation signatures of OSK reprogramming without the genes, the AAVs, or the doxycycline. A pill, the press release implied, might do the work that ER-100 does with surgery.

The first thing to notice about this paper is the timeline. Most peer review at reputable journals takes weeks to months. Five days is fast even for a paper waved through by a friendly editor. The second thing to notice is that David Sinclair is co-Editor-in-Chief of *Aging*.

Charles Brenner, Ph.D., Alfred E. Mann Family Foundation Chair in Diabetes and Cancer Metabolism at the City of Hope Beckman Research Institute, posted on X within hours of publication (Figure 7, top). Brenner went on to point out, in follow-up posts and interviews, that three of the compounds in the cocktails—CHIR99021, tranylcypromine, and valproic acid—are “generally not safe alone or in a combination,” and that the paper performed no animal work and no single-cell sequencing of cellular identity. The latter being the basic safety check anyone working on reprogramming has been doing since Ocampo 2016.

Matt Kaeberlein, Ph.D., professor (now emeritus) at the University of Washington and CEO of Optispan, was more economical (Figure 7, bottom). Speaking to the *Daily Mail*, Kaeberlein added that the authors “overstate the importance of their results by using phrases like ‘rejuvenate’ and ‘reverse biological aging’ but provide no actual evidence to support these claims.”

Paul Knoepfler, Ph.D., professor of cell biology at the UC Davis School of Medicine and proprietor of *The Niche*, the most widely read independent stem-cell blog, devoted a long post titled “Review of new David Sinclair paper, supplements & anti-aging glitz” to a section-by-section dismantling. Knoepfler counted the word “rejuvenation” thirty-five times in the paper, asked pointedly whether

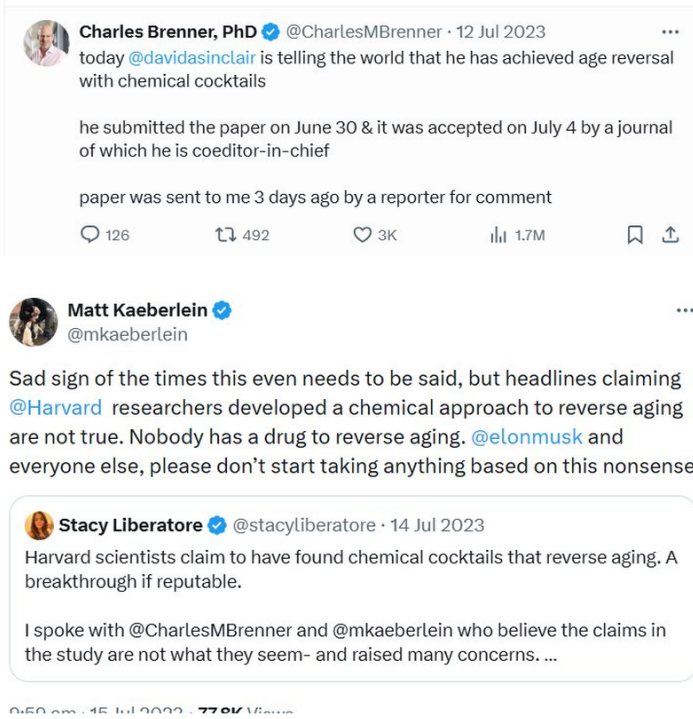


FIGURE 7. The receipts. Top: Charles Brenner, July 12, 2023. Bottom: Matt Kaeberlein, July 15, 2023. The paper had been online for one week.

the term had any scientifically rigorous definition, and noted the absence of any discussion of teratoma risk—a basic concern for any work that pushes cells toward pluripotency, since pluripotent cells unconstrained by their normal biology form teratomas, which are the kind of tumor we already met a section ago. Knoepfler included an image of a teratoma generated in his own laboratory and observed, with characteristic dryness: “Though it’s a beautiful picture, you wouldn’t be happy if these were in your body due to premature anti-aging clinical trial efforts.”

Jeffrey Flier, M.D., the former dean of Harvard Medical School and the Caroline Shields Walker Professor of Medicine at Harvard,

contributed: “Many reasons to be concerned about these claims by a certain scientist at @harvardmed.”

The paper has not been retracted. As of April 2026 it carries no editor’s note. *Aging* has issued no statement. Harvard has issued no statement. The paper continues to accumulate citations, some from Sinclair’s own subsequent papers, in the manner of all papers that exist in the literature whether or not they should.

March 2024: The Cascade

The chemical-cocktail paper was a tremor. The earthquake came nine months later.

On March 5, 2024, Animal Biosciences—a company co-founded by Sinclair and his brother Nick—issued a press release announcing “Leap Years,” a daily dog chew, as “the first supplement proven to reverse aging in dogs.” The active ingredients were variations on the kind of nutraceuticals that have been sold for years in the human longevity market: a fisetin derivative, a quercetin variant, a polyphenol blend.

Kaeberlein responded the same week, on March 3, 2024 (the announcement was timed for embargo lift): “The press release from Animal Biosciences is dishonest. This supplement has not been ‘proven to reverse aging in dogs.’ That is a lie.” A few days later, he expanded the indictment:

I find it deeply distressing that we’ve gotten to a point where dishonesty in science is normalized to an extent that nobody is shocked when a tenured @Harvard professor falsely proclaims in a press release that a product he is selling to pet owners has “reversed aging in dogs.” To me, this is the textbook definition of snake oil salesman.

The phrase “textbook definition of snake oil salesman,” used by a fellow tenured professor about a colleague who had spent a decade as the public face of longevity research, was not the kind of phrase anyone walks back. The press release was edited within forty-eight hours, replacing “proven to reverse aging in dogs” with the softer “reverses age-related deterioration.” But the change did not put the toothpaste back in the tube. *The Wall Street Journal* sent reporters Amy Dockser Marcus and Dominique Mosbergen to talk to anyone in the field who would speak on the record, and many did. The resulting article, published that May, made the problem the field’s problem: not Sinclair’s claim alone but the willingness of an entire industry to allow its most prominent spokesperson to keep selling things he had not proven.

On March 13, 2024—eight days after the Leap Years press release—David Sinclair resigned the presidency of the Academy for Health and Lifespan Research, the field’s flagship professional society, of which he had been a founder. The resignation letter, by all accounts, was brief.

What ER-100 Actually Promises, and Doesn’t

Now the two stories come together.

Twenty-two months after Sinclair resigned the AHLR presidency, the laboratory he runs—working through Life Biosciences, the company he co-founded—received FDA clearance to inject Yamanaka factors into human eyes. The science underneath that clearance is the 2020 *Nature* paper and five years of follow-up. The methylation data are real. The non-human primate work is, by Life Bio’s account in their IND, comprehensive. The trial is registered, the protocol is public, the endpoints are conventional. Twelve patients will get ER-100, and we will know, by 2027 or so, whether it is safe enough

to push to a Phase 2.

What ER-100 promises, very precisely, is this: a small, controlled test of whether a tool from aging research can rescue retinal ganglion cells in two specific eye diseases. It does not promise to make anyone young. It does not promise to lower intraocular pressure, the actual driver of glaucoma progression. It does not promise to reach the optic nerve at all in any patient whose ganglion cells have already finished dying. It is, as Knoepfler puts it, decidedly not the kind of thing your doctor will be writing a prescription for. It is a Phase 1 trial of a gene therapy for two retinal diseases, and that is enough.

Hold the chemical-cocktail paper and the Leap Years dog chew next to ER-100 and ask which one represents the field. The honest answer is: all of them, simultaneously. Sinclair's CV is the field in miniature. The 2003 Sirtris work was real biology that produced an artifact and a \$720 million mistake. The 2020 *Nature* paper is real biology that produced an FDA-cleared IND. The 2023 *Aging* paper and the 2024 dog chew are the parts of the CV that the field has now—belatedly, awkwardly, but really—refused to endorse. The science and the scandals run on parallel tracks because they always have. The field is finally separating them.

Whether ER-100 works is a question for the trial. Whether the field can sustain the next decade of partial-reprogramming research while its most visible researcher continues to publish in his own journal and sell things he has not proven is a different question, and it is the one the next two essays are about. The bird, in the meantime, keeps coming back. The mountain has not gotten any smaller.

Sources and Further Reading

Brenner, C. [@CharlesMBrenner]. (2023, July 12). Today @davidasinclair is telling the world that he has achieved age reversal with chemical

cocktails. He submitted the paper on June 30 & it was accepted on July 4 by a journal of which he is coeditor-in-chief [Post]. X.

Kaeberlein, M. [@mkaeberlein]. (2023, July 15). Sad sign of the times this even needs to be said, but headlines claiming @Harvard researchers developed a chemical approach to reverse aging are not true. Nobody has a drug to reverse aging [Post]. X.

Kaeberlein, M. [@mkaeberlein]. (2024, March 3). The press release from Animal Biosciences is dishonest. This supplement has not been “proven to reverse aging in dogs.” That is a lie [Post]. X.

Knoepfler, P. (2023, July 18). Review of new David Sinclair paper, supplements & anti-aging glitz. *The Niche*. <https://ipscell.com/2023/07/review-of-new-david-sinclair-paper-anti-aging-glitz/>

Kazemian, M., Hu, Y., Liu, Y., Su, R.-C., Daly, T. M., Dwivedi, B., & Caplan, M. L. (2021). The COVID-19 virus may not insert genetic material into human DNA. *Journal of Virology*. [See also Purdue University News, “The COVID-19 virus may not insert genetic material into human DNA, research shows.”]

Life Biosciences, Inc. (2026, January 28). *Life Biosciences announces FDA clearance of IND application for ER-100 in optic neuropathies* [Press release]. <https://www.lifebiosciences.com/>

Lu, Y., Brommer, B., Tian, X., Krishnan, A., Meer, M., Wang, C., Vera, D. L., Zeng, Q., Yu, D., Bonkowski, M. S., Yang, J.-H., Zhou, S., Hoffmann, E. M., Karg, M. M., Schultz, M. B., Kane, A. E., Davidsohn, N., Korobkina, E., Chwalek, K., ... Sinclair, D. A. (2020). Reprogramming to recover youthful epigenetic information and restore vision. *Nature*, 588(7836), 124–129. <https://doi.org/10.1038/s41586-020-2975-4>

Marcus, A. D., & Mosbergen, D. (2024, May). Star scientist’s claim of ‘reverse aging’ draws hail of criticism. *The Wall Street Journal*.

Ocampo, A., Reddy, P., Martínez-Redondo, P., Platero-Luengo, A., Hatanaka, F., Hishida, T., Li, M., Lam, D., Kurita, M., Beyret, E., Araoka, T., Vazquez-Ferrer, E., Donoso, D., Roman, J. L., Xu, J., Rodriguez Esteban, C., Nuñez, G., Nuñez Delicado, E., Campistol, J. M., ... Izpisua Belmonte, J. C. (2016). In vivo amelioration of age-associated hallmarks by partial reprogramming. *Cell*, 167(7), 1719–1733.e12. <https://doi.org/10.1016/j.cell.2016>

Regalado, A. (2026, January 27). The first “aging reversal” drug just cleared the FDA. *MIT Technology Review*.

- Takahashi, K., & Yamanaka, S. (2006). Induction of pluripotent stem cells from mouse embryonic and adult fibroblast cultures by defined factors. *Cell*, 126(4), 663–676. <https://doi.org/10.1016/j.cell.2006.07.024>
- Yang, J.-H., Petty, C. A., Dixon-McDougall, T., Vina Lopez, M., Tyshkovskiy, A., Maybury-Lewis, S., Tian, X., Ibrahim, N., Chen, Z., Griffin, P. T., Arnold, M., Li, J., Martinez, O. A., Behn, A., Rogers-Hammond, R., Angeli, S., Gladyshev, V. N., & Sinclair, D. A. (2023). Chemically induced reprogramming to reverse cellular aging. *Aging (Albany NY)*, 15(13), 5966–5989. <https://doi.org/10.18632/aging.204896>

The Merchant's Century

From Laboratory to Lifestyle Brand

Dr. Jordan Shlain runs five longevity clinics. He has seen the blood panels, the biomarker dashboards, the supplement stacks that cost more than most people's rent. When *The New Yorker* asked him to distill decades of clinical observation into a single insight, he offered this: "Everything you do to improve your health span can improve your life span. Everything you do to improve your life span is fucking bullshit."

Treat that sentence as the operating definition of the line we are about to walk along. There is, for the time being, a real distinction between interventions that help you live better and interventions that claim to help you live longer. The former have evidence. The latter have venture capital. Previous essays in this series established the scientific foundations: why evolution stopped maintaining our bodies after reproduction; how the twelve hallmarks cascade from molecular damage to collapse; how epigenetic clocks translate invisible biochemistry into measurable age. Now we examine what happens when those foundations meet money.

The Genome Racer

Craig Venter's scientific contributions are beyond dispute. In the late 1990s, while the publicly funded Human Genome Project methodically sequenced DNA using laborious clone-by-clone methods, Venter proposed shotgun sequencing—shattering the genome into millions of fragments and reassembling them computationally. The NIH rejected his grant proposal. Venter proved them wrong in 1995

with the first complete bacterial genome, and the method is now standard practice worldwide.

The race that followed became legendary. James Watson (co-discoverer of the structure of DNA) reportedly called Venter “Hitler”; the public consortium supported free data release while Venter’s Celera planned subscriber access and gene patents. His 2001 *Science* paper has accumulated over 10,600 citations. His synthetic biology achievements followed: the 2010 creation of “Synthia,” the first synthetic cell. Caltech’s David Baltimore called it “a technical tour de force” but added: “He has not created life, only mimicked it.” *Scientific American’s* John Horgan was blunter: “Craig Venter is the Lady Gaga of science. . . Hype is what Craig Venter does, and he does it extremely well.”

Then came Human Longevity, Inc. Founded in 2014 with over \$300 million, HLI promised to combine genomics, imaging, and machine learning to extend healthy lifespan. Mayo Clinic’s Michael Joyner offered an assessment: “The whole thing is an example of technology run amok from a belief that if you can measure it, it must be meaningful.” By 2018, Venter had been fired; a trade-secrets lawsuit followed; the valuation crashed from \$1.6 billion to \$310 million. Venter’s 280+ articles and 137,000 citations establish him as one of genomics’ most influential figures. His longevity ventures demonstrate how that influence can be leveraged beyond what the science supports.

The Engineer of Negligible Senescence

Aubrey de Grey entered aging research sideways. Trained in computer science at Cambridge, he earned his PhD by publication after marrying a geneticist—an unusual path outside the traditional hierarchy. His framework, SENS (Strategies for Engineered Negligible

Senescence), identifies seven categories of age-related damage (not too dissimilar to the 12 Hallmarks of aging we visited in a previous essay) with engineering solutions for each: clear cellular garbage, replace lost cells, remove senescent zombies, break crosslinks that stiffen tissues. Where traditional researchers sought to slow damage, de Grey proposed repairing it. Where geriatricians treated diseases, de Grey argued aging itself was the disease.

In 2005, 28 prominent biogerontologists—including Steven Austad, Leonard Guarente, and Linda Partridge—published an extraordinary assessment in *EMBO Reports*. They called SENS “a farrago—a confused mixture” whose claims are “so far from plausible that it commands no respect at all within the informed scientific community.” Yet science rarely delivers clean verdicts. *MIT Technology Review* offered \$20,000 to prove SENS “unworthy of learned debate.” No winner emerged; the judges concluded SENS “does not compel the assent of many knowledgeable scientists; but neither is it demonstrably wrong.”

Individual SENS concepts have since gained mainstream acceptance—senolytics now constitute over thirty clinical trials. The comprehensive program and dramatic timelines remain contested: de Grey predicts a 50% chance of “longevity escape velocity” by the late 2030s. Jan Vijg of the Institute for Aging Research at Albert Einstein College of Medicine calls such predictions “sheer nonsense.”

The Billionaire's Dilemma

Chip Walter's *Immortality, Inc.* chronicles the moment Silicon Valley decided aging was a problem worth solving. Bill Maris, founder of Google Ventures and architect of Calico Labs, told *Bloomberg*: “If you ask me today, ‘Is it possible to live to 500?’ The answer is yes.” Google committed somewhere between \$750 million and \$1.5 billion.

Walter, who began thinking the quest was “crazy,” underwent a conversion. Critics noted his book provides access but pays “far less attention to the views of underlings, dissenters and outsiders with critical insights.”

The dissenting voices paint a different picture. Eric Topol of Scripps described Calico as “hyper-secretive.” Nir Barzilai of the aforementioned Albert Einstein observed: “It seems they are not doing that [delaying aging]. It’s weird they don’t come to us... Nobody from Calico talks to us.” The two most public Calico programs have now posted negative readouts: fosigotifator (ABBV-CLS-7262) failed the HEALEY ALS Phase 2/3 platform trial in January 2025, and AbbVie ended the broader Calico collaboration in November 2025. Altos Labs launched in 2022 with \$3 billion and remains pre-clinical for any human therapeutic; Unity Biotechnology’s senolytic drugs failed Phase 2 in 2020 and Phase 2b in 2025 (more on UNITY in the next essay). The translational gulf swallows venture after venture.

The Quantified Selves

“I’m not a biohacker. I’m not an optimization person. I’m an explorer, about the future of being human.”

– Bryan Johnson

Bryan Johnson sold Braintree Venmo for \$800 million and invested \$300 million into becoming “the most measured person in history.” Project Blueprint involves 100+ daily supplements, precisely 1,977 vegan calories before 11 AM, and—briefly—blood plasma transfusions from his teenage son. Johnson discontinued the transfusions after seeing “no benefits,” stopped rapamycin due to lipid abnormalities, abandoned growth hormone when glucose rose. The pattern suggests iterative elimination of interventions that don’t work—which is how science proceeds, albeit rarely with this publicity.

Johnson released *Don't Die*, a Netflix documentary, on January 1, 2025. By March, a *New York Times* investigation had documented twenty-page non-disclosure agreements binding former employees, complaints about adverse reactions to the Blueprint "Longevity Mix," and Johnson's eventual concession on Dorian Yates's podcast that he had been on testosterone replacement therapy since age 45—a fact previously absent from the documentary's metabolic-purity narrative. The same year he closed a \$60 million investment round that included Kim Kardashian and Paris Hilton.

" Being ruthlessly honest is risky, awkward, and uncomfortable - for nearly all of us. There's really no getting around it."

– Bryan Johnson

Peter Diamandis operates adjacent: co-founder of Human Longevity Inc., Fountain Life, Lifeforce; organizer of the \$101 million XPRIZE Healthspan; host of premium retreats promising "private conversations with Harvard researchers and Nobel laureates." His daily routine involves biometrics from an Oura ring, Apple Watch, and continuous glucose monitor, then meditating while using three red-light-therapy devices. Diamandis functions less as a scientist (although he certainly likes portraying himself as such) than as a catalyst—connecting capital to claims.

Peter Attia is the figure who has executed this arc most successfully. His book *Outlive* (2023) sold over a million copies; his podcast crossed 100 million downloads; his new clinic *Biograph* opened in San Francisco in February 2025 and New York shortly after, with a \$7,500 annual base tier. Attia is the most clinically rigorous of the longevity-as-lifestyle figures and, in part for that reason, the most expensive.

The Boring Truth

S. Jay Olshansky's 2024 *Nature Aging* paper found life-expectancy improvements have decelerated since 1990, from approximately 2.5 years per decade in the 1990s to roughly 1.5 by the 2010s, and effectively to zero in the United States. Olshansky's projection: probability of survival to 100 is unlikely to exceed 15% for women and 5% for men in most long-lived countries during this century. "Most people alive today at older ages are living on time manufactured by medicine," Olshansky observes. "But these medical Band-Aids are producing fewer years of life."

What actually works? A physician told *The New Yorker*: implementing "the no-brainer stuff"—sleep, diet, exercise—can produce a 32-year difference. Supplements add perhaps four. Eric Topol was direct: "If you really want to know something proven to change biologic aging, it's exercise."

A Brisk Trade in Beaks, and the Conquest Deferred

Senolytics is the case study where the marketplace dynamic is cleanest. The science is real. The 2011 paper from Jan van Deursen's lab at the Mayo Clinic showed that genetically clearing senescent cells extended healthspan in progeroid mice; the 2016 follow-up showed the same in naturally aged wild-type mice. James Kirkland's group, also at Mayo, coined the word "senolytic" in 2015 and published the first human pilot trial of dasatinib plus quercetin in idiopathic pulmonary fibrosis in 2019. The molecular biology is well-supported. The biology of cellular senescence is the subject of the next essay.

The supplement industry, however, did not wait for the next essay. Fisetin, a flavonoid from strawberries that showed senolytic activity in mouse studies at roughly 20 mg/kg, is now sold in capsules typically dosed at 100–500 mg per day—approximately one-tenth

to one-twentieth of the dose used in the trials. Quercetin is sold the same way. Neither has a published human trial supporting the doses on the bottle. Both are marketed with images of clean kitchens and phrases like “cellular renewal.”

The most striking instance of the genre is the Animal Biosciences “Leap Years” dog chew, marketed by David Sinclair and his brother Nick. The March 5, 2024 press release described it as “the first supplement proven to reverse aging in dogs.” Matt Kaeberlein responded on X: “The press release from Animal Biosciences is dishonest. This supplement has not been ‘proven to reverse aging in dogs.’ That is a lie.” He added, in a follow-up post: “I find it deeply distressing that we’ve gotten to a point where dishonesty in science is normalized to an extent that nobody is shocked when a tenured @Harvard professor falsely proclaims in a press release that a product he is selling to pet owners has ‘reversed aging in dogs.’ To me, this is the textbook definition of snake oil salesman.” The press release was subsequently edited to “reverses age-related deterioration.” On March 13, 2024, Sinclair resigned the presidency of the Academy for Health and Lifespan Research.

Leonard Hayflick, who discovered the cellular replication limit bearing his name, called anti-aging ventures “a great annoyance.” Before his death in 2024: “Individuals are going to the bank with enormous sums of money gained by persuading people that they’ve found a way to extend your life or make you immortal. . . Everything in the Universe ages with time, and to think that you can reverse it is nonsense.”

Hayflick was wrong, narrowly: a number of the molecular interventions discussed in this collection do, in fact, slow some aspect of cellular aging in some model system. The next essay opens on one such class—the PROTAC senolytics, which target the same biology UNITY chased and the supplement industry counterfeited, and

which work in mice with a cleaner mechanism than anything sold at retail. The chemistry community responded to the same biology by going in the opposite direction from Animal Biosciences—toward more selective, more potent, more clinically testable molecules.

Hayflick was also right, broadly: the conquest remains elusive, the most reliable intervention remains exercise, and the field's enormous accumulation of capital has not yet purchased a proven year of human life. Both readings are correct because they are reading different timescales. On the timescale of any individual investor—five years, ten years, the Series B, the IPO, the write-down—the ventures fail. On the timescale of the field, each failed venture pays for the next round of methodology, instruments, and validated negative results. UNITY's ASPIRE failure narrowed the viable BCL-xL design space; Calico's HEALEY ALS readout taught the field something about what fosigotifator does not do; Sinclair's Leap Years embarrassment may yet do for senolytic supplements what the Sirtris write-down did for resveratrol consumer products.

Sources and Further Reading

- Friend, T. (2017, April 3). Silicon Valley's quest to live forever. *The New Yorker*.
- Friend, T. (2025, August 11). How to live forever and get rich doing it. *The New Yorker*.
- Horgan, J. (2010, May 21). Craig Venter has neither created—nor demystified—life. *Scientific American*.
- Marcus, A. D., & Mosbergen, D. (2024, May). Star scientist's claim of 'reverse aging' draws hail of criticism. *The Wall Street Journal*.
- Olshansky, S. J., Newman, A. B., Guo, J., & Goldman, D. P. (2024). Implausibility of radical life extension in humans in the twenty-first century. *Nature Aging*, 4, 1635–1642.
- Walter, C. (2020). *Immortality, Inc.: Renegade science, Silicon Valley billions, and the quest to live forever*. National Geographic.

THE MERCHANT'S CENTURY

Warner, H., Anderson, J., Austad, S., Bergamini, E., Bredesen, D., Butler, R., Carnes, B. A., Clark, B. F. C., Cristofalo, V., Faulkner, J., Guarente, L., Harrison, D. E., Kirkwood, T., Lithgow, G., Martin, G., Masoro, E., Melov, S., Miller, R. A., Olshansky, S. J., . . . Sierra, F. (2005). Science fact and the SENS agenda. *EMBO Reports*, 6(11), 1006–1008.

Williams, J. (2025, March). Inside the gilded cage of Bryan Johnson's Blueprint. *The New York Times Magazine*.

A More Surgical Scalpel

PROTAC Senolytics, and What UNITY Biotechnology Left Behind

What Cellular Senescence Is, and Why You Want It (Sometimes)

PROTACs

UNITY Biotechnology and its mess

Why Any of This Has to Do with Aging

A Bridge Essay

1. When a distinguished but elderly scientist states that something is possible, he is almost certainly right. When he states that something is impossible, he is very probably wrong.

2. The only way of discovering the limits of the possible is to venture a little way past them into the impossible.

3. Any sufficiently advanced technology is indistinguishable from magic.

— ARTHUR C. CLARKE

Life is a machine that processes information. The claim travels, in the form Schrödinger gave it in *What is Life?* in 1944, through the entire molecular biology revolution: the double helix as a physical encoding, the genetic code as a literal code, transcription and translation as literal acts of reading and writing. The slogan “life is computation” is, in 2026, essentially correct.

If life is information processing, then the disciplines that study how information is processed in the abstract — theoretical computer science, machine learning, quantum computation — will have something to say about biology. If those disciplines have something to say about biology, they will have something to say about aging, which is the slow and systematic failure of the information-processing machinery in question. And if they have something to say about

aging, they will be part of whatever conquest of death humanity ever manages.

This is the argument of Part III. It is not a strong argument, in the sense that it does not specify which technologies will matter, or when, or how much. It is a weak argument in the sense that it only tells you which direction to look. The four essays that follow look in that direction with the same skepticism the rest of the collection has tried to practice elsewhere — which is to say they are about both the substance and the scaffolding of the claim that AI and quantum computing will revolutionize biology.

Why Computers, Why Now

Consider what happens in the human body in the time it takes to read this paragraph. Several million cells will divide, each of which will copy the entire genome, each of which will make a few mistakes that the repair machinery will mostly catch. Several thousand cells will misfold a protein badly enough to flag the chaperones; some smaller number will fail the flag and either get cleared or quietly senesce. The mitochondria of the senescent ones will leak more reactive oxygen species than usual, which will damage the DNA of their neighbours, which will trigger the next round of repair, which will fail at some rate, which will produce the next round of senescence. So. Many. Moving. Parts.

Nobel Laureate Daniel Kahneman, in *Thinking, Fast and Slow*, observed that human working memory holds about four items at a time—the figure comes from Nelson Cowan’s revision of Miller’s older 7 ± 2 —and that anything beyond three or four interacting variables exceeds what unaided cognition can hold without a pencil. We certainly need to keep track of more than four moving parts in biological systems.

The instruments that handle several thousand variables simultaneously are the statistical learning tools that produced the epigenetic clocks of Essay 3, solved the protein structures that AlphaFold2 made available for free in 2021, and now power the protein language models that can predict, for the first time, how multi-protein complexes—including the transcription factor assemblies that the Yamanaka factors recruit to reset the epigenetic clock—fit together and behave. The drugs themselves—senolytics, epigenetic modulators, NAD⁺ precursors, metformin—act on many targets at once, partly because aging acts on many things at once, and partly because the universe doesn't owe medicinal chemistry the favour of clean mechanisms.

At the same time, each of those individual drug-target interactions, taken on its own, is a problem in quantum chemistry: how electrons in a small molecule rearrange when the molecule slots into a protein binding pocket, and what that rearrangement does to binding affinity. The tools that solve those problems—density functional theory, coupled-cluster methods, QM/MM—have been doing it on classical hardware for thirty years, getting steadily better, with no quantum computer involved. The phrase “quantum chemistry” nevertheless appears in pharmaceutical press releases far more often than the word “classical.” Essay 10 is about the conflation.

In June 2025, Insilico Medicine published the GENESIS-IPF Phase 2a trial in *Nature Medicine*: rentosertib, their TNK2 inhibitor for idiopathic pulmonary fibrosis, showed a forced-vital-capacity gain of +98.4 mL versus −20.3 mL on placebo over twelve weeks, in 71 patients across 22 sites in China. The molecule came from Insilico's generative chemistry platform; the target came from their AI target-discovery pipeline; the disease, depending on how one frames it, is age-related. This is the first time an AI-designed molecule against an AI-identified target has shown a positive efficacy signal in a Phase 2a trial of an age-related disease. It is also a small Phase 2a in a single

country, awaiting confirmatory work, in a disease whose history of failed late-stage trials is long enough to fill its own book. Both framings are correct, and the rest of the industry has clearly decided to bet on the optimistic one.

In October 2025 Eli Lilly extended its engagement with Insilico to up to \$2.75 billion. In January 2026, at the J.P. Morgan Healthcare Conference, NVIDIA and Lilly announced a five-year, billion-dollar AI co-innovation lab in the Bay Area. The same week, Anthropic launched Claude for Healthcare, building on its Claude for Life Sciences product from October 2025; partners include AstraZeneca, Sanofi, Genmab, and Novo Nordisk. OpenAI announced ChatGPT for Health within days of the Anthropic launch. In April 2026 Anthropic acquired Coefficient Bio, a stealth drug-discovery startup founded by ex-Genentech computational biologists, for roughly \$400 million in stock — a price that works out to about fifty million dollars per employee, which tells you everything about how the field is currently pricing the closing of that loop.

Classical machine learning has, by 2026, transformed aging biology. It built the clocks. It solved the structures. It identified rentosertib's target. The transformation is uneven, contested, and in places oversold, but it has happened, and several billion dollars of capital have arrived to make sure it keeps happening. Quantum computing has not yet transformed anything in biology, despite a decade of press releases insisting otherwise. Both technologies will take longer than their marketing departments promise to deliver what their marketing departments promise. The gap between what is being delivered and what is being promised — in 2026, in both fields — is what Part III is about.

Sources and Further Reading

Abramson, J., Adler, J., Dunger, J., Evans, R., Green, T., Pritzel, A., Ronneberger, O., Willmore, L., Ballard, A. J., Bambrick, J., Bodenstein, S. W.,

- Evans, D. A., Hung, C.-C., O'Neill, M., Reiman, D., Tunyasuvunakool, K., Wu, Z., Žemgulytė, A., Arvaniti, E., . . . Jumper, J. M. (2024). Accurate structure prediction of biomolecular interactions with AlphaFold 3. *Nature*, 630(8016), 493–500.
- Callies, M., & colleagues. (2025). Causality-informed reproductive clocks reflect the impact of childhood maltreatment on epigenetic aging. *Communications Medicine*.
- Clarke, A. C. (1973). *Profiles of the future: An inquiry into the limits of the possible* (Rev. ed.). Harper & Row.
- Hayes, T., Rao, R., Akin, H., Sofroniew, N. J., Oktay, D., Lin, Z., Verkuil, R., Tran, V. Q., Deaton, J., Wiggert, M., Badkundri, R., Shafkat, I., Gong, J., Derry, A., Molina, R. S., Thomas, N., Khan, Y. A., Mishra, C., Kim, C., . . . Rives, A. (2025). Simulating 500 million years of evolution with a language model. *Science*, 387(6736), 850–858.
- López-Otín, C., Blasco, M. A., Partridge, L., Serrano, M., & Kroemer, G. (2023). Hallmarks of aging: An expanding universe. *Cell*, 186(2), 243–278.
- Schrödinger, E. (1944). *What is life? The physical aspect of the living cell*. Cambridge University Press.
- Vakili, M. G., & colleagues. (2025). Quantum-classical generative models for the discovery of KRAS inhibitors. *Nature Biotechnology*.
- Xu, Z., Ren, F., Wang, Z., Hou, Y., Wang, H., Wang, Y., Liu, J., Zhao, Y., Pun, F. W., Aliper, A., Zhavoronkov, A., & Dai, H. (2025). A generative AI-discovered TNIK inhibitor for idiopathic pulmonary fibrosis: A randomized phase 2a trial. *Nature Medicine*, 31, 2602–2610.

All Theories are Wrong, but Quantum Improvement is Real

What Quantum Computing Actually Is, and What It Isn't Yet

In December 1981, at MIT's Endicott House, Richard Feynman told a room of physicists what he thought a computer should be. "Nature isn't classical, dammit, and if you want to make a simulation of nature, you'd better make it quantum mechanical, and by golly it's a wonderful problem, because it doesn't look so easy."

Forty-five years later, the quote shows up in everyone's slide deck, including the slide decks of CEOs warning about Q-day — a vague, evil future where your bank account is read by a refrigerator in Santa Barbara. So there is no shortage of essays, videos, and TED talks on what a quantum computer is. This is mine. It is opinionated about the framing because I think the framing is where most of the confusion lives. We will start, deliberately, not with the problems quantum computers solve, but with what they *are*, computationally.

What Quantum Computation Actually Is

If one wishes to be formally pretentious about computation — which is to say, correct — one borrows from Alan Turing and Alonzo Church. To compute is to map one set of symbols to another set of symbols, in a finite number of steps, by an explicit rule. The rule is the algorithm. John S. Conery, professor emeritus of *biology* at the University of Oregon (yes, biology), is a great educator in computer science and writes in one of his textbooks: "A more formal definition is that . . . a step is a symbol manipulation that transforms one set of symbols into a new set of symbols."

Asked to compute 57×3 , you are mapping the string “ 57×3 ” to a string “171”. Asked to render this paragraph in **red**, your laptop is mapping bits that represent text to other bits that represent the same text, marked red.³ Thus, a computer is anything that can read and write symbols according to some algorithm — ideally a “useful” one. It must also do so in a finite number of “steps.” Turing himself writes: “We may compare a man in the process of computing a real number to a machine which is only capable of a finite number of conditions $q_1, q_2, \dots, q_l \dots$ The machine is supplied with ... a ‘tape’ (the analogue of paper) running through it, and divided into sections (called ‘squares’) each capable of bearing a ‘symbol.’” This 1936 paper — “On Computable Numbers, with an Application to the Entscheidungsproblem” — did several things at once. It defined the machine. It demonstrated that there are functions no such machine can compute (the halting problem). It established the equivalence between Turing’s machines and Church’s λ -calculus, giving us the Church-Turing thesis. And it provided a formalization simple enough that Gödel would later call Turing’s analysis “a kind of miracle.”

The abstraction of information into 0s and 1s is, notably, not an afterthought of the Turing-Church model. It is built into the definition of a computing machine from the start:

“If an a-machine prints two kinds of symbols, of which the first kind (called figures) consists entirely of 0 and 1 [...], then the machine

³Note that deciding the best way to represent my text as bits; the best way to save those bits in physical memory; the best way to allocate memory for the old and new representation; the best way to schedule the call to the algorithm that transforms one representation to another given all the other algorithms that are running in my computer; the best way to show me that representation in my screen, etc., are all generally regarded as slightly nontrivial problems, which is why we have the field of computer engineering, among others.

will be called a computing machine. If the machine is supplied with a blank tape and set in motion, starting from the correct initial m -configuration, the subsequence of the symbols printed by it [...] will be called the sequence computed by the machine."

This is where the idea of bits, and eventually qubits, comes from.

Anyways, after Turing finished his decade of revolutionizing mathematics (at the very least inasmuch as his work simplified and enabled Gödel's incompleteness proofs), computer science (arguably, he *created* the discipline), philosophy (according to some)⁴, and biology (both in the Promethean sense of "biology is computation, and he formalized computers, so he formalized biology," and in the literal sense that he wrote papers on mathematical biology during his lifetime — his 1952 paper on morphogenesis is the founding document of pattern formation theory), he also helped defeat the Nazis by working as a cryptographer for the British. For this, of course, he was rewarded with a well-deserved life of academic recognition, and comodity.

pause for laughs

But I digress.

The Pen That Writes on Every Line

Quantum computers also just map a set of symbols (usually, but not necessarily, 1s and 0s) to another set of symbols (usually, but not necessarily, 1s and 0s). The main difference is that for classical computers, the initial set of symbols is essentially just treated as symbols on a piece of paper, which we can only read and modify one at a time; in a quantum computer the input symbols spring out

⁴See Scott Aaronson, "Why Philosophers Should Care About Computational Complexity" — scottaaronson.com/papers/philos.pdf.

of the paper and get treated as elements of a *state vector*. This new treatment gives us an exponentially larger piece of paper to work on. That new piece of paper and its properties get very fancy names like Hilbert space, qubit space, complete inner product vector space, etc., but the two most salient properties are the following:

- This new space of vectors — vector space, if you will — has a dimension that grows exponentially with the number of symbols we started with. For the reader who pictures vectors as nice little arrows in a Cartesian plane, think about this: in a Cartesian plane each vector can be completely specified by its coordinates, which are just two numbers (x-coordinate and y-coordinate). This is a 2-dimensional vector. Simple enough. Now imagine we instead needed 2^n numbers. Can you picture it? No? Neither can physicists, but that doesn't matter. We don't need the picture to be *true*. We need the picture to be *empirically adequate*: to give us the right predictions when we measure. Hilbert space is not a place. The state vector is not a thing. They are tools we use to predict what a measurement will read. Whether they “exist” is a question for philosophers and Twitter.
- Right after springing out of the paper, our new state vector obeys all the linear algebra of quantum mechanics: in particular, it can be put into a superposition of states, and multiple components of the state vector can be entangled with each other — “entangled” meaning, roughly, that the joint state cannot be factored into a product of independent parts.

Entanglement is formally a very obvious consequence of the tensor product structure of Hilbert spaces. This phrase is not strictly relevant for the next explanation, but everyone should hear it once in their lifetime.

How is that any better?

At the end of the day, all of these sophisticated properties of state vectors are new tools to write and erase symbols. Imagine you had a pen that could write on every line of your piece of paper at once. Imagine that *what* it writes depends on what is currently written on each line: if there is a sentence about raccoons, it might complete it with a fun fact about polar bears; if there is a line written in iambic pentameter, it will add a new phrase in trochaic octameter. Imagine you have many such pens and they each act on all lines in a different way. Some pens will make every line reference something that was said in an earlier line (entanglement). If you want to write a poem, you might start with a blank page, apply pen A to set up the meter, apply pen B to seed the rhymes, apply pen C to entangle the stanzas, and then read what came out.

This is, more or less, what David Deutsch did in 1985. He took Turing’s machine — the head, the tape, the squares, the symbols — and replaced the tape with a Hilbert space and the head’s transitions with unitary operators. The squares of the tape are now qubits; the head’s instruction set is now a finite collection of unitary gates; the act of computation is now applying those gates, which is to say applying those pens, in some sequence. Deutsch’s universal quantum computer is the Turing machine in which the paper is exponentially larger than the symbols it carries, and the writing tools are correlated. This is the architecture. Everything else — Shor, Grover, the variational eigensolvers, the entire NISQ-era zoo — is a question of which sequence of pens to use.

Surely, if we had such writing tools, they would not be even remotely useful for our current writing practices. You cannot dictate a grocery list to a pen that mostly speaks in entanglement. Accordingly, quantum properties are not at all immediately useful for “typical” algorithm design paradigms. As Leslie Valiant once put

it, in characteristic deadpan, “most of what we want computers to do, we already know how to make computers do.” The quantum pens are a wild and beautiful answer to a question almost no one was asking until Feynman asked it.

There are a few quirky requirements that our new writing tools must fulfil — they must be unitary transformations, and all state vectors must be L^2 -normalized — and all those quirkinesses are why quantum computing is the kind of field where the entry exam is a graduate course in linear algebra and the consolation prize is a graduate course in error correction.

The Input Problem and the Output Problem

It is also highly nontrivial to make the first set of symbols “spring out of the page” (i.e., input state preparation), as is extracting information from the final form of our state vector — we need to make it collapse back onto the page. State vectors are first and foremost mathematical formalisms; we don’t have direct access to them, and they don’t exist in the sense that we cannot extract all the information at once. There are rules for this, stipulated by the same pesky formalism: Heisenberg’s uncertainty, Born’s rule, the no-cloning theorem, and many others. The main point is this: when we “measure” a state we lose information. Yes, we get our final n symbols back, but by this point we had been working with 2^n amplitudes. Almost all of them are gone the moment we look.

This is why quantum algorithm design is hard. The job is not to fill the exponentially large piece of paper with information. The job is to arrange that information so that the single classical readout we get at the end — one outcome, drawn from a probability distribution the algorithm shaped — tells us what we wanted to know. Shor’s algorithm is exactly this trick, executed for the problem of finding

the period of a modular exponential, which is exactly the trick that breaks RSA.

The Quantum Decoherence Problem

Niels Bohr, asked once what quantum mechanics says about reality, reportedly replied that there is no quantum world, only an abstract quantum description; the job of physics is not to find out how nature is, but only what we can say about it. This is, again, the constructive empiricist line, and it is also the cleanest way to introduce decoherence.

The state vector lives, briefly, in the abstract description. It does not live, ever, in the world. The world is a thermodynamic system at room temperature (or a cryostat at fifteen millikelvin, which is colder than outer space but still not zero) full of phonons and stray photons and microscopic interactions that the state vector — which is, again, our tool, not a thing — cannot remain isolated from. Each interaction is, formally, a measurement. Each measurement collapses some of the amplitudes back into the page. The page asserts itself.

This is decoherence, and it is the engineering enemy of quantum computing. A qubit's coherence time is the duration over which the page has not yet asserted itself enough to corrupt the writing. On Google's Willow processor, in 2024, the average coherence time was 68 microseconds, plus or minus 13. A surface-code cycle takes about 1 microsecond. So you have, very roughly, 68 cycles of writing before the universe edits your draft.

This is why error correction matters, and why “below threshold” — a phrase we are about to spend some time with — is the only milestone that has so far justified the price of the field's coffee.

Feynman to Shor: A Forty-Five-Year Argument

The argument that this strange machinery might be useful was made, in stages, by three people.

In 1981 Feynman pointed out that simulating a quantum system on a classical computer requires resources that scale exponentially with the system size, which is why we cannot simulate a modest molecule on a modest cluster. His proposed fix: build a computer that is itself quantum mechanical.

In 1985 Deutsch wrote down what a universal quantum computer would look like — the Turing machine with the Hilbert-space tape — and proved that it could simulate any physical process. The architecture became, in principle, a thing.

In 1994 Peter Shor showed that this thing could factor large integers in polynomial time. RSA encryption rests on the assumption that factoring is hard. Shor's algorithm broke that assumption, in theory. The intelligence services of every country took notice. Funding followed.

This is the argument: nature is quantum, so a quantum machine simulates it efficiently, and along the way, that machine happens to break the cryptography that protects most of the internet. Three results, thirteen years apart, all still load-bearing.

What's Real

The genuine progress is quieter than the press releases.

In December 2024, Google's Quantum AI team published a result on its Willow processor that the field has been chasing since the mid-1990s. By arranging 105 superconducting qubits into a surface-code

lattice, they showed that scaling the code from a 3×3 to 5×5 to 7×7 patch of physical qubits cuts the logical error rate roughly in half each time. The encoded logical qubit lives more than twice as long as the best of the underlying physical qubits.

This is what “below threshold” means: the more qubits you throw at a logical qubit, the better it gets, exponentially. Below threshold is what error correction was always supposed to do, and it has now, after thirty years, finally done it on real hardware. Whether the engineering can be scaled to the millions of qubits a useful machine would need is a separate question. But the regime exists.

This sits inside what John Preskill called the NISQ era — noisy intermediate-scale quantum — and it is the first crack in the wall between NISQ and fault tolerance. Below threshold is not a useful quantum computer. It is the proof that one is, in principle, buildable.

What's Noise

Almost everything else.

In 2019 Google’s Sycamore processor performed a random circuit sampling task in 200 seconds and claimed it would have taken a classical supercomputer 10,000 years. Within months, IBM announced a classical algorithm that would do it in 2.5 days. Sycamore did something. It did not do something useful.

In 2023 IBM published a Nature paper titled “Evidence for the utility of quantum computing before fault tolerance.” The processor simulated a 127-qubit kicked Ising model. Within weeks, Begušić and Chan reproduced the result on a single core of a laptop using sparse Pauli dynamics; the more thorough follow-up with Gray, in *Science Advances*, did so with several methods at once. The headline

aged badly.

Random circuit sampling and boson sampling — Jiuzhang’s photonic version — are real demonstrations that a quantum device is doing something a classical computer struggles to mimic exactly. They are not demonstrations that the quantum device is doing anything anyone wanted done. Supremacy is not utility. The two get conflated in the press because both involve the word “quantum” and a large number, and most readers are not going to read the rebuttal preprint.

This is not a complaint about the field. The field knows. The conflation lives in the secondary coverage.

What’s Probably Wrong

In February 2025 Microsoft published a Nature paper announcing Majorana 1, a chip the company described as the first powered by topological qubits. Topological qubits, if they exist, would be intrinsically protected from local noise — the holy grail of qubit design. The press release outpaced the paper. Nature’s editors attached a note clarifying that the published results did not, in fact, demonstrate Majorana zero modes; they demonstrated a platform on which such modes might one day be sought.

A month later, at the APS Global Physics Summit, Henry Legg of St Andrews said the quiet part out loud: “Any company claiming to have a topological qubit in 2025 is essentially selling a fairy tale that undermines the field of quantum computation.” He had two preprints arguing that Microsoft’s topological gap protocol — the test it uses to identify Majoranas — delivers false positives, and that the raw conductance data are too disordered to support the claim.

There is also history. In 2018 a Microsoft-funded group at Delft published a Nature paper claiming a Majorana signature. In 2021 they retracted it. The 2025 device may be different. The protocol is in dispute. The case is still open. The point is that the field has antibodies, and they are working.

The Shor Gap

Here is the gap between fairy tale and fact, in numbers.

In 2019, Craig Gidney and Martin Ekerå estimated that breaking RSA-2048 with Shor's algorithm would require roughly 20 million noisy physical qubits running for eight hours. In May 2025, Gidney revisited the problem and brought the estimate down to fewer than one million noisy qubits running for less than a week. Algorithmic and error-correction improvements — approximate residue arithmetic, yoked surface codes, magic state cultivation — did the work. A twentyfold reduction in qubit cost in six years, by one researcher, on the same hardware assumptions.

The current frontier is around a hundred physical qubits and a dozen logical ones. The gap between a hundred physical and a million is four orders of magnitude. The gap between a logical-qubit memory that lives twice as long as a physical qubit and one that runs Shor's algorithm to completion is more than that.

In August 2024, NIST finalized its first three post-quantum cryptography standards. Google has announced plans to phase RSA and ECC out of its products by 2030. The cryptographic transition has begun, on a timescale that assumes quantum computers will not arrive in time to read this decade's traffic. This is the field disciplining itself. It is also the field hedging.

Gidney's number can move again. It has moved twentyfold in six years. It can move twentyfold more, or it can stop moving, or the hardware can stall before the algorithms catch up. Nobody knows. Acting as though nobody knows is what serious people are doing.

Every Hundred Years

All theories are wrong. Some are useful. Quantum mechanics is the most useful wrong theory ever written down. Quantum computing is the long, expensive, beautiful experiment of finding out how useful.

Sources and Further Reading

Aaronson, S. (2013). *Quantum computing since Democritus*. Cambridge University Press.

Aaronson, S. (2024, September). *Quantum computing: Between hope and hype*. Talk delivered at UCLA.

Arute, F., Arya, K., Babbush, R., Bacon, D., Bardin, J. C., Barends, R., Biswas, R., Boixo, S., Brandao, F. G. S. L., Buell, D. A., Burkett, B., Chen, Y., Chen, Z., Chiaro, B., Collins, R., Courtney, W., Dunsworth, A., Farhi, E., Foxen, B., ... Martinis, J. M. (2019). Quantum supremacy using a programmable superconducting processor. *Nature*, 574(7779), 505–510. <https://doi.org/10.1038/s41586-019-1666-5>

Begušić, T., Gray, J., & Chan, G. K.-L. (2024). Fast and converged classical simulations of evidence for the utility of quantum computing before fault tolerance. *Science Advances*, 10(3), eadk4321. <https://doi.org/10.1126/sciadv.adk4321>

Conery, J. S. (2021). *Explorations in computing: An introduction to computer science and Python programming*. Chapman and Hall/CRC.

Deutsch, D. (1985). Quantum theory, the Church-Turing principle and the universal quantum computer. *Proceedings of the Royal Society A*, 400(1818), 97–117. <https://doi.org/10.1098/rspa.1985.0070>

Feynman, R. P. (1982). Simulating physics with computers. *International Journal of Theoretical Physics*, 21(6/7), 467–488.

- Gidney, C. (2025). How to factor 2048 bit RSA integers with less than a million noisy qubits. arXiv:2505.15917.
- Gidney, C., & Ekerå, M. (2021). How to factor 2048 bit RSA integers in 8 hours using 20 million noisy qubits. *Quantum*, 5, 433. <https://doi.org/10.22331/q-2021-04-15-433>
- Google Quantum AI and Collaborators. (2025). Quantum error correction below the surface code threshold. *Nature*, 638(8052), 920–926. <https://doi.org/10.1038/s41586-024-08449-y>
- Kim, Y., Eddins, A., Anand, S., Wei, K. X., van den Berg, E., Rosenblatt, S., Nayfeh, H., Wu, Y., Zaletel, M., Temme, K., & Kandala, A. (2023). Evidence for the utility of quantum computing before fault tolerance. *Nature*, 618(7965), 500–505. <https://doi.org/10.1038/s41586-023-06096-3>
- Microsoft Azure Quantum. (2025). Interferometric single-shot parity measurement in InAs–Al hybrid devices. *Nature*, 638(8051), 651–655.
- Nielsen, M. A., & Chuang, I. L. (2010). *Quantum computation and quantum information* (10th anniversary ed.). Cambridge University Press.
- Preskill, J. (2018). Quantum computing in the NISQ era and beyond. *Quantum*, 2, 79. <https://doi.org/10.22331/q-2018-08-06-79>
- Shor, P. W. (1997). Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer. *SIAM Journal on Computing*, 26(5), 1484–1509. <https://doi.org/10.1137/S0097539795293172>
- Turing, A. M. (1937). On computable numbers, with an application to the Entscheidungsproblem. *Proceedings of the London Mathematical Society*, s2-42(1), 230–265. <https://doi.org/10.1112/plms/s2-42.1.230>
- Turing, A. M. (1952). The chemical basis of morphogenesis. *Philosophical Transactions of the Royal Society B*, 237(641), 37–72.
- van Fraassen, B. C. (1980). *The scientific image*. Oxford University Press.

Quantum Everything? No. Quantum Bio-chemistry? Yes.

Cytochrome P450, QM/MM, and the Industrial Press-Release Layer

The Hype

Pfizer-IBM, Roche-Quantinuum, Boehringer-PsiQuantum, Sanofi-SandboxAQ. VQE demos stuck at H₂, LiH, BeH₂, H₆-H₁₂ chains; benzene is already out of reach. Reiher et al. PNAS 2017 — nitrogenase at 1M physical qubits, days of runtime, 2030s earliest.

The Sharpening Tools

The 2024 Nobels, Protein Language Models, and What They Don't Yet Tell You

In 2024, the Nobel committees did something unusual. They gave both the Chemistry and Physics prizes to people who never touched a test tube.

David Baker at the University of Washington shared the Chemistry prize with Demis Hassabis and John Jumper at Google DeepMind for “protein structure prediction and design.” Geoffrey Hinton at the University of Toronto and John Hopfield at Princeton University took Physics for neural networks.

“It’s a remarkable recognition,” Hassabis said after the announcement, “that computational methods have become central to understanding the physical world.”

The awards do signal a real shift. The tools that crack biology’s hardest *prediction* problems now come from computer science. Whether they will crack biology’s hardest *design* problems—or its hardest *generalization* problems—is a different question, and the rest of this essay is about that gap.

AlphaFold: Structure From Sequence

Proteins are chains of amino acids—20 different molecular building blocks, each with distinct chemical properties. Unlike DNA and RNA, which are built from nucleotides (sugars, phosphates, and bases), proteins fold into intricate three-dimensional shapes that determine their function.

“The folding problem is this,” John Jumper explained in a 2021 *Nature* paper. “You have a sequence of amino acids. How does it fold into a specific three-dimensional structure?”

The problem has haunted biologists since the 1960s. A typical

protein has billions of possible shapes it could adopt. Yet in cells, proteins fold correctly in milliseconds. Cyrus Levinthal pointed out the paradox in 1969: if a 100-amino-acid protein sampled every possible shape at a billion tries per second, the universe would end before it found the right one.




FIGURE 8. Predicted protein structure, taken from Google DeepMind’s official website.

For decades, solving a single protein structure required years of painstaking work—growing crystals, firing X-rays, solving mathematical phase problems. Entire careers were built on determining one structure.

AlphaFold2 solved 200 million of them.

How AlphaFold Works

“We didn’t try to simulate the physics of protein folding,” Jumper told *Science* in 2022. “We learned the rules from evolutionary data.”

AlphaFold2 uses what researchers call an Evoformer architecture—a neural network design that processes evolutionary information. In machine learning, “architecture” refers to how a model is structured: what computational layers it uses and how information flows through them.

The Evoformer is based on transformers, the same neural network

architecture that powers ChatGPT. According to Ashish Vaswani, who led the team that invented transformers at Google in 2017, the key innovation is attention: “The model learns which parts of the input are relevant to each other, even if they’re far apart.”

For proteins, this means AlphaFold can identify that amino acid 42 might interact with amino acid 187, even though they’re distant in the linear sequence.

The model’s input is a multiple sequence alignment—evolutionarily related proteins lined up position by position. Think of it like comparing cookie recipes from different bakers. Human hemoglobin is one recipe, mouse hemoglobin is another. Each baker substitutes ingredients, but the cookies taste similar.

ESM: Learning Protein Grammar

Meta’s Evolutionary Scale Modeling project took a different route. Instead of learning structure directly, it learned the language of proteins.

“ESM-2 is trained on 65 million protein sequences using masked language modeling,” said Alexander Rives, who led the project at Meta. “The same technique behind GPT and BERT.”

The training works like this: take a protein sequence, randomly hide some amino acids, and force the model to predict what’s missing. In a sequence like MKTAYIAKQR, the model sees MK?AYI?KQR and learns to fill the gaps.

“It’s exactly how language models learn,” Rives explained. “By predicting masked words in sentences, GPT learns grammar. By predicting masked amino acids, ESM learns the biochemical rules of proteins.”

ESM-2 never saw protein structures during training. It learned only from sequences. Yet according to a 2023 *Science* paper by Zeming Lin and colleagues at Meta, “the model’s learned representations encode rich structural information.”

They proved it by building ESMFold—adding a lightweight structure-

prediction module on top of ESM-2's frozen representations. Lin's team reported that ESMFold "predicts structures at accuracy competitive with AlphaFold2, but 60 times faster."

More striking: ESMFold works on orphan proteins—proteins with few evolutionary relatives. AlphaFold struggles with these because it needs multiple sequence alignments. ESMFold needs only the sequence itself.

"The transformer learned long-range dependencies from statistics alone," Lin said. "Certain amino acid combinations appear together not because evolution selected them in one protein, but because they form stable structural motifs everywhere."

The model discovered folding rules by reading sequences, without ever seeing what folded proteins look like.

Designing What Evolution Never Tried

David Baker's lab at the University of Washington inverts the problem. Instead of predicting structure from sequence, they generate sequences that will fold into desired structures.

"We can now design proteins for functions that evolution never encountered," Baker said in his Nobel lecture. His group has created novel folds, protein cages for drug delivery, and therapeutic binders. Several are in clinical trials.

"This isn't incremental," Baker added. "It's orthogonal to four billion years of natural selection."

A useful corrective is also worth stating. The *de novo* binders that the Baker lab has reported have published success rates in the neighborhood of 15–19% in their best campaigns—meaning roughly four out of five computationally designed binders, screened experimentally, do not work as designed. This is a genuinely remarkable hit rate by historical standards. It is not a solved problem. The Nobel was awarded for prediction; design remains an active and humbling research frontier.

Aging: A Plausible Frontier

Aging researchers have watched these developments closely.

“Aging is fundamentally a systems problem,” said Carlos López-Otín at the University of Oviedo, whose 2013 *Cell* paper defined the hallmarks of aging. “Genomic instability, telomere attrition, cellular senescence—these aren’t independent failures. They form cascading networks where each mechanism accelerates the others.”

The causal graph connecting molecular damage to organismal collapse involves thousands of interacting proteins. Human intuition breaks down at that scale.

“This is exactly where protein language models dominate,” said Alex Zhavoronkov, founder of Insilico Medicine. “Aging is the most complex disease imaginable—which is exactly why it requires AI.”

Take this kind of statement carefully. The companies building these tools have a commercial interest in describing them as essential. That does not make them not essential. It does mean the strongest claims warrant the strongest evidence, and the strongest evidence remains thinner than the strongest claims.

Targeting Zombie Cells

Take senescent cells. Judith Campisi at the Buck Institute discovered that these cells stop dividing but remain metabolically active, secreting inflammatory signals that damage surrounding tissue.

“We call it the senescence-associated secretory phenotype, or SASP,” Campisi explained in a 2008 *PLoS Biology* paper. “In young organisms with robust immune systems, SASP signals attract immune cells that clear the damaged cell. But in older organisms, the immune response fails. Senescent cells accumulate.”

Traditional drug discovery screens compounds against known protein targets. But many senescence markers are membrane proteins or intrinsically disordered regions—protein segments that lack stable structure and are historically intractable to crystallography.

“AlphaFold changes the game,” Zhavoronkov said. “We can

predict structures for previously unsolved markers, identify binding pockets computationally, and design molecules targeting them—without crystallizing anything.”

The Loop, Not the Leap

Insilico Medicine’s clinical pipeline is the test case. The company’s lead asset is rentosertib (formerly INS018_055), a TNK2 (TNIK) inhibitor for idiopathic pulmonary fibrosis discovered through its generative chemistry platform. In June 2025, Insilico published the GENESIS-IPF Phase 2a trial in *Nature Medicine*: 71 patients across 22 sites in China, the 60 mg once-daily group showed a forced-vital-capacity gain of +98.4 mL versus –20.3 mL on placebo over twelve weeks. This is, depending on how you frame it, either the first AI-designed molecule against an AI-identified target with a positive efficacy signal in an age-related disease, or a small Phase 2a in a single country awaiting confirmatory work. Both framings are correct.

It is the right kind of result for the right kind of essay. The beak is sharper than it was in 2019, when Insilico claimed a 21-day fibrosis lead. The ascending pipeline is real. None of it has yet metabolized into an approved drug, in IPF or in any age-related disease, and the field’s track record suggests one should withhold the celebration until the mountain has actually been chiseled.

What Changed

Fifty years ago, structural biology was artisanal. One structure, one career.

Then AlphaFold solved 200 million structures and released them for free. ESMFold generates predictions in seconds on standard hardware.

“The bottleneck shifted,” López-Otín said in a 2023 *Cell* review. “Structure is no longer rate-limiting. The constraint now is functional validation—which predictions actually matter for aging?”

That's where the next generation of models is headed. Not just structure prediction, but function prediction. Not proteins in isolation, but proteins in context—bound to ligands, modified by cellular machinery, interacting with membranes. AlphaFold3, released in May 2024, is a step in this direction; whether it generalizes the way AlphaFold2 did is a question whose answer will not arrive on the timeline its press releases prefer.

The 2024 Nobels were awarded for tools that predict structure from sequence. The open problems—design that works, function prediction, generalization to new chemistry, integration with experiment—are larger than the problems that were solved. The bird's beak is genuinely sharper. The mountain has not noticed.

Sources and Further Reading

- Baek, M., DiMaio, F., Anishchenko, I., Dauparas, J., Ovchinnikov, S., Lee, G. R., Wang, J., Cong, Q., Kinch, L. N., Schaeffer, R. D., Millán, C., Park, H., Adams, C., Glassman, C. R., DeGiovanni, A., Pereira, J. H., Rodrigues, A. V., van Dijk, A. A., Ebrecht, A. C., . . . Baker, D. (2021). Accurate prediction of protein structures and interactions using a three-track neural network. *Science*, 373(6557), 871–876. <https://doi.org/10.1126/science.abj875>
- Coppé, J.-P., Patil, C. K., Rodier, F., Sun, Y., Muñoz, D. P., Goldstein, J., Nelson, P. S., Desprez, P.-Y., & Campisi, J. (2008). Senescence-associated secretory phenotypes reveal cell-nonautonomous functions of oncogenic RAS and the p53 tumor suppressor. *PLoS Biology*, 6(12), 2853–2868. <https://doi.org/10.1371/journal.pbio.0060301>
- Jumper, J., Evans, R., Pritzel, A., Green, T., Figurnov, M., Ronneberger, O., Tunyasuvunakool, K., Bates, R., Žídek, A., Potapenko, A., Bridgland, A., Meyer, C., Kohl, S. A. A., Ballard, A. J., Cowie, A., Romera-Paredes, B., Nikolov, S., Jain, R., Adler, J., . . . Hassabis, D. (2021). Highly accurate protein structure prediction with AlphaFold. *Nature*, 596(7873), 583–589. <https://doi.org/10.1038/s41586-021-03819-2>
- Lin, Z., Akin, H., Rao, R., Hie, B., Zhu, Z., Lu, W., Smetanin, N., Verkuil, R., Kabeli, O., Shmueli, Y., dos Santos Costa, A., Fazel-Zarandi, M., Sercu, T., Candido, S., & Rives, A. (2023). Evolutionary-scale prediction of atomic-level protein structure with a language model. *Science*, 379(6637),

1123–1130. <https://doi.org/10.1126/science.ade2574>

López-Otín, C., Blasco, M. A., Partridge, L., Serrano, M., & Kroemer, G.

(2013). The hallmarks of aging. *Cell*, 153(6), 1194–1217. <https://doi.org/10.1016/j.cell.2013.05.041>

López-Otín, C., Blasco, M. A., Partridge, L., Serrano, M., & Kroemer, G.

(2023). Hallmarks of aging: An expanding universe. *Cell*, 186(2), 243–278. <https://doi.org/10.1016/j.cell.2022.11.001>

Radford, A., Narasimhan, K., Salimans, T., & Sutskever, I. (2018). Improving language understanding by generative pre-training. OpenAI.

Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, L., & Polosukhin, I. (2017). Attention is all you need. *Advances in Neural Information Processing Systems*, 30.

Xu, Z., Ren, F., Wang, Z., Hou, Y., Wang, H., Wang, Y., Liu, J., Zhao, Y., Pun, F. W., Aliper, A., Zhavoronkov, A., & Dai, H. (2025). A generative AI-discovered TNIK inhibitor for idiopathic pulmonary fibrosis: A randomized phase 2a trial. *Nature Medicine*, 31, 2602–2610.

The Administrative Revolution

Anthropic, OpenAI, and What Healthcare AI Actually Did in 2025–26

Coda: That's a Hell of a Bird

Author's Note

Acknowledgments