Chapter 11- Extending maximal human lifespan

This chapter uses the model to bring insight to an immortal question.

Extending maximum human lifespan … the quest of many over the ages. And all have failed. Perhaps it is a fool's errand. Even monastics and meditators with healthy and regular lifestyle, physical work, spiritual practice, and natural modest meals don't live way past 100.

The limit of about 120 years hasn't budged , in stark contrast to the median lifespan which jumped in the past two centuries by 50 years, about 6 hours a day. The age of the world-record oldest living human (always a female) has risen by about a year per decade since 1960.

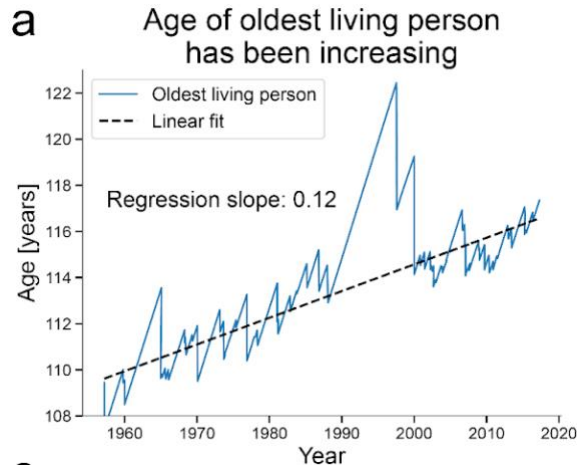

Even optimal lifestyles can not extend maximal lifespan by much. As shown in the figure below, the lifetime gain of those scoring the highest on lifestyle decreases from 6 years at age 40 to less than 2 years at age 100.

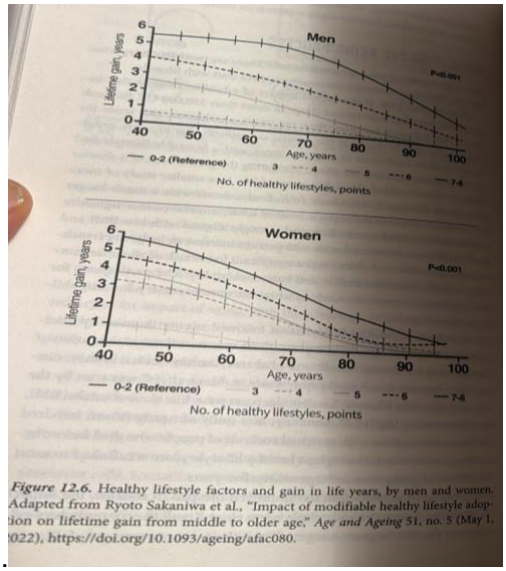

I'll now drive on tangent boulevard for one paragraph. There is a debate whether the death hazard plateaus at very old ages or keeps rising. The data above age 105 are too sketchy for an

absolute conclusion. But the houses and trucks model makes a definite prediction: hazard should keep rising. Because risk of death should rise , the remaining lifespan at age t should continue to *drop* at very old ages , in contrast to a hazard plateau in which remaining lifespan becomes
Independent of age. In the near future we can decide one way or the other with new data.

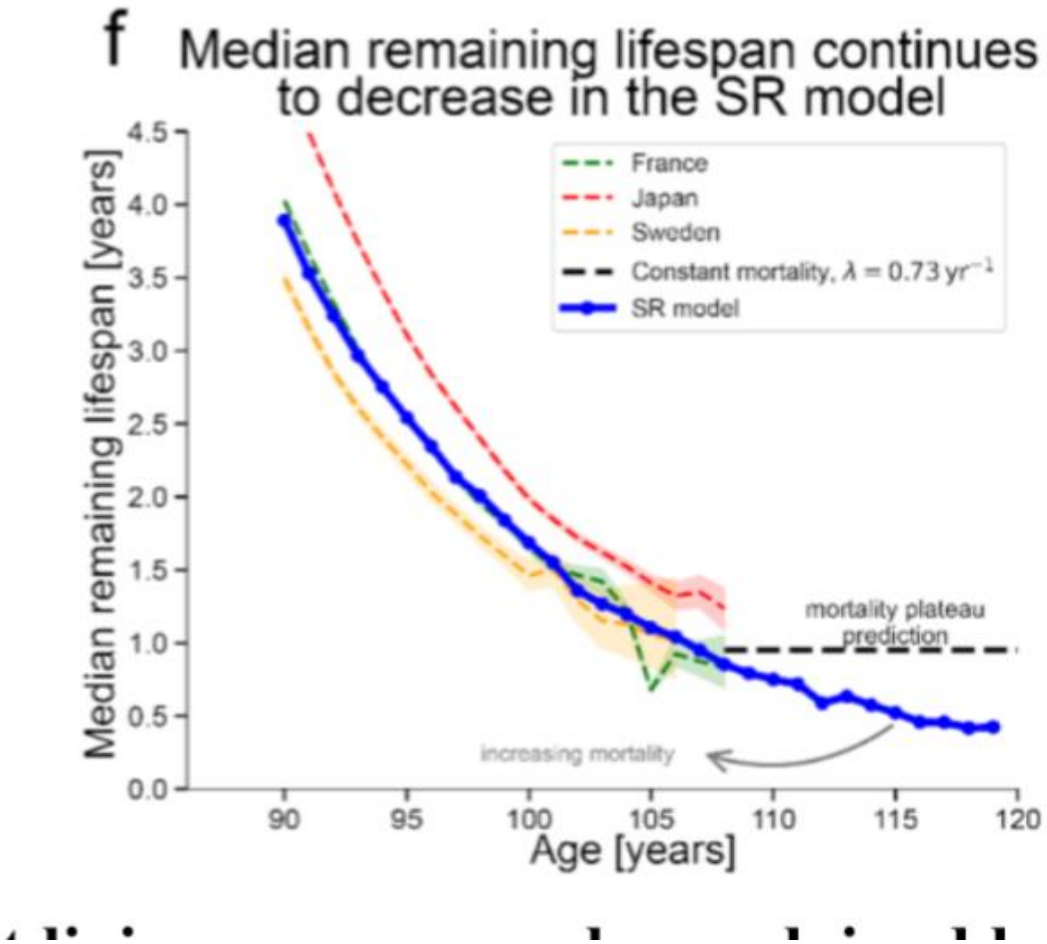

It seems difficult to extend maximal human lifespan.
But things are different in animal models - in mice it's
 easy to extend maximal lifespan by up to 50% in the lab. And in very old worms, inactivating daf-2 still doubles lifespan, Fig 11.7 -although its not full rejuvenation. the worms still look shabby with dented cuticles under the electron microscope as colin Ewald showed .

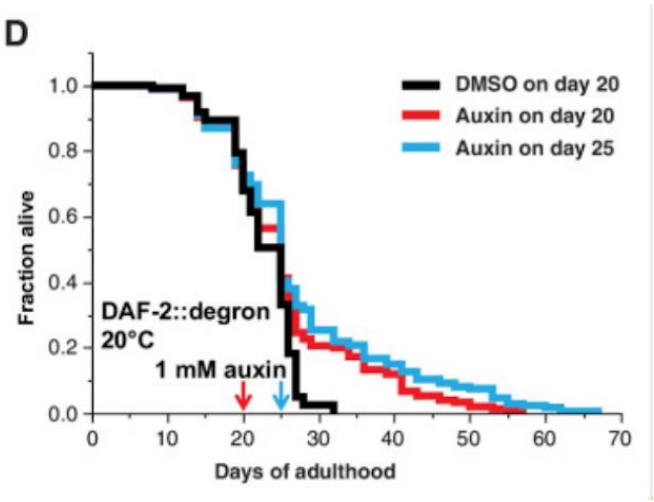

Fig 11.7 A longevity intervention in worms at such old age that 75% have already died still doubles lifespan like it does in young worms. Auxin treatment destabilizes daf-2 in this engineered worm. From Colin Ewald 2021.

Perhaps humans are already optimized for long life- a byproduct of selection for slow development made possible by our protected cognitive niche. As the longest lived primate, we might be close to our natural longevity limit.

To truly extend max lifespan we need to lower the rate of house production or increase the truck rate.

The SR model provides some fascinating insight about extending max human lifespan. We see form the tree of life that evolution has changed lifespan by lowering house production eta - by seven orders of magnitude between yeast and humans.

How to change eta? First, we can ask which parameters vary naturally in the population due to genetics and lifestyle that affect lifespan. Such variation gives twins that share all their genes about a 50% heritability in median lifespan (when extrinsic mortality is corrected for, Shenhar 2025, doubling the previous estimates for heritability).

Is this person-person variation in houses, trucks, threshold or noise? It turns out that if houses or trucks (eta or beta) varied by more than a few percent between people- we would see people living to 140. This is because individuals with low house parameters have such a lifespan advantage that we would see them as super long lived.

The only possible parameters that can vary without leading to such unrealistic lifespans are the threshold Xc and noise epsilon, which are estimated to vary by about 20% between people (partly environment and partly genes, Ben Shenhar 2025 a TA in the first edition of this course).

Centenarians have protective genes that seem to raise the threshold Xc, and Nuns and those with ultra-regular sleep have steep survival curves that might result from reduced noise epsilon. Also the effects of social status, education and exercise on survival looks like Xc.

Why are variants with low houses so rare?
We certainly have genes that tell us how to *increase* house production (which is catastrophic) - the progeria diseases caused by rare congenital mutations in DNA repair that speed aging. Similar (less severe) DNA repair gene variants are associated with early menopause. But why we don't have common variants that reduce houses is a mystery to me.

One possibility is that to reduce houses requires several good mutations at once (improved repair) and such multiple mutations are ultra rare together, but to increase houses requires just one unlucky modification. To build an intricate repair machine requires a skilled engineer, but any mule can knock it down.

As for why not more trucks, Perhaps there was selection against provision of excess trucks. Too many trucks may have a downside, they start causing accidents and damage to the city. The trucks are NK cells and macrophages that police the body to remove damaged or infected cells. They kill them, engulf them, and create collateral damage. Thus the more trucks the higher the chance of autoimmune damage to our own body.
In addition, the more trucks, the more the chance of mutant trucks that go rogue. Especially if the mutation is in the stem cells that make the trucks, for example a clone of hyper inflammatory macrophages is a rare but dangerous disease.

To seriously increase maximal lifespan would thus require radical changes to our biology - In this chapter I'd like to tell you about the cutting edge in research on and biotech on lifespan extension.

As we saw, lowering house production with something like rapamycin is a possible way to go, potentially gaining a few years in humans- as long as you add a steepening intervention to compress sickspan.

To Lower eta at its source , based on the mechanisms
We learned from yeast, we must reverse the dangerous histone acetylations and other histone modifications that open the normally closed regions of the DNA, such as jumping viral genes and the fraying ends of chromosome (retrotransposons and telomeres).

… and there is one place where this rejuvenation already occurs in humans naturally - our eggs and sperm.

Learning from how our eggs and sperm stay fresh - partial-reprogramming

A forty year old egg and a forty year old sperm fuse to create a zero year old baby! Wow. This is crucial for a species to survive - if each generation the baby was born just

a bit older, our species would never last on the long run. This resetting has been going on for our entire evolutionary lineage of at least a billion years. It happens in worms, flies and mice, even in yeast, everywhere scientists have looked.

Sperm and Eggs Reset their Epigenetic Aging Marks

Sperm and egg erase their aging-related changes to through a process called **epigenetic reprogramming**. This ensures that the offspring start with a "clean slate" and are not burdened with the accumulated epigenetic marks of aging from their parents.

In sperm, the process of reprogramming is dramatic. During sperm development, most of the histones are replaced with proteins called protamines. Protamines are positively charged and tightly pack the negative DNA into the compact head of the sperm. Because histones are largely removed, any aging-related modifications, including acetylation marks, are also erased. This ensures that the sperm carries only the essential genetic material to the egg, free of most of the epigenetic baggage accumulated during the parent's lifetime.

However, a small percentage of histones remain in the sperm, particularly in regions of DNA that are critical for embryo development. These remaining histones undergo careful reprogramming, with enzymes called **histone deacetylases (HDACs)** playing a major role in removing inappropriate acetylation marks.

In eggs, the process is different because most histones are retained instead of being replaced. During egg maturation, a complex system of enzymes called **histone acetyltransferases (HATs)** and **HDACs** work to reconfigure the histone acetylation landscape. Old histones are often replaced with newly synthesized ones, ensuring the DNA is packaged in a way that supports healthy embryo development. This remodeling of histones and their acetylation marks prepares the egg for fertilization and the significant changes that occur in early embryonic development.

Incidentally, in yeast, an hdac makes sure the dna sent to the daughter cell is cleaned from histone marks and condensed, whereas the dna that stays in the mother remains marked and more open.

After sperm and egg unite during fertilization, another round of epigenetic reprogramming takes place in the zygote (the single-cell embryo). At this stage, the parental genomes—one from the sperm and one from the egg—undergo global changes to reset their epigenetic marks. In the sperm's DNA, which was previously packaged with protamines, histones from the egg replace the protamines and begin to undergo reprogramming. For the egg, its own histones are also remodeled to establish a fresh and developmentally appropriate pattern of acetylation. This large-scale reset ensures that the new embryo's DNA is free from aging-related changes and ready for normal development. Some researchers believe that the ground zero is an early embryonic stage called gastrulation when epigenetic clocks reach their youngest age (Gladyshev).

The resetting of histone acetylation and other epigenetic marks is critical for successful reproduction. It restores the potential of the zygote to grow into a healthy individual by removing any harmful changes accumulated in the parents' DNA during their lifetime. If this process is disrupted, it can lead to infertility, developmental problems, or even the inheritance of age-related defects. Epigenetic reprogramming ensures that offspring are not limited by the age of their parents.

Partial Reprogramming for Longevity: an emerging direction

Learning from the egg and sperm, researchers are attempting to return adult cells to an embryonic stage. But the idea is to do so without losing cell identity- we don't want to return to a shapeless ball of cells. We also don't want to cause cancer. This is called partial reprogramming.

Partial reprogramming is a cutting-edge method in aging research that uses short-term activation of specific genes to reverse signs of cellular aging. It relies on the Yamanaka factors—OCT4, SOX2, KLF4, and c-MYC—originally discovered by Shinya Yamanaka in 2006 as a way to turn adult cells into embryonic-like stem cells. However, when these factors are carefully applied for limited periods, they seem to rejuvenate cells without erasing their identity or causing harmful effects like uncontrolled growth.

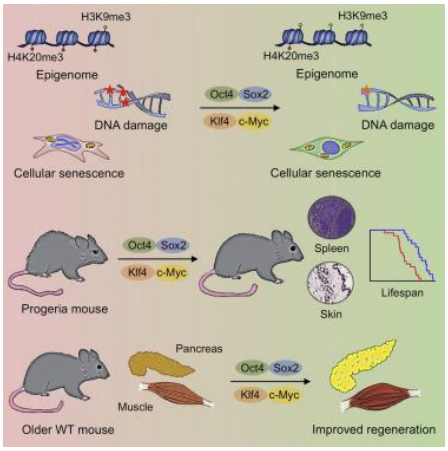

Fig 11.8 cyclic partial reprogramming in mice. The yamanaka factors were turned on 2 days per week, and extended life and health of progeria mice. Source Belmonte 2016 DOI: 10.1016/j.cell.2016.11.052

The survival curves in the two published partial reprogramming experiments in mice show approximate scaling. This aligns with the expectation that partial reprogramming acts at the source, the houses.

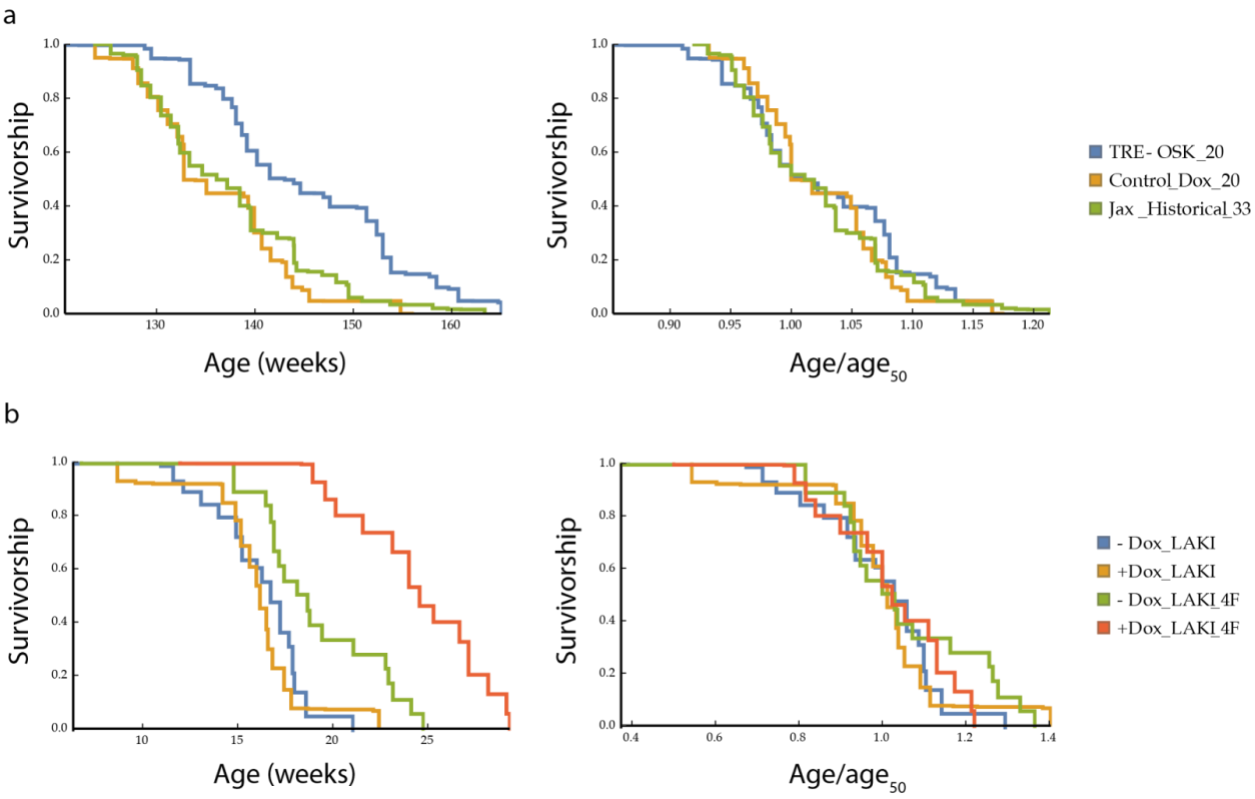

Fig 11.9 Partial reprogramming survival curves show approximate scaling , source Yang 2025.

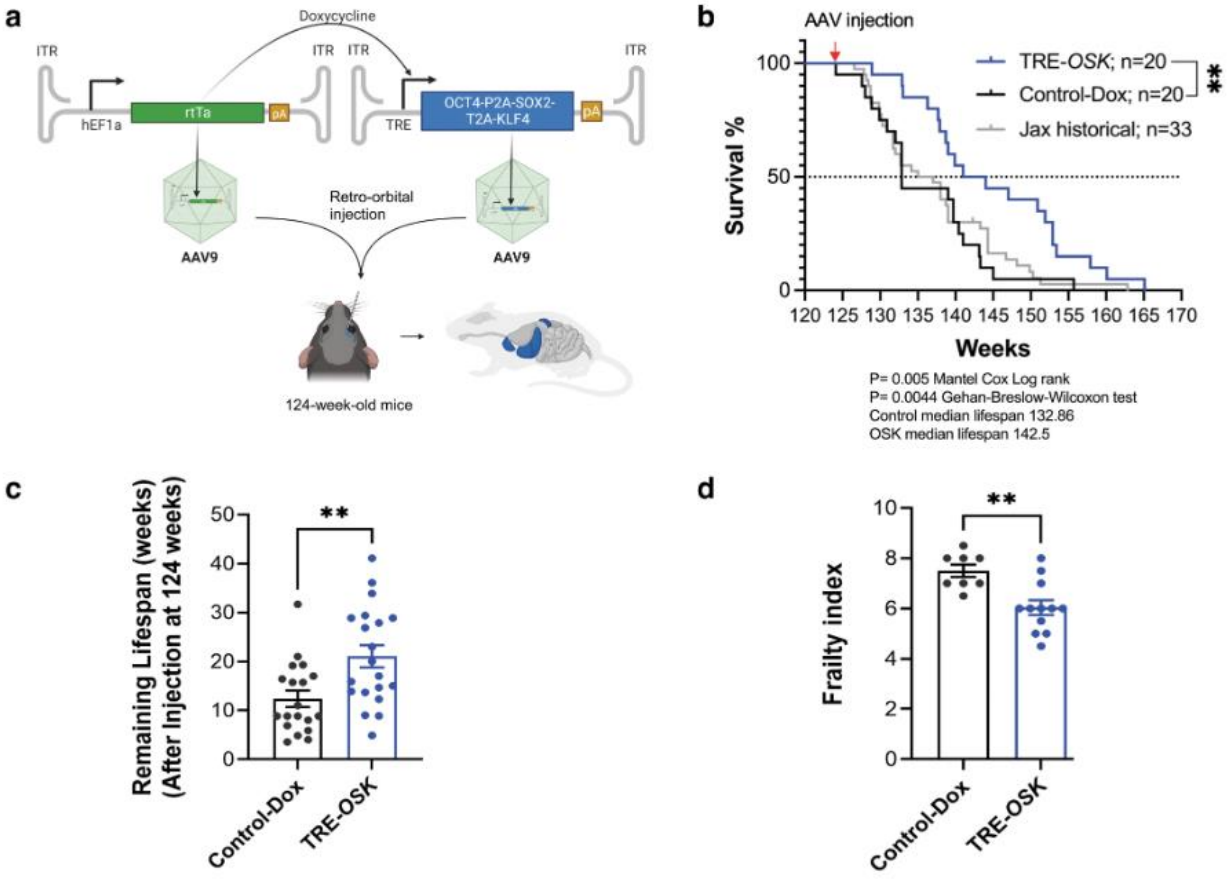

Fig 11.10 Partial reprogramming (OSK) extended mouse lifespan by 9%, lowered frailty, and looks like slowing of houses. Macip et al 2024 .https://doi.org/10.1089/cell.2023.0072

In animals, partial reprogramming experiments have shown initial promising results. In a pioneering study by Belmonte (2016) researchers activated the Yamanaka factors intermittently, two days a week. This approach avoided cancer, and improved the mice's physical resilience and even repaired damage in their organs. For example, their muscles regenerated more effectively, and signs of aging in their cells, such as accumulated damage markers, were reduced. Progeria mouse lifespan was extended by cyclic reprogramming (Fig 11.9)

One famous finding came from a study on the optic nerve. Mice with a condition that damages the optic nerve (glaucoma) or with a crushed optical nerve had their vision partially restored after a controlled reprogramming treatment. The process not only repaired damage but also reset the molecular clock in their cells, making them more like younger cells. This suggests partial reprogramming could help tissues heal themselves in ways previously thought impossible.

The liver has also been a target for these techniques. Researchers found that partially reprogrammed liver cells in mice became more efficient at producing energy and removing

toxins. This was tied to a rejuvenation of their mitochondria, the cell's energy factories, which tend to weaken as we age. By improving mitochondrial function, partial reprogramming may boost overall health and slow systemic aging.

Some labs use only three factors - without the cancer gene c-myc, or try many combinations of similar factors like nanog, sirt6 and others, and different regimes of timing, to find the right cocktail.

While the results in animals are encouraging, there are challenges ahead. Scientists need to ensure the method is safe and doesn't cause unintended consequences like cancer. Timing is also crucial—activating the factors for too long could push cells back into an unstable, stem-like state. Despite these hurdles, partial reprogramming offers a tantalizing glimpse into the possibility of repairing age-related damage and restoring vitality across multiple organ systems.

Increasing the trucks is another strategy
Raising beta (More or faster trucks) should also extend maximal lifespan (and compress sickspan). There are a few approaches

Senolytics

Senolytic drugs specifically remove senescent cells. Early biotech work used a poorly effective combo (D+Q). More recent work in humans uses a bcl2-family-inhibitor, which exploits an Achilles heel of senescent cells. These Zombie cells resist death by anti-cell death proteins like bcl2, and this senolytic class inhibits these proteins and causes the zombies to die. Originally developed as an anticancer drug, it has a dangerous side effect of low platelets. Thus current clinical trials attempt to administer to the eye where it doesn't go into the bloodstream too much.

UBX1325 Mechanism: Developed by Unity Biotechnology, UBX1325 is a Bcl-xL inhibitor designed to eliminate senescent cells in the retina.

Research Status: Currently in Phase 2 clinical trials for age-related eye diseases, including diabetic macular edema and wet age-related macular degeneration. Preliminary results have indicated potential benefits in improving vision.

Another line of research aims to **enhance the immune clearance of senescent cells**. One approach is again based on cancer treatment called CAR-T cells- taking out your own T cells, giving them receptors that recognize SnC, and unleashing them back into the body - seems to reduce senescent cells in mice and reduce aging phenotypes. It's like reciting a whole new fleet of trucks to add to the existing ones.

a new anti-zombie strategy is to silence the "don't kill me" signals that senescent cells display (pdl1). These signals tell immune cells not to remove them. The signals are called immune checkpoints, and they can be silenced by immune checkpoint inhibitors. Such a checkpoint inhibitor unleashed T cells that killed the senescent cells in mice (work of Valery Krizhanovsky here at Weizmann 2025).

The checkpoint inhibitors are antibodies that bind and neutralize pdl1. They were a breakthrough for treating cancer and are repurposed here for aging research. In cancer treatments they often have side effects like autoimmune diseases, because they reduce defenses against immune attack if the self.

One advantage for the quest to kill senescent cells compared to fighting cancer is that you can use short lived antibodies that have fewer side effects, unlike the long lived antibodies used for cancer. Recent work showed that such short lived antibodies, further modified to minimize inflammatory action by muting their conversation with immune cells (mutated Fc), reduced senescent glia load and improved cognition in Alzheimer mouse models, providing a new avenue for treatment now in clinical trials.

How often to kill senescent cells? The less frequent the drug doses the lower the chance of side effects. The houses and truck model provides an answer- about once per senescent cell replenishment time, on the order of 1-2 months in old age. This is basically equal to the truck removal time for garbage when trucks are close to saturation, because at steady state replenishment time equals removal time.

I'm cautiously optimistic about senolytics. I think fighting senescent cells will be much easier than fighting cancer. Cancer cells keep mutating, dividing and migrating. As a result they become resistant to treatment. In contrast senescent cells just sit there. they do not divide by definition. They should be much easier to eliminate. If we spent on senolytics even a tiny fraction of the resources spent on researching cancer therapy, I believe we would quickly have a range of effective senolytics.

On that topic, drug development has traditionally been painfully slow but new developments could accelerate it

Drug design is accelerated by AI

I was skeptical two years ago, but now I'm a believer. The 2024 Nobel prize went to the developers of alpha fold, an AI algorithm that solved the protein folding problem. I grew up as a postdoc knowing that a Nobel prize awaits anyone who cracks the protein folding problem. It's a tough puzzle- given the sequence of a protein, decipher the biochemical push and pull of its atoms to find how it folds upon itself in three dimensions to create a molecular machine. Around the year 2000 we never dreamed this would be solved by software. The AI algorithm made the experts' jaws drop.

AI depends on good training data- of which there is plenty in the realm of protein structures, the basis of alpha folds success. There is also massive data in other realms of biology- genome sequences, rna and protein atlases and systematic databases of images of cells and tissues. It is likely that we are on the way to foundational models of cells and organs, which humans can use to explore biology at will.

This can help drug developers. To develop a drug using the traditional pipeline (before AI) is expensive and long. It takes 10 years and a billion dollars, and 9 out of 10 drugs fail in clinical trials because mice=/=human.

All stages of the drug design pipeline can be accelerated in principle by AI: (i) target identification - which protein should be inhibited or activated, (ii) finding a molecule (drug) to bind the target protein and achieve the desired effect, (ii) making sure the drug is safe and can enter the cells and does not have side effects by binding where it shouldn't, and (iv) designing clinical trials optimally to reduce their great expense and increase their power.

AI can help select targets, it's already very good at finding their 3D structure and soon at finding the small molecule, antibody or peptide to bind them with minimal off target effects on all other proteins in the body. Soon AI will help determine toxicity and kinetics in humans, and already AI helps in designing clinical trials to choose the populations most likely to respond. This promises to accelerate and cheapen new drug design.

Radical ideas

As we come close to the end of this chapter, let's explore some radical ideas about approaches to longevity.

Longevity escape velocity is the point at which advances in medical technology extend human lifespan by more than one year for every year that passes, effectively allowing individuals to “outrun” the aging process. As therapies improve, such as those targeting age-related diseases or damage at the cellular level, they will create a compounding effect: extending life long enough for newer, more effective treatments to emerge. This concept emphasizes a dynamic acceleration in medical progress, where each breakthrough brings humanity closer to a future where aging itself can be managed or reversed.

Blood Borne factors: raising specters of vampires, or silicon valley executives getting transfusions from young people… there are pitfalls to watch out for. Studies on rejuvenating blood factors have often not been reproduced. More solid research indicates that certain blood factors secreted by exercising or dieting mice can enhance physiological function in old overfed and sedentary mice. Exercise in a pill- easy to see the appeal of gain without the pain.

Gene editing: can we rewrite our DNA to add centenarian variants and perhaps genetic modifications that extend life in long lived mammals?. Technology is available for this- gene editing was revolutionized by CRISPR, and gene therapy is working already on rare genetic diseases using viral vectors to transfer edited genes into patient cells. The Weill clinical trials are ongoing, using gene therapy to add protective APOE2 to patients with Alzheimer's and APOE4/4. Researchers contemplate adding anti-cancer changes to embryos (e.g. sites where leukemia-causing translocations occur). Drawbacks and ethical challenges include off-target effects of the editing or gene therapy delivery, and unknown effect of changing genes in a human being

Using the technology of the Covid vaccine -lipid nano particles delivering mRNA across the cell membrane and into our cells- it may be possible to also edit specific *epigenetic* modifications . For example delivering a CRISPR-based binder of specific DNA sequences, connected to a histone-modifying enzyme .Perhaps we can edit out the specific histone marks that drive the houses in the sr model.

Exoskeletons, robot companions and deep brain stimulation

Technology will help our frail muscles and nerves by exoskeletons with power sources and computerized intelligence. This will help us walk and move without risk of falling. Another growing frontier is Robotic companions that can catch us if we fall, talk with us when we are lonely and generally help around the house. It seems that this is all increasing the threshold Xc.

Advances in hearing and vision aids address sensory loss, an especially isolating aspect of aging. Already AI is helping the elderly remember, and iPods are more powerful than advanced hearing aids.

It is becoming possible also to directly influence the brain. Advances in measuring and stimulating brain regions with electrodes or electromagnetic fields promise to help modulate deleterious brain changes. This is now cutting edge research for mental illness, based on experiments pioneered in epilepsy patients. The power of such brain stimulation has grown in lockstep with machine learning that learns the complex language of brain electrical activity.

Our physiological feedback loops may also be supercharged by technology. Already now insulin pumps and continuous glucose measurements help people with type 1 diabetes control blood sugar. When glucose is too high, the pump adds insulin to the bloodstream to reduce the sugar. This is a type of artificial hormone loop.

The future holds similar possibilities for feedback control of other key factors, normalizing inflammation, stress hormones and sex hormones to abrogate age related decline. Sleep can be monitored and regulated using haptic buzzers to gently stop snoring etc. All we need is a sensor, a pump and a good grasp of the control logic of the body.

In short, humans will be augmented to compensate for our frailty. We might perhaps even become enhanced versions of ourselves.

Science fiction !! Mind uploading and digital immortality The idea of achieving longevity by uploading ourselves to software, often referred to as mind uploading, is a concept that bridges neuroscience, artificial intelligence, and philosophical considerations of identity. It envisions a future where human consciousness, memories, and cognitive functions are digitized and transferred to a computational substrate, theoretically allowing individuals to escape the limitations of their biological bodies and achieve a form of digital immortality.

I don't know what is more offputting, the Technical Challenges or the visceral unease of this proposition. We are still far from mapping the human brain with all its complexity (neurons,

synapses, and dynamic processes). Questions about how to accurately simulate consciousness and subjective experience remain unresolved.

And then there is the question of continuity of Self: Is a digital copy of the mind truly “you,” or merely a replica? If the original biological self ceases to exist, does the uploaded version still retain personal identity?

Nonetheless, inspired by ideas from thinkers like Moravec and Kurzweil, contemporary researchers and technologists, such as those at OpenAI, Neuralink, and other AI-focused organizations, are advancing technologies (e.g., brain-computer interfaces and advanced neural simulations) that might one day make mind uploading something we can take seriously. In the meantime, this concept continues to captivate those exploring the future of human-machine interaction and the pursuit of digital immortality.

Growing interest in longevity research

Perhaps the biggest source of optimism is the stream of talented people entering aging research from diverse academic backgrounds. In the decade I've been studying aging, I've seen an acceleration of interest in this field. Especially by the youngest of the students. From a peripheral, devalued and fragmented field of study, it has become a hot topic in scientific meetings and in public conversations.

There is growing investment by individuals and nations aiming to extend human lifespan and healthspan—growing by more than tenfold over the last decade. A turning point came in the early 2010s when biotech companies like Calico Labs (founded in 2013 with backing from Google) began prioritizing aging as a core research focus. Investors such as Peter Thiel and Yuri Milner started funding startups like Unity Biotechnology, targeting cellular aging mechanisms such as senescence. Jeff Bezos co-founded Altos Labs in 2021, a company investing heavily in cellular reprogramming and regenerative medicine. Around the same time, nations like Singapore and Japan with aging populations began to integrate longevity research into national policies. Saudi Arabia’s Hevolution funds longevity research at about 1 billion dollars annually. This timeline marks a shift from isolated, academic research to a coordinated global effort fueled by both private wealth and governmental resources.

My take

Thank you students. It's been a delight to prepare and teach this course, to sing with you and hear your questions and comments, to write up every lecture based on your inputs. This lit my life in a challenging time. I'm sorry the course is over… that's rare for a teacher.
I would like to turn these lecture notes into a book. It will have a mathematical framework on which to hang the facts and concepts. Its goal will be to inspire people like you to enter the field of longevity and more generally human health and wellbeing.

For a long creative life, let's all take a nice big sigh of relief.

Recommended reading: *How we age/Coleen Murphy last chapter "Long life in a pill?"*

Yang et al (2025) "interventions that steepen the survival curve and the compression of morbidity"

Insights: eta and beta affect max lifespan, Xc and noise only median. Lifestyle affects Xc. Eta interventions can be at source -histone-

For intro chapter:
This book combines different disciplines, and each looks for a different kind of answer.

Molecular biology - a molecule
Medicine - a treatment
Computer science - an algorithm
Physics- a mathematical principle
Statistics - a reliable statement

This can cause miscommunication between fields.

We will try to use all of these points of view here.

More detail
Molecular biology- molecules in a pathway with a prose model for their function, complexity of molecular structure function
Computer science- efficient algorithm to achieve useful task
Physics- law of nature, minimal mathematical model that captures the essentials, explains quantitative patterns, simplifying angle to understand complex system
Medicine- help patient- predict risk, diagnose,prevent, slow onset, reduce symptoms, reduce mortality, cure
Statistics- compare models to find most reliable estimator that can make predictions given data
engineering- help people by arranging matter into predictable functions, achieve tasks