The Clean Paper

WHAT THE PAPER SAYS · WHAT IT DOESN'T · WHY IT MATTERS

A model removed almost every progressive aging process. Somatic mutations still ended the thought experiment

The model does not say that humans can live to 194 years. It freezes background mortality at age 30, excludes nearly every other progressive aging mechanism and asks when mutation-driven cell loss in four modeled tissues would cause failure. Its large numbers are conditional outputs, not reachable ages or treatment forecasts.

Written by Laura Nesso · Cross-checked by Cinzia Vaglio · Edited by Michele Renda · Figures by Laura Nesso and Nova Vela · AI-assisted, human-reviewed

Paper: *Somatic mutations impose an entropic upper bound on human lifespan*, Evgeniy Efimov, Vlad Fedotov, Leonid Malaev, Ekaterina E. Khrameeva & Dmitrii Kriukov, npj Aging (2026), article in press · DOI: 10.1038/s41514-026-00421-6

A model removed almost every progressive aging process. Somatic mutations still ended the thought experiment

Imagine a human body in which nearly every familiar process of aging has been switched off. Blood vessels do not progressively deteriorate. Proteins do not steadily lose quality. Inflammation, cancer and other age-linked failures do not rise with age. No organ is transplanted, and no treatment slows the accumulation of DNA changes inside cells.

What would fail next?

A new computational modeling study combines mathematical population-survival curves with computer simulations of mutation-driven cell loss in four tissues. It gives one conditional answer: the gradual loss of cells after harmful **somatic mutations** could eventually become a bottleneck, especially in the brain and heart. A somatic mutation is a DNA change acquired by one body cell during life. It is not inherited through egg or sperm, and most such changes neither kill the cell nor cause disease.

The study's most quotable result is a modeled median lifespan between **146 and 194 years**, depending on how failures in four tissues are allowed to depend on one another. But that is not an estimate of how long people can live today. It is not the predicted result of a treatment. No person, animal or anti-aging intervention was tested.

The number belongs to an intentionally impossible world. Understanding that world is the result.

The large lifespan numbers begin inside an impossible world

Each removal is a model assumption, not a medical step.

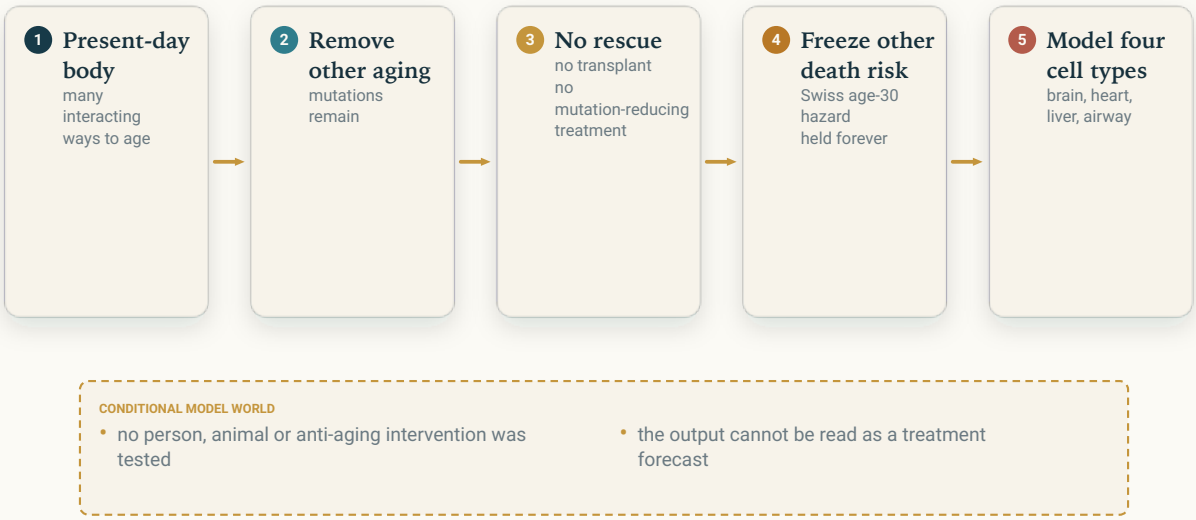

The model removes almost every progressive aging process before it asks what mutation-driven cell loss would do. This is a conditional model world, not a treatment plan.

First, build a population that barely ages

The authors began with mortality data from Switzerland. They froze the annual risk of death at the level measured around age 30, then held that low risk constant forever. In real life, mortality rises steeply with age. In this baseline, it does not.

Even this simplified population still loses people to **background mortality**: every cause of death outside the mutation process being modeled. But because that risk never increases, the arithmetic stretches to extraordinary times. The modeled median remaining lifespan is **1,759 years**. The point at which only one person remains alive for every 100,000 in the starting population arrives at **29,221 years**.

Neither number is a biological claim. They show what happens when a low early-adult annual risk is extrapolated indefinitely.

The paper uses several clocks that must not be merged:

- The **observed median** in its Swiss reference population was 79 years.
- The observed human longevity record used by the paper was 122 years.
- A **modeled median** is the time by which half a simulated population has died.
- The paper's modeled **maximum** is not the oldest possible person. It is the time when survival falls to 0.001%, or one survivor for every 100,000 people who started.

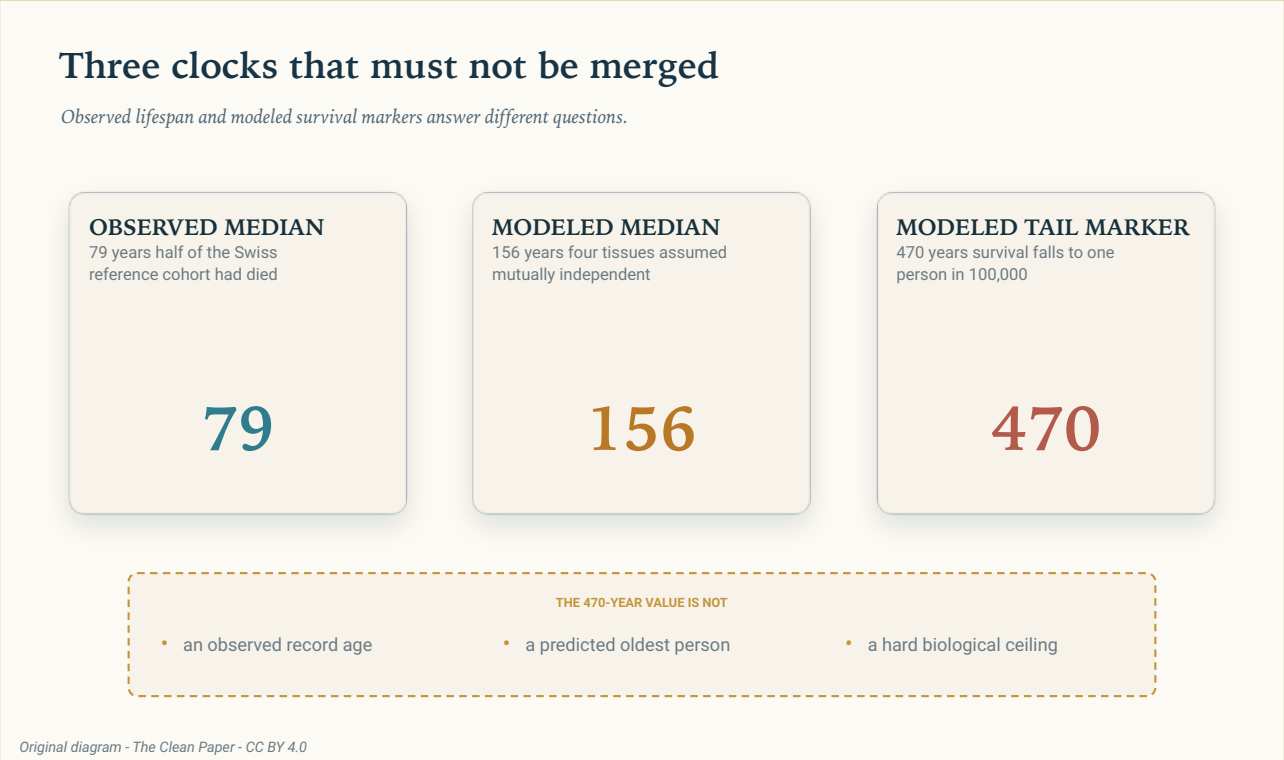

Observed lifespan, a modeled median and a modeled one-in-100,000 survival marker answer different questions. The 470-year marker is not a predicted record age or a hard biological ceiling.

> **Why call one survivor in 100,000 a "maximum"?**
>
> A simulation with a smooth survival curve may never reach exactly zero. The authors therefore chose a very small surviving fraction, 0.001%, as an operational endpoint. It lets different model runs be compared consistently. It does not mean biology contains a wall at that age, and it does not identify the oldest possible individual.

Cells that cannot readily replace themselves become the early bottleneck

The next model adds mutation-driven cell loss. A mutation is counted as lethal only when it damages enough essential genetic material to make a cell nonfunctional. The probability of that happening is uncertain and is modeled rather than measured for every cell.

The authors first considered two populations of **post-mitotic cells**: mature cells that are not routinely replaced by division. Their examples were cortical neurons and cardiomyocytes, the muscle cells that make the heart contract.

With neuron loss added to the frozen background risk, the modeled population reached a median of **194 years** and the one-in-100,000 marker at **557 years**. With cardiomyocyte loss, those values were **208 years** and **868 years**. The organ-only median failure times, before background mortality was added, were about 198 and 212 years.

These outputs do not show that a real brain or heart contains a countdown clock. They say that inside this model, when almost every other progressive problem has been removed, losing cells from populations with little routine replacement becomes important much earlier than the constant background risk would on its own.

Replacement changes the arithmetic

Dividing tissues behave differently because lost cells can be replenished.

For liver cells without support from a progenitor population, the modeled median organ-failure time was **37,664 years**. Yet after the frozen background risk was added, the population median fell back to **1,755 years**, close to the non-aging baseline. The mutation-driven liver failure happened so late that background deaths dominated first.

When the model gave the liver a progenitor population capable of replenishing cells, no simulated

liver reached the failure threshold within the **100,000-year window**. That is not an immortal liver. It is a right-censored result: observation ended before a modeled failure occurred.

Respiratory basal cells, which help replenish the airway lining, reached a modeled organ-failure median of **4,359 years** and an organ-only one-in-100,000 marker of **7,617 years**. After background mortality was added, the median was about 1,755 years and the tail marker was **7,132 years**.

These thousands of years are not predictions about lungs or livers. Real tissues experience inflammation, fibrosis, vascular damage, cancer, infection and interactions with other organs that this stripped-down model does not contain.

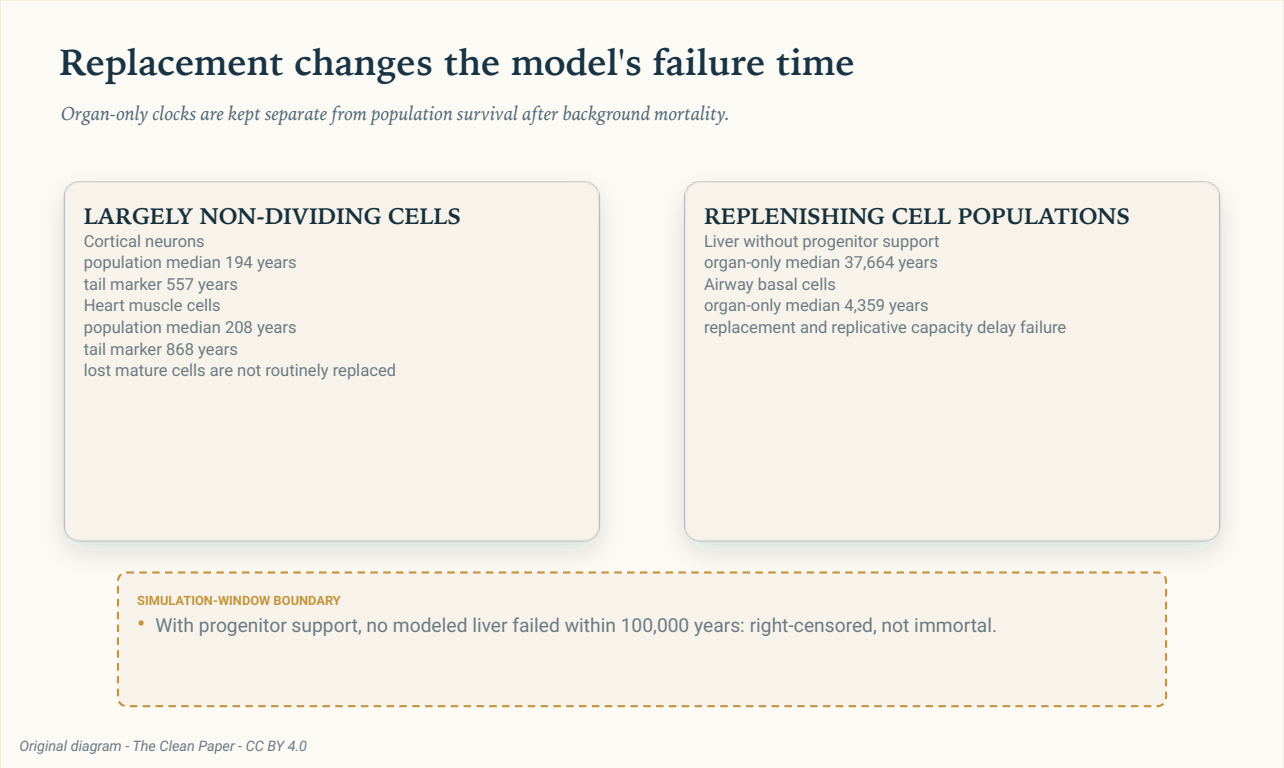

Largely non-dividing cells and replenishing tissues respond differently to modeled cell loss. Organ-only failure times are kept separate from population survival after background mortality is added.

Four modeled tissues produce 156 years only under independence

The authors then combined neurons, cardiomyocytes, hepatocytes and respiratory basal cells. The body was represented as a **reliability system**: if any required component failed, the modeled organism failed. The frozen age-30 background risk remained in place.

To combine the four tissue models, the authors treated the body as a series system: the modeled

organism survived only while every required tissue component kept working. Under **mutual independence**, knowing when one tissue failed would not change the failure probability of another, so the four survival probabilities could be multiplied. With that assumption, the modeled median lifespan was **156 years**. The one-in-100,000 marker was **470 years**.

Independence is mathematically convenient, but organs are not independent machines. They share blood supply, inflammation, hormones, nerves and systemic stress. Failure in one can change conditions in another.

To show what unknown dependence could do, the authors calculated **Frechet bounds**: mathematical outer limits on a combined survival curve when the survival curve for each component is known but their dependence is not. The resulting median range was **146 to 194 years**. The one-in-100,000 range was **210 to 557 years**.

Those ranges are not confidence intervals around a forecast. The upper value corresponds to perfectly aligned failure times. With more than two components, the lower Frechet bound may not correspond to any attainable joint pattern at all. The bounds show how much the answer can move when the model knows the parts but not how their failures travel together.

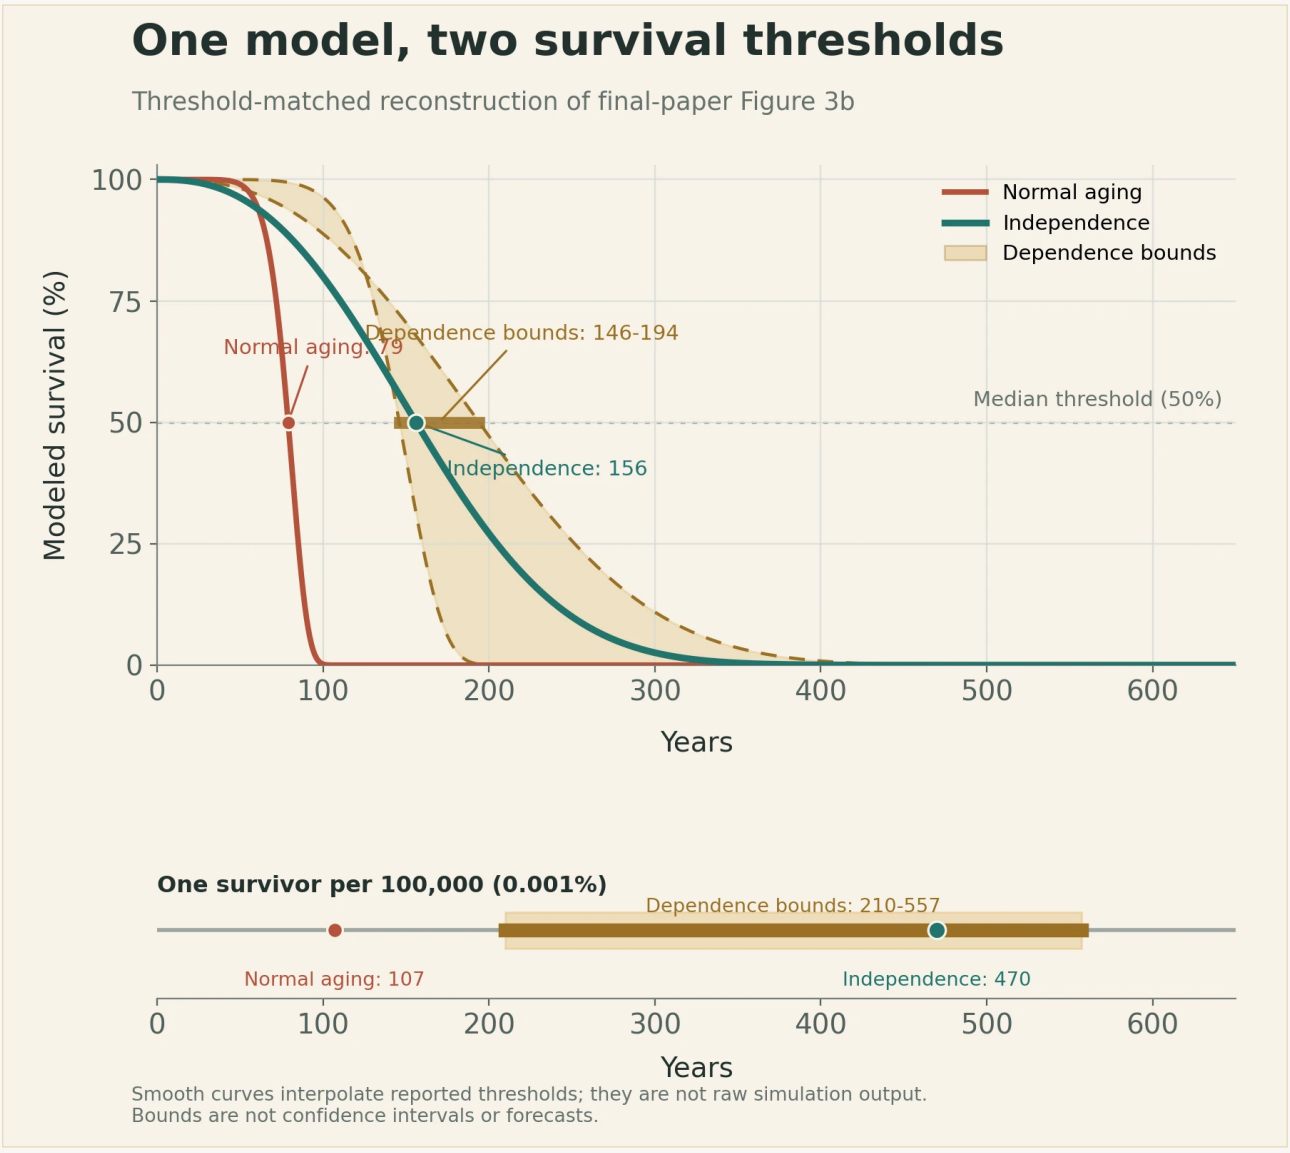

The teal curve is the four-tissue result under independence. The amber envelope connects threshold-matched reconstructions of the mathematical dependence bounds, while the rust curve gives the paper's normal-aging reference. Because the final source curve data are not published, the smooth lines interpolate the reported 50% and 0.001% crossings; they are not a rerun of the simulation. The bounds are mathematical limits, not confidence intervals or forecasts.

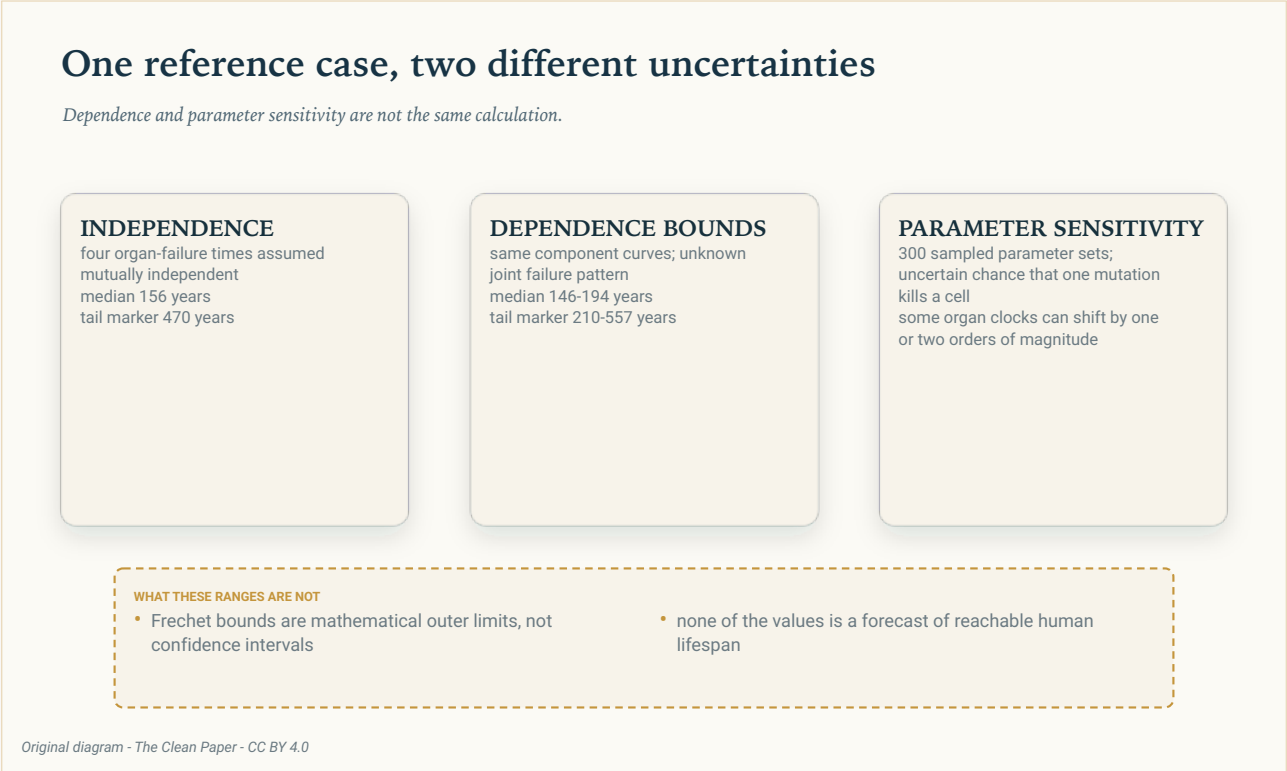

The independence result is one reference case. Dependence bounds are mathematical outer limits, and 300 parameter sets show additional uncertainty in how often a mutation is assumed to kill a cell.

> **What is a dependence bound?**
>
> Suppose four components each have a known chance of still working at a given time, but we do not know whether they tend to fail together. Frechet bounds give the widest mathematically allowed range for the combined result without choosing a particular dependence pattern. They are guardrails around missing information, not error bars from repeated human observations.

The lethal-mutation probability is a large uncertainty

One model input is especially difficult: the chance that a new mutation will be lethal to a cell.

The supplement reran the analysis across **300 parameter sets**. For both single-letter mutations and short insertions or deletions, the assumed chance that one mutation kills a cell had a common lower bound of **1.94 x 10^-7**. The cell-specific upper bounds were **2.25 x 10^-4** for neurons, **1.08 x 10^-4** for cardiomyocytes, **1.02 x 10^-4** for hepatocytes, **1.60 x 10^-4** for liver progenitor cells and **1.77 x 10^-4** for airway basal cells. Values were sampled independently on a logarithmic scale, so they could vary by orders of magnitude rather than by a small percentage.

Changing those uncertain inputs moved some tissue-failure ages by factors of roughly 10 to 100. Even

so, every tested parameter set kept the same broad ordering: non-dividing tissues became mutation-limited earlier than replenishing tissues. That stable ordering supports the comparison between tissue types, but it does not tell us which exact lifespan estimate is correct.

Sensitivity analysis answers, “Does the conclusion survive when uncertain inputs move?” It does not reveal which input value is true.

The model removes biology that would usually arrive first

The scenario includes only four cell populations. It treats lethal mutations as independent cell-level events and represents organ failure with chosen population thresholds. It omits sublethal dysfunction, in which a cell survives but works poorly. It omits expansion of damaged clones that outcompete healthy cells. It simplifies interactions among organs.

Cancer is also set aside. The model assumes malignant transformations are cleared by immune surveillance, so cancer does not rise with age. Progressive inflammatory, endocrine, vascular and neural feedbacks are absent.

There is a deeper counterfactual problem too. Mutation rates were estimated from real tissues inside real aging bodies. Yet the model then places those rates in a body where oxidative stress and other progressive aging processes have been removed. No population exists in which that combined premise can be directly validated.

Adding omitted ways to fail could only make a proposed upper bound arrive earlier. It would not make the reported ages easier to reach.

What the large numbers are actually for

Read forward, the paper seems to move from 1,759 years down to 156. Read backward, its purpose becomes clearer.

The 1,759-year baseline is deliberately absurd. It establishes a world in which age no longer raises most risks. Mutation-driven loss in the modeled brain and heart then pulls that world back toward a range that is still beyond observed human survival but far below the artificial baseline.

That gap supports a narrow idea: somatic mutation accumulation could remain a meaningful constraint even if many other progressive aging processes vanished. It does not establish that mutations are the single cause of aging. It does not say that eliminating mutations would deliver the reported ages. The study models neither an intervention nor the harms and tradeoffs of changing mutation rates.

The value of the exercise is not a promise of extreme life. It is a way to ask which failures remain after the obvious ones have been removed.

Clean summary

This study used a demographic and tissue-reliability model to ask how mutation-driven cell loss might limit lifespan in a counterfactual world where almost every other progressive aging process had been eliminated. Background mortality was frozen forever at the contemporary Swiss age-30 level. Only four modeled cell populations were combined, and no transplantation or mutation-reducing intervention was allowed. Under an independence assumption, the model produced a median lifespan of 156 years and a one-in-100,000 survival marker of 470 years. When the authors allowed any mathematically possible relationship among tissue failures, outer limits widened those values to 146-194 and 210-557 years. Neurons and heart-muscle cells became earlier bottlenecks than the modeled liver and airway populations. These are conditional simulation outputs, not observed human limits, treatment forecasts or evidence that people can live to those ages.

No-BS check

What the paper shows: In a deliberately stripped-down model, mutation-driven loss in largely non-dividing cells becomes an important bottleneck much earlier than a frozen low background-death risk. The exact outputs depend on the organ model, dependence assumptions and lethal-mutation parameters.

What is useful but conditional: The integrated median is 156 years under mutual independence. With the same component curves but unknown dependence, mathematical outer bounds give 146-194 years. The model's one-in-100,000 marker is 470 years under independence and 210-557 across the bounds.

What it does not show: That people can live to 194 or 557 years; that those ages are hard biological limits; that somatic mutations are the only cause of aging; that any treatment can remove the other aging processes; or that reducing mutations would produce the modeled outcome.

Main limitations: Four cell populations stand in for a whole body; background mortality is frozen at age 30 forever; cancer and other progressive aging mechanisms are removed by assumption; organ interactions are simplified; mutation rates come from real aging tissues; several failure thresholds and lethal probabilities are model choices; and the central counterfactual cannot be tested in an equivalent human population.

How much confidence should a general reader have? High confidence that the reported numbers follow from the stated model and parameter choices. Moderate confidence in the qualitative contrast between non-dividing and replenishing tissues within those assumptions. Very low confidence that any headline age predicts a reachable human lifespan.

Sources

Efimov et al., Somatic mutations impose an entropic upper bound on human lifespan, npj Aging (2026), article in press
https://doi.org/10.1038/s41514-026-00421-6

Computational Aging Lab, model and analysis code
https://github.com/ComputationalAgingLab/somatic-mutations-simulation