

Philip Clark | Senior Consulting Actuary

What happens to mortality improvement when science starts moving faster — and what happens to an annuity book if it moves all at once?

A recent *International Business Times* piece picked up on a striking claim from Anthropic CEO Dario Amodei: sufficiently capable AI could increase the rate of biological discovery tenfold, compressing 50 to 100 years of biological progress into five or ten (Amodei, 2024). Amodei presents it as a prediction, not a promise: biology still has to be tested, trials still take time, regulators still need to be satisfied.

But for those of us who think about longevity professionally, there is an interesting question underneath the headline: **could AI materially change the trajectory of mortality improvement?** There are at least four ways it could.

How AI may accelerate mortality improvements

1. Earlier detection Finding disease before it becomes disease — AI reading imaging, pathology, genomics and wearables.	2. Faster medical innovation AI compresses the drug-discovery search problem, bringing treatments to market sooner.
3. Genetic tailoring Shifting from population medicine to individual medicine via genomics and pharmacogenomics.	4. Treating ageing itself The most disruptive possibility: therapies that target the biology of ageing directly, not just its diseases.

1. Earlier detection: finding disease before it becomes disease

The most immediate opportunity is diagnosis. AI finds patterns across imaging, pathology, genomics and wearables that humans may miss — and the earlier a disease is caught, the more opportunity there is to intervene.

AI is already used to analyse large, complex datasets and identify disease risk (Sharma et al., 2024). Earlier detection doesn't require revolutionary treatments, it's medicine getting better at applying the tools it already has.

2. Faster medical innovation

A more interesting possibility is that AI increases the rate at which new treatments are discovered. Drug development is largely a search problem, and AI can search much of that space computationally before candidates reach the lab.

A 2024 review in *Med* highlighted major AI advances across small molecules, RNA, antibodies and clinical trials, while stressing that AI-designed drugs still require conventional validation (Zhang et al., 2024). AI doesn't remove biology from the equation — it may simply make the search faster.

A treatment that might otherwise arrive in 2045 arriving in 2035 brings its mortality improvement ten years earlier too.

3. Genetic tailoring: from population medicine to individual medicine

AI and genomics could also shift medicine from what works best for people with a disease towards what works best for this particular person — interpreting genomic variants and predicting treatment response (Sharma et al., 2024).

Pharmacogenomics is a good example: genetic variants affect drug efficacy and adverse-reaction risk, so better identification of who will and won't respond should improve outcomes (Chenchula et al., 2024). None of this needs a revolutionary treatment, better matching of existing treatments to individuals could itself reduce mortality.

4. The really interesting possibility: treating ageing itself

This is where things stop being a story about better medicine and start being a story about a different world. Treating diseases of ageing is familiar; treating ageing itself is something quite different... and it is no longer purely hypothetical.

A recent paper by Guy Coughlan and Richard Faragher, *Extreme Longevity: A science-based scenario analysis* (Coughlan & Faragher, 2026), sets out a deliberately concrete version of this scenario, and it's worth walking through because of how unspeculative the underlying biology is.

The scenario centres on a repurposed drug combination: broad-spectrum senolytics, which selectively clear senescent cells, paired with low-dose rapalogues (derived from rapamycin), targeting dysregulated nutrient sensing. Both sit within the “hallmarks of ageing” framework (López-Otín et al., 2023) and have already extended lifespan and healthspan in animal models, separately and combined. Coughlan and Faragher model the treatment using the FDA's TAME trial design — built to treat ageing itself as the clinical target — and note that faster regulatory pathways already exist, including the MHRA's 2026 reforms and the FDA's real-time trial initiative (Coughlan & Faragher, 2026).

500+

AI-driven discovery platforms are already re-mining decades of data to resurface anti-ageing candidates previously missed — with one system alone flagging over 500 such candidates (Ying et al., 2025, as cited in Coughlan & Faragher, 2026).

The longevity actuary's problem

As longevity actuaries, we know many past mortality improvements have been driven by medical progress. In the SELECT trial, semaglutide reduced the primary composite cardiovascular endpoint by 20%, with all-cause mortality lower in the treatment group (hazard ratio 0.81) (Lincoff et al., 2023).

GLP-1s may become an important contributor to future mortality improvement, but I wouldn't call them a fundamental disruption — more another step in the existing story. **They probably help sustain the existing direction of travel rather than bend it.**

Smoking cessation is a useful analogy: it contributed enormously to falling cardiovascular mortality, but accumulated over decades rather than arriving as one breakthrough. That's generally how longevity works.

AI could be different, because it potentially changes not just what we discover, but how quickly we discover it.

So, is AI different from previous advances?

Possibly. AI could simply produce another generation of conventional improvements: better diagnosis, faster discovery, more personalised treatment... all part of the familiar process. The more disruptive possibility is the one Coughlan and Faragher discuss.

Translate the life-expectancy shift into liabilities and the picture sharpens fast: the cost of providing a typical annuity would rise by 33% for a 65-year-old cohort, and by 30–45% across a typical annuity book once take-up and adherence are factored in (Coughlan & Faragher, 2026).

Now compare that to the capital insurers actually hold: LAGIC stresses mortality by a permanent, instantaneous 20% reduction. A 30–45% valuation shock walks straight through that buffer, and through the PCA behind it.

The regulatory stress test annuity writers are held to is nowhere near the size of one scenario a serious paper now calls plausible rather than theoretical (Coughlan & Faragher, 2026).

The Carrington Event analogy

I'd go back to an analogy I've used before: the Carrington Event of 1859, the geomagnetic storm that set telegraph systems on fire (NASA, 2020). Nobody built the 1850s telegraph network to survive that storm, and nobody builds infrastructures today that would survive another of such events. A genuine cure for ageing presents a similar problem for lifetime income. Why worry about managing a risk that will collapse a whole industry anyway?

From the longevity actuary's perspective

There are two mistakes we should avoid.

The first is treating every exciting medical development as a longevity revolution.

GLP-1 drugs are a good example. They may well have a material impact on mortality and morbidity. But that does not necessarily mean they represent a break in the trajectory of human longevity.

The second mistake is the opposite one: assuming that the next advance must be as incremental as the last.

If we assume that mortality improvement will always look like what we have observed historically, then we risk becoming very comfortable applying the same stress framework year after year, without asking whether the underlying risk has fundamentally changed.

So, what should life insurers do?

1. Move towards a genuinely prospective approach to mortality improvement.

Mortality assumptions should not simply extrapolate the average. They should incorporate a forward-looking assessment of emerging drivers of mortality improvement and the uncertainty around them.

2. Monitor the longevity landscape systematically.

Life insurers need to know what is happening in areas such as obesity treatments, gene therapies, cancer treatments, personalised medicine, regenerative medicine and, ultimately, interventions that could alter the biology of ageing itself.

The objective is to identify when the probability of a major breakthrough is increasing, early enough to do something about it.

3. Use the traditional tools of risk management.

Scenario analysis. Stress testing. Capital buffers. Reinsurance. Diversification. Limits.

We don't need to invent an entirely new risk-management framework for longevity risk.

4. And perhaps we should start thinking differently about the contract itself.

This is the most provocative possibility.

Insurance has always recognised that behaviour can change risk. If a policyholder takes up an activity that materially increases their risk, the insurer may charge more or exclude the risk.

So why should longevity insurance be completely different?

Could an annuity contract contain a provision allowing the insurer to adjust the economics of the contract in response, not by abandoning a **baseline income** that policyholders rely on, but by reflecting, fairly and transparently, a change in risk that simply didn't exist when the contract was originally priced?

After all, this would not be entirely without precedent. There is already at least one annuity product in the Australian market where the income guarantee is contingent, in small dose, on population mortality experience remaining within expectations.

There are also several products where the payout is linked to the performance of a market index.

AI-driven advances in medicine may provide the catalyst for a new approach to managing longevity risk, one in which annuitants retain some limited exposure to the risk of disruptive medical breakthroughs that have yet to be discovered but could materially extend human life expectancy.

References

Amodei, D. (2024). Machines of Loving Grace: How AI Could Transform the World. <https://darioamodei.com/essay/machines-of-loving-grace>

APRA (2026). Prudential Standard LPS 115 — Capital Adequacy: Insurance Risk Charge. <https://www.apra.gov.au/standards/lps-115>

Bank of England (2025). Life Insurance Stress Test: 2025 Results, 24 November 2025. <https://www.bankofengland.co.uk/prudential-regulation/publication/2025/november/list-2025-results-report>

Chenchula, S., Atal, S. & Uppugunduri, C.R.S. (2024). A review of real-world evidence on preemptive pharmacogenomic testing for preventing adverse drug reactions: a reality for future health care. *The Pharmacogenomics Journal*.

Coughlan, G. & Faragher, R. (2026). Extreme Longevity: A science-based scenario analysis. 20 July 2026.

Lincoff, A.M. et al. (2023). Semaglutide and Cardiovascular Outcomes in Obesity without Diabetes. *New England Journal of Medicine*, 389, 2221–2232.

López-Otín, C. et al. (2023). Hallmarks of ageing: An expanding universe. *Cell*, 186(2), 243–278.

Meng, D., Zhang, S., Huang, Y., Mao, K. & Han, J.-D.J. (2024). Application of AI in biological age prediction. *Seminars in Cell & Developmental Biology*.

NASA (2020). The Carrington Event. <https://www.nasa.gov/history/spots-waves-and-wind-a-solar-science-timeline-full-text/>

Sharma, A., Lysenko, A., Jia, S., Boroevich, K.A. et al. (2024). Advances in AI and machine learning for predictive medicine. *Journal of Human Genetics*, 69, 487–497.

Srour, L., Bejaoui, Y., She, J., Alam, T. & El Hajj, N. (2025). Deep aging clocks: AI-powered strategies for biological age estimation. *Ageing Research Reviews*, 112, 102889.

Ying, et al. (2025). ClockBase Agent — AI-driven reanalysis of ageing intervention data, as cited in Coughlan & Faragher (2026).

Zhang, Y., Mastouri, M. & Zhang, Y. (2024). Accelerating drug discovery, development, and clinical trials by artificial intelligence. *Med*, 5(9), 1050–1070.