https://doi.org/10.1038/s43587-024-00788-9

Balancing the promise and risks of geroscience interventions

Alan A. Cohen, John R. Beard, Luigi Ferrucci, Tamàs Fülöp, Vadim N. Gladyshev, Mahdi Moqri, Marcel G. M. Olde Rikkert & Martin Picard

Although the emerging field of geroscience holds great promise for identifying new approaches to improve healthspan, several risks of the current framework are underappreciated. Long time horizons, challenges in identifying causality-driven surrogate biomarkers of aging, and the potential for biological trade-offs and antagonistic effects across various timescales mean it will be hard to know when such interventions have a net benefit. We propose eight strategies to mitigate these risks going forwards.

The geroscience hypothesis posits that by intervening in the biological mechanisms of aging, we can simultaneously reduce the risks of many age-related diseases by counteracting the aging process. As the field has grown, serious efforts have emerged to consider what geroprotective interventions (Box 1) might look like and how we might develop a clinical framework for evaluating the safety and efficacy of proposed interventions — pharmacological, epigenetic reprogramming, lifestyle or other[1]. If successful, such interventions could increase healthspan and markedly reduce the age-associated burden of disease.

This unprecedented promise comes with equally unprecedented challenges, including gaining regulatory approval. Here, we focus on one of the least appreciated challenges: that the complexity and unique nature of the aging process creates a risk that initially promising putative geroprotectors may generate long-term harm, which even a rigorous testing framework may not detect.

Unlike most chronic diseases, aging is not a mechanistically distinct pathology of a single organ or system. It arises at least in part from the impossibility of perfectly balancing competing priorities in the maintenance of the organism as a complex adaptive system, which results in trade-offs within and among systems[2]. Additionally, the aging process is heterogeneous and multifactorial, and many aging mechanisms are also functional or adaptive[3]. This complexity makes it uniquely challenging to identify which interventions may be net beneficial versus net harmful over the full life course, particularly using typical clinical trials over several years. Here, we outline these challenges and propose solutions and mitigation strategies to minimize potential risks.

The path to success and/or harm

It is already clear that some geroprotective interventions are net positive: for example, many lifestyle interventions (moderate exercise, sleep, healthy diet, social integration and so on) have few, if any, adverse effects. The challenge is to ensure that biologically targeted interventions are also free of side effects, or are at least net-positive. Figure 1 outlines how current testing frameworks might inadvertently lead to net harms. In this scenario, biomarkers of aging are developed that fulfill all criteria as currently defined[4]: they correlate with age and upstream determinants of health, predict future health outcomes better than chronological age alone, have plausible mechanistic links to the biology of aging, respond in the expected direction to putative geroprotective interventions, and so on[4]. Using these biomarkers, we identify and deploy multiple interventions to counteract aging. Millions of people — including young adults — start using them, on or off label, and many see short-term benefits. However, in the long-run, unforeseen harms become apparent, and these far outweigh the benefits. Hypotheticals are countless, but two plausible scenarios include (a) an increased long-term cancer risk due to perturbation of anticancer functions of telomeres, stem cell exhaustion or other aging mechanisms and (b) changes in inflammatory biology that perturb a delicate balance in innate immune regulation, which renders individuals more susceptible to acute infections and/or acute respiratory distress syndrome.

Of course, the risk of long-term harm is present with any novel therapy, as is the risk that pharmaceutical compounds or other interventions that logically seem as if they should work turn out in fact to have counterproductive effects. The anticancer treatment doxorubicin is effective against cancer but generates long-term cardiotoxicity[5]. Aggressively feeding critically ill patients via parenteral nutrition to prevent weight loss actually increases infections and mortality[6]. Antioxidant supplements such as vitamin E or β-carotene may increase all-cause mortality[7]. It is impossible to test any intervention across all time frames, context-dependent effects and potential side effects. However, there are several reasons that such risks are particularly salient in the geroscience context.

Unique features of aging that increase the risk of unintended consequences

First, human aging proceeds over decades, which is longer than can be followed in a clinical trial. In other medical domains, both benefits and harms can generally be perceived clearly in the short-to-medium term. However, for geroprotectors that might be given earlier in life, the success of our evaluation of both efficacy and risk hinges much more strongly on the quality of surrogate biomarkers that will be assessed in midlife, and on their ability to predict consequences we care about decades into the future — there is little else that can feasibly stand in as an outcome in clinical trials[1]. Multimorbidity and survival data that can be used in animal models (for example, of metformin[8]) may not be feasible in humans. Once deployed, it will be a long time before real-world evidence confirms or refutes the assumption that these surrogate markers accurately predict the future effect of interventions.

BOX 1

Scope and glossary

We use terms in a way consistent with the framework proposed by the Biomarkers of Aging Consortium[4] and our own perception of how they are broadly used in the field (not necessarily our own preferred definitions). Specifically, we restrict ourselves to cases in which the current regulatory and clinical testing framework is not sufficient to evaluate efficacy and safety of a proposed intervention, and for which evaluation of the benefits or harms would require use of a surrogate end point for the aging process itself. Although this is currently outside of regulatory frameworks such as of the US Food and Drug Administration (FDA) or European Medicines Agency (EMA), it is a direction actively being pursued by many academics, industry partners and advocates. The time to discuss benefits and risks is now, before major changes are made to regulatory frameworks, and indeed with an eye towards influencing any changes. Cases covered by current clinical end points and their approved surrogates are not necessarily distinct from any other medical intervention and thus are beyond the scope of our discussion. For example, we would not consider an intervention specific to atherosclerosis to be a geroprotector that acts on aging, even if the net benefit was to prevent multiple chronic diseases related to atherosclerosis.

Aging
The process of accumulation of consequences of life, such as molecular and cellular damage, that leads to functional decline, chronic diseases and ultimately mortality[4]; broadly thought to be due to the hallmarks of aging and related mechanisms.

Biological age
Conceptually, an individual's age defined by the level of age-dependent biological changes, such as molecular and cellular damage accumulation. In practical use, this is often summarized as a number (in units of time) matching the chronological age where the average person in a reference population shares the individual's level of age-dependent biological changes[4].

Biomarker of aging
A quantitative parameter of an organism that either alone or in a composite predicts biological age and ideally its changes in response to interventions[4].

Geroprotector or geroprotective intervention
A therapy or intervention that acts on one or more of the hallmarks of aging or similar processes to slow, prevent or reverse aging broadly, and thereby reduces the risk of most or all aging-related chronic diseases and/or improves healthspan.

Healthspan or healthy lifespan
The period of life before the onset of chronic disease and disabilities of aging (that is, the period of life in good health)[4].

Testing framework
A framework to evaluate the clinical efficacy and safety of putative geroprotectors for different potential target populations.

Trade-off
A situation in which improving one aspect of biology or physiology inevitably results in the worsening of another; generally considered in an evolutionary framework in which natural selection has optimized multiple conflicting needs simultaneously, which results in suboptimal status for some needs.

Surrogate biomarker or end point
A biomarker or other end point that can be used in clinical research as a sufficient proxy for a harder-to-measure true target of an intervention; in principle, improvement on the surrogate (for example, apolipoprotein B) can be understood as evidence that the intervention would improve the true target (for example, atherosclerosis or myocardial infarction).

Second, sufficiently high-quality surrogate biomarkers of aging may be hard to identify because aging itself is particularly challenging to define and quantify. Almost all researchers now agree that aging is multifactorial rather than a single, discrete biological process. Even in a best-case scenario, there are likely to be multiple aspects of aging that are weakly correlated with each other[2]. For example, presbyopia is caused by universal, age-dependent lens growth and accumulation of protein aggregates due to processes largely independent of other aging mechanisms; it is unlikely to be fully captured in standard proposed biomarkers. Unless biomarkers of aging identified as surrogate end points include all such specific cases, they may be incomplete, and we may not realize this.

Third, the combination of antagonistic pleiotropy and the interconnectedness across systems means that aging biology is particularly susceptible to unexpected cost–benefit trade-offs that appear at different timescales[1]. For example, in genetically diverse mice, caloric restriction and diet have variable effects across ages and different dimensions of health and lifespan, which is indicative of trade-offs[9,10]. Unlike in many other biomedical fields, aging interventions are expected to affect many interrelated physiological domains, and there is greater risk that benefits in one domain have negative effects in another. Rapamycin, for example, has side effects on infection risk and metabolism that might not be picked up by an aging clock and that might be masked in animals studied under germ-free conditions[11]. The biological pathways most related to aging are often those related to balancing the competing demands of the organism: growth or reproduction versus maintenance, low cancer risk versus replicative potential, current versus future function, and so on. We emphasize potential trade-offs between causes of mortality — most notably, between cancer risk and the hallmarks of aging[12]. At least three aging hallmarks may have evolved as cancer protection mechanisms (cellular senescence, telomere shortening and stem cell exhaustion); human aging biology may thus have been calibrated by selection to equilibrate the risks of cancer and aging[12]. Some geroprotectors could tip the scales toward

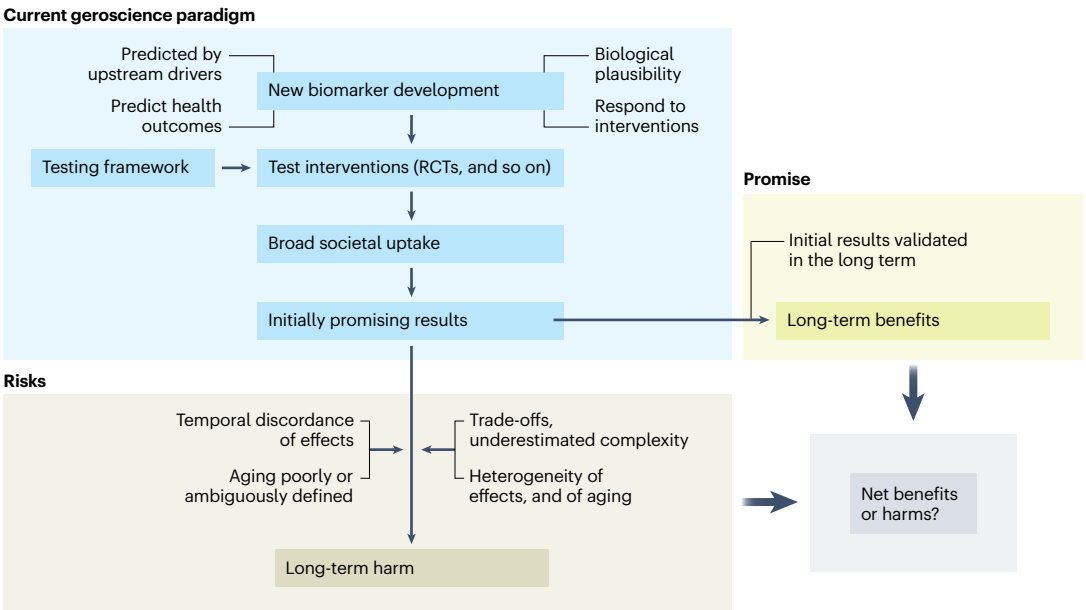

Fig. 1 | How a putative geroprotector that causes long-term harm could slip past the current biomarker-oriented testing framework. The testing framewok proposed under the current geroscience paradigm could produce evidence for interventions that generate both benefits and harms, but is not equipped to evaluate the relative balance between the two.

one side of such a trade-off, the other side of which may manifest on another timescale.

Fourth, aging is particularly heterogeneous, reflecting the cumulative effects of influences across a person's life interacting with individual traits. Indeed, many of the age-related diseases that geroscience hopes to target (for example, cardiovascular disease, Alzheimer's disease and diabetes) are rare in some non-industrialized populations, such as Tsimane horticulturalists[13]. This suggests that biomarkers of aging developed in our industrialized societies might — and perhaps should — reflect the cumulative biological effects of industrialized lives, in addition to universal human aging processes. Interventions that affect such biomarkers may be those that counter factors such as overnutrition and sedentary lifestyles, rather than biological aging. However, many of the individuals who most aggressively seek antiaging therapies (including off label) may be exactly those individuals whose exposures and lifestyles are already better, and who thus have the smallest potential to benefit from such interventions and the highest risk of net harm. This is because maintaining health almost certainly requires getting the correct balance on implicated pathways, rather than pushing any to an extreme. Indeed, many healthy individuals are already taking metformin, rapalogs, senolytics, NAD^+ precursors and so on to slow their aging rate before there is evidence even at the population level for benefits or safety — much less before there is a sufficient body of evidence for precision prescription. Studies published in basic biology are thus already having public health consequences.

These four facets interact to compound the potential risk of applying putative geroprotectors in the absence of a sufficient understanding of the complex interplay among pathology, compensation, homeodynamics and context during aging. The likely incompleteness of biomarkers of aging will result in more trouble assessing trade-offs versus global benefit, particularly when harms and benefits appear on very different timescales. Interventions chosen as promising are, by definition, those that have short-term benefits; owing to trade-offs, these impressive benefits could actually be warnings of even greater harms down the road. The promise is great, but the risks are not intuitive. The threshold for implementation of a geroprotective intervention should be particularly high in this context, especially for geroprotectors that require regular administration over years.

A proposed agenda for mindfully developing and testing putative geroprotectors

Although there may not be a complete solution to the challenges outlined above, we have identified eight strategies to guide the development of putative geroprotectors and at least partially mitigate risks (Fig. 2). First, lifestyle or contextual interventions present less risk than targeted biological interventions; because they holistically stimulate endogenous networks of adaptive processes rather than a single molecular target that forces compensatory recalibrations, the biological changes produced by lifestyle changes are probably true benefits and not just one side of a trade-off. Validation of biomarkers of aging should thus prioritize biomarkers that are capable of showing meaningful biological responses to diverse lifestyle or contextual interventions, and biomarkers that only respond to biologically targeted interventions (rapamycin, senolytics, therapeutic plasma exchange and so on) should be used with caution.

Second, we should consider targeting health in addition to or instead of aging. Broad biomarkers of health are under development and if we can succeed in stabilizing basic homeostatic processes, this avoids most of the concerns raised above.

Third, we should continue to develop, and deploy in tandem, biomarkers of early cancer risk[14] and other aspects of safety in parallel with aging biomarkers. Cancer is probably the most obvious potential trade-off and long-term risk for interventions that aim to preserve cellular youthfulness and replicative potential[12].

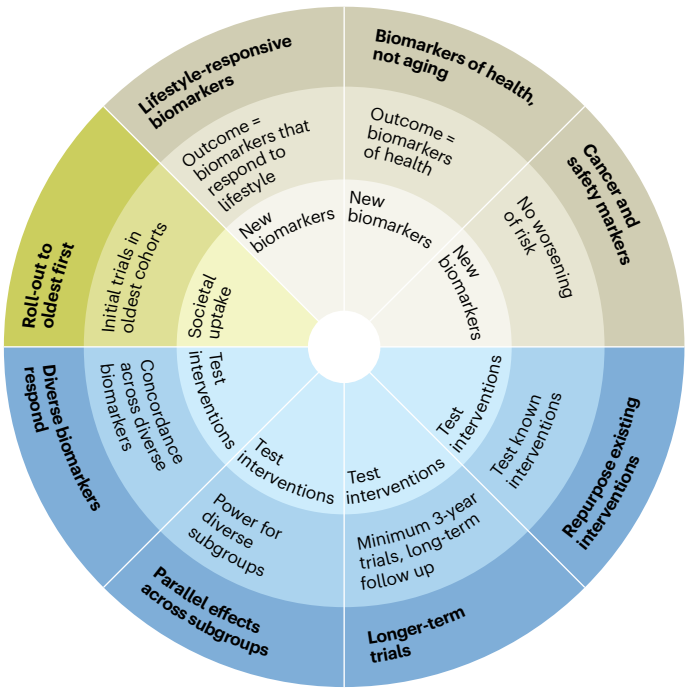

Fig. 2 | Eight strategies to guide the development and testing of putative geroprotectors. The outer ring is the strategy. The middle ring represents the clinical trial design elements to respond to the strategy. The inner ring reflects the stage of the geroscience paradigm at which the strategy operates (from Fig. 1).

Fourth, we should first assess the effect of existing drugs with known safety profiles, such as metformin, and of lifestyle or contextual interventions such as exercise. For example, the effect of new drugs could be benchmarked against that of exercise. This does not guarantee long-term net benefit and does not address safety or benefits in individuals who would not typically receive the drug, but it may mitigate the most serious safety concerns.

Fifth, benefits should be seen over the medium term — unlike other clinical trials, antiaging interventions should be tested over many years (five years should not be considered long term), and multidecade monitoring of clinical trial participants after the intervention and broader post-market monitoring should be mandated.

Sixth, at least initially, effects should not diverge in subgroups (most importantly by age, but also by frailty status, resilience and in diverse human populations), as divergent effects could indicate complex biological processes and trade-offs at work. This point would be less important as a body of evidence for precision geroscience accrues.

Seventh, putative geroprotectors should be tested against a full array of aging biomarkers. Benefits should be present for many of the best-validated biomarkers, including multiple domain-specific biomarkers. Having any aging biomarkers that worsen could indicate an underlying trade-off, and having a benefit in only one domain should flag doubts on net benefits. Ratios that are susceptible to denominator effects should be used with caution.

Lastly, and perhaps most importantly, rollout should start with the oldest adults: those with the shortest life expectancy and thus the lowest follow-up time and most easily measurable risk of long-term harms. As longer-term safety becomes apparent in older individuals, access could gradually be granted to younger and younger individuals, rolling out over a period of decades. Note that, similar to interventions on Alzheimer's disease too late in the process, absence of benefits in older individuals does not preclude benefits in younger individuals; the goal of starting with older adults is to evaluate safety rather than efficacy. However, many geroprotectors may have age-specific periods of benefit, as found for other drugs. For example, the Frail-AF trial found important adverse effects of direct oral anticoagulants in frail older adults underrepresented in registration trials in younger cohorts[15]. Age targeting will be crucial, and the safe strategy is to work backwards from the oldest ages.

The effectiveness of these and other mitigation strategies as an ensemble will need to be carefully monitored moving forward. One barrier is that the current regulatory framework in the USA permitting off-label (or accelerated) use might make it hard to tightly control access, and cultural trends might drive widespread, dangerous use of any FDA-approved products for which there is even partial evidence of benefit.

Conclusion

Human biology has evolved for a reason, and the negative features of aging are presumably accompanied by evolutionary advantages. At least since Hippocrates warned of the risk of harming patients, doctors have struggled to balance trust in the homeostatic mechanisms of the body and the need for interventions to counter very real pathology and dysfunction. The context of clinical research on putative geroprotectors has numerous particularities that make it ill-suited to the classical evidence-based medicine framework, and many leaders in our field have done admirable work in identifying and addressing these challenges[1,4]. Nonetheless, we must be careful not to let hope and enthusiasm outrun evidence and the entangled balance of biological complexity, and result in substantial long-term consequences. We hope that this framework for geroprotector development and testing will maximize the probability that we identify, develop and translate interventions that are unambiguously protective and health-promoting across timescales. However, it is vital that the risks and benefits should be considered continually. Lessons learned will probably apply more broadly to testing and implementing preventive approaches in medicine and public health. Critically, we must improve our communication with the lay public about the nascent state of our science and the risks of taking unproven and potentially harmful therapies on the basis of the current, insufficient evidence. As scientists, we have an obligation to be humble about our current state of knowledge and not overstate the potential or the evidence in ways that could generate popular hype and endanger the public.

Alan A. Cohen [1] ✉, John R. Beard [2], Luigi Ferrucci [3], Tamàs Fülöp [4,5], Vadim N. Gladyshev [6], Mahdi Moqri [6], Marcel G. M. Olde Rikkert [7] & Martin Picard [8,9]

[1]Department of Environmental Health Sciences and the Butler Columbia Aging Center, Mailman School of Public Health, Columbia University, New York, NY, USA. [2]Departments of Health Policy Management and Epidemiology and the Butler Columbia Aging Center, Mailman School of Public Health, Columbia University, New York, NY, USA. [3]Translational Gerontology Branch, National Institute on Aging, NIH, Baltimore, MD, USA. [4]Research Center on Aging and Research Center of Centre Hospitalier Universitaire de Sherbrooke, Sherbrooke, Quebec, Canada. [5]Faculty of Medicine and Health Sciences, University of Sherbrooke, Sherbrooke, Quebec, Canada.

[6]Division of Genetics, Department of Medicine, Brigham and Women's Hospital, Harvard Medical School, Boston, MA, USA. [7]Department of Geriatrics, Radboud University Medical Centre, Nijmegen, The Netherlands. [8]Departments of Psychiatry and Neurology, Columbia University Irving Medical Center, New York, NY, USA. [9]The Butler Columbia Aging Center, Mailman School of Public Health, Columbia University, New York, NY, USA.
✉e-mail: aac2277@cumc.columbia.edu

Published online: 03 January 2025

References

1. Rolland, Y. et al. *Nat. Commun.* **14**, 5038 (2023).
2. Cohen, A. A. et al. *Nat. Aging* **2**, 580–591 (2022).
3. Ogrodnik, M. & Gladyshev, V. N. *Nat. Aging* **3**, 766–775 (2023).
4. Moqri, M. et al. *Cell* **186**, 3758–3775 (2023).
5. Linders, A. N. et al. *NPJ Aging* **10**, 9 (2024).
6. Casaer, M. P. et al. *N. Engl. J. Med.* **365**, 506–517 (2011).
7. Miller, E. R. III et al. *Ann. Intern. Med.* **142**, 37–46 (2005).
8. Martin-Montalvo, A. et al. *Nat. Commun.* **4**, 2192 (2013).
9. Harper, J. M., Leathers, C. W. & Austad, S. N. *Aging Cell* **5**, 441–449 (2006).
10. Di Francesco, A. et al. *Nature* **634**, 684–692 (2024).
11. Lee, D. J. W., Hodzic Kuerec, A. & Maier, A. B. *Lancet Healthy Longev.* **5**, e152–e162 (2024).
12. Pavard, S. & Metcalf, C. J. E. in *Human Evolutionary Demography*. (eds Burger, O. et al.) 715–740 (OpenBook, 2019).
13. Lea, A. J. et al. *PLoS Biol.* **21**, e3002311 (2023).
14. Liu, M. C., Oxnard, G. R., Klein, E. A., Swanton, C. & Seiden, M. V. *Ann. Oncol.* **31**, 745–759 (2020).
15. Joosten, L. P. T. et al. *Circulation* **149**, 279–289 (2024).

Acknowledgements

We thank D. Belsky for helpful comments on the manuscript.

Author contributions

A.A.C. conceived and drafted the manuscript. All authors contributed ideas, revised and edited the manuscript.

Competing interests

A.A.C. is founder and CEO at Oken Health. The other authors declare no competing interests.

Additional information

Peer review information *Nature Aging* thanks James Peyer, Heike Bischoff-Ferrari, and the other, anonymous, reviewer(s) for their contribution to the peer review of this work.