1: Introduction

So far, humanity has surrendered to the horrors of aging and death as unavoidable certainties.

But in the age of biological science, unlimited healthy lifespans are bound to happen.

The only question is whether we will act to make them happen in our lifetime.

- The Vitalism Whitepaper, April 2023

If Roman Emperor Marcus Aurelius wanted to build a jet plane so he could conquer the world, he would likely have failed even if his "Grand Strategy" were to put all of Rome's GDP and millions of slaves to the task. The Romans were simply too early in the technology tree.

Having said that, could Marcus Aurelius have realized total mobilization would likely fail?

When a big goal is distant, complex, and requires coordinated action, choosing the right Grand Strategy is critical. To conceive of, evaluate, and select such a strategy requires a mental model that explains why it is likely to work.

Such a mental model is a **Theory of Change**.

A **Theory of Change** explains how actions are expected to lead to outcomes through a sequence of causal steps, assumptions, and intermediates. Nearly all sustained actions occur within the context of such theories, sometimes consciously but often unconsciously. Even when no theory of change is articulated, one is always implicit.

Therefore, any individual who wants to have a big impact, should have a clear Theory of Change to help answer their very practical question - "What should I do?"

In Marcus Aurelius' case, the mere act of articulating a Theory of Change would have confronted him with the many unknowns between Roman science and technology and a jet. Maybe instead of building a jet he would have turned his attention to better catapults or discovered gunpowder.

When it comes to Aging, we face a similar conundrum. Should one train as a biologist or other type of scientist? Start a company? Start a longevity-focused fund? Join a lab? Run for political office?

To find an answer, it's useful to start with the two key factors that underpin every Theory of Change; a Theory's falsifiability and its distinct point of view about the binding constraint[1].

[1] A binding constraint is the single most limiting factor that prevents a system from achieving a higher level of performance, such that relieving it would produce the largest immediate improvement.

It is "binding" because it actively holds the system back. If you remove it, output increases. If you optimize anything else, nothing meaningful changes.

Let’s examine some familiar cases where either a Theory of Change or a Grand Strategy (or both) were clearly defined and put to the test. The Second World War provides good examples.

As our first example, we know UK Prime Minister Neville Chamberlain's Theory of Change preceding World War II clearly failed, because it was fully tested and falsified.

Chamberlain’s theory was that limited concessions to Germany would satisfy it and preserve peace. The binding constraint in his view was unresolved injustice from World War I combined with insufficient diplomatic flexibility. If those factors were addressed, escalation could be avoided.

The resulting Grand Strategy was straightforward: Address territorial grievances (e.g., the Sudetenland) → reduce nationalist resentment → moderate German behavior → avoid catastrophic war.

Chamberlain was successful at carrying out his strategy; but it failed because it was based on the wrong Theory of Change. Germany was never going to be mollified by concessions.

Conversely, in World War II, Nazi Germany’s Theory of Change was never fully tested, but its Grand Strategy failed.

Germany’s Theory of Change was that if it could achieve continental dominance fast enough, the combined industrial and resource base of Europe + western Russia would deter challengers, leading to long-term global dominance.

This directly birthed the Nazi Grand Strategy of Blitzkrieg, demographic engineering and aggressive oppression, with the goal of resource self-sufficiency and an unchallengeable geographic position.

The Nazis knew the adage “Never fight a land war in Russia in winter” but did it anyway, because it was a critical part of their Theory of Change.

Fortunately, we will never know if the Nazi Theory of Change was right, because their Grand Strategy failed to achieve it. They failed to conquer the Soviet Union, never consolidated Eurasia, and never tested whether a German-dominated continental bloc could deter American intervention.

Importantly, America took Germany’s Theory of Change seriously; it acknowledged that Nazi consolidation of the continent could dramatically threaten the US and therefore took strong action.

This shaped America’s Grand Strategy as one of full industrial mobilization (the “Arsenal of Democracy”) combined with tight alliance coordination. The aim was sustained production superiority to deny Nazi Germany the ability to consolidate gains or outlast the coalition, and ultimately drive them to unconditional surrender.

Fast forward to 2026 and the deadly threat we still face: Aging.

Arguably, when my father was born in 1945 it would have been folly to assume we could solve death and its primary cause, aging, by merely throwing resources and attention at the problem. At the time, we had no idea how human heredity happens, let alone clear evidence that aging is malleable.

Today, the situation is very different.

Lifespan extension has been repeatedly achieved in model organisms, from caloric restriction in mice to dramatic genetic interventions in worms. We discovered species that defy conventional aging patterns and successfully transferred these traits to other species[2]. Yamanaka and others showed it is possible to make human cells rejuvenate, and we've cloned various species, demonstrating normal lifespans in the respective offspring[3]. Machine learning has made progress on longstanding challenges such as protein folding, and our ability to edit DNA has dramatically advanced since the discovery of CRISPR. These and other advances in our ability to understand and engineer biology are why many leading scientists are convinced that solving aging is foreseeable, potentially in our lifetimes.

If we do, we will enable billions to live on in a future of health and abundance. However, while public sentiment is gravitating towards longevity and demand for interventions grows, our progress, resourcing and attention to the matter have been catastrophically small.

1.1 We need to choose the right Theory of Change to move as fast as possible
A poor theory of change could misallocate efforts and delay progress by decades, even if it is internally coherent and technically sophisticated. Conversely, a well-chosen theory of change will identify the true constraints on progress and can thus dramatically accelerate outcomes.

That's why understanding and scrutinizing Theories of Change in longevity is a critical **prerequisite to taking effective action**.

For example, one common Theory of Change assumes we don't need regulatory or even much funding change, but rather one landmark commercial success through the existing FDA pathway. On its own, the theory goes, such a success would activate a cascading capitalist flywheel which will, in short order (i.e. fast enough for those alive today) cause a tsunami of funding, attention and other necessary affordances to make swift progress.

An alternative theory is that regulations are the binding constraint, and that we must either fully reform the FDA or find alternative jurisdictions which will allow bringing new therapies to humans faster and, in their mind, better.

---

[2] https://pmc.ncbi.nlm.nih.gov/articles/PMC137187/, https://pubmed.ncbi.nlm.nih.gov/37612507/
[3] "Dolly the Sheep" was famously short lived, but K. D. Sinclair, S. A. Corr, C. G. Gutierrez, P. A. Fisher, J.-H. Lee et al., "Healthy ageing of cloned sheep," Nature Communications, 2016 examined 13 aged cloned sheep, including four derived from the same cell line as Dolly, and found no evidence of long-term detrimental health effects or premature aging when compared to naturally conceived sheep. The authors conclude that the clones aged normally and did not exhibit accelerated age-related diseases.

Yet another approach asserts we need a scientific breakthrough, which can be achieve with existing resources, if only we assemble the correct teams and allow them to work on unconventional strategies.

1.2 To choose the right Theory of Change, we need to define what makes it good
Given varied, oft competing, Theories of Change and Grand Strategies in longevity; having a framework to define and evaluate them is essential.

A Theory of Change is a bet about what's truly in the way. Is the real problem that we still don't understand aging deeply enough, meaning we need more basic science and steady government funding?

Or is the problem momentum and evidence of success, meaning we need a clear commercial breakthrough to unlock capital, talent, and regulations?

Different people see different bottlenecks, and they commit to a path accordingly. They spend time, money, and energy to remove that constraint at the expense of others.

While Theories of Change can coexist, they often compete for attention and resources given finite time, capital, talent, and institutional capacity. The central strategic question is therefore not whether a given action is helpful in isolation, but whether it unlocks the most progress.

In longevity, this framing makes Theories of Change especially important, since causal chains are long, feedback is slow, and incorrect diagnoses of the bottleneck may persist for decades before failure is evident.

And it doesn't take many people to make a Theory of Change dominant. Often, a small group of determined and resourceful people can create an outsized impact by redirecting attention, policy and resources. Ultimately, once attention is redirected at the global level, people react to external incentives; they work on things they can get paid for, and are societally legitimate.

1.3 We can learn from historical precedent
Major social, scientific and technological achievements are often described retrospectively as the natural or inevitable result of progress. But they are almost always preceded by disagreement over which Theory of Change should guide action.

These disagreements reveal that success often depends less on the clarity of the end goal than on correctly diagnosing the dominant constraints and levers, and selecting a strategy to impact them.

Let's examine some examples.

Example I: The Apollo Program

From Kennedy's own words, the primary goal of the Apollo Program was: *To prevent the Soviet Union from establishing perceived technological–ideological superiority that could shift global alliances and deterrence credibility.*

Not "space exploration", "science" or "economic benefit"

By 1961, Kennedy knew Apollo would cost ~$250 billion (in 2026 dollars). NASA Administrator James Webb repeatedly warned him of the scale and Kennedy agonized over it.

In private recordings (Oval Office tapes, 1962), Kennedy says:
"We're talking about a lot of money, for what?" and "If we didn't have the Cold War, I wouldn't be spending this kind of money."

Kennedy's internal cost-benefit test was never: "Is Apollo worth $250B?" but rather: "Is there a cheaper or safer Theory of Change to achieve **the same geopolitical goal**?". And indeed, there were cheaper, competing Theories of Change.

Competing Theory of Change #1: Incremental space superiority (cheaper, slower)

Theory of Change: *Gradually outpace the USSR across many space milestones (satellites, rendezvous, space stations), gaining credibility over time.*

The Grand Strategy to support this Theory of Change would have been much cheaper year-to-year, built sustained capability and allowed dynamic diversion of resources to dual-use efforts.

Why it lost:

- Too slow relative to Soviet propaganda cycles
- Too ambiguous: the USSR could claim parity or symbolic wins
- No single "checkmate" moment

Competing Theory of Change #2: Robotic missions (much cheaper)

Theory of Change: *Demonstrate technological superiority through robotic probes, which are cheaper, safer, and scientifically rich.*

This strategy was scientifically compelling, orders of magnitude cheaper and had lower political risk (e.g. astronaut deaths).

Why it lost:

- Robots were less engaging to the public; Humans mattered symbolically
- The USSR could plausibly match outcomes

Competing Theory of Change #3: Military and economic dominance alone (much cheaper)

Theory of Change: *Rely on nuclear deterrence, economic growth, and conventional power to demonstrate superiority; space is a distraction.*

The U.S. already had advantages here, requiring no new institutions and enjoying direct amplification of investment in existing security concerns.

Why it lost:

- Space was conceptually linked to ICBM capability, even if the Apollo Program was not directly meant to defuse such risks; Sputnik suggested the US may already be behind
- Non-aligned nations interpreted space success as future power
- Space had become a symbolic battleground

Key Takeaway
Kennedy didn't necessarily believe Apollo was the only viable path, but he believed it was the lowest-risk path to a decisive outcome under radical uncertainty. While it had high cost, it also had high robustness, was hard to fake, match or reinterpret. And most importantly, Kennedy believed it would be a decisive win if successful.

Example II: The HIV Crisis of the 1980s

For most, the primary goal was: *Reduce AIDS mortality as fast as possible.*

But in the epidemic's early years, the dominant institutional posture (NIH/FDA/academic medicine) effectively treated AIDS as a problem where progress would necessarily be slow, because the science was immature and the risks of acting on weak evidence were high.

The Theory of Change was that we lacked basic science. Only once we had it could we identify targets leading to drugs which, after proper trials, would save lives. Sadly, in a fast-moving epidemic speed defines body counts, and this theory of change was unacceptable to victims and their loved ones.

In the same way Apollo forced Kennedy to ask whether there was a cheaper way to achieve the same geopolitical effect, AIDS patients and activists asked:

Is there a faster, less institutionally conservative way to achieve survival outcomes, even with uncertainty and risk?

They believed the answer was yes and built an alternative Theory of Change around it.

The Dominant Theory of Change = Orthodoxy: Scientific uncertainty is the binding constraint (therefore slow progress is unavoidable)

Status quo: *Proceed through conventional biomedical pipelines. Prioritize certainty and safety. Assume that without deep understanding, accelerated access won't help much and may harm.*

This seemed plausible because HIV/AIDS was novel; with its causal agent and mechanisms unclear. Also, drug development and antiviral success historically had low base rates and regulators were shaped by past disasters and risk aversion.

The status quo assumed the binding constraint was knowledge, not institutional speed. But also, there was a posture that timelines were negotiable if the result was scientific certainty, even if some deaths were then unavoidable.

Competing Theory of Change #1: Find ways to move faster because time is the core constraint (people are dying!)

Theory of Change: *Treat speed and iteration tempo as the primary constraints. Increase attention and resources, force institutional adaptation, and learn faster through expanded trials and real-world use.*

If urgency was the top concern, resources + institutional reform could be combined to deliver faster trials, broader access, faster learning and better therapies.

Patient activists (e.g., ACT UP and allied groups), clinicians, community organizations, and later reformist insiders pushed this theory of change. Proponents believed that when baseline mortality is extremely high, the risk calculus changes: patients are willing to accept uncertainty because the alternative is death.

Their sustained effort included pressure campaigns aimed at the NIH and FDA, challenging protocols, inclusion criteria, endpoints, and timelines, and betting that real-world use and expanded access can generate knowledge faster than a slow, exclusionary Randomized Control Trial (RCT) pipeline.

Why it won (eventually):

- It directly targeted a binding constraint the system had treated as fixed: speed of discovery.
- It created new mechanisms (accelerated review norms, expanded access/parallel-track style thinking, patient voice in trial design) that increased the effective learning rate of the entire ecosystem.

Competing Theory of Change #2: Prevention-first as a primary lever (reduce incidence; therapy will lag)

Theory of Change: *Focus resources on prevention and behavioral/public health interventions (education, harm reduction, screening), because drug discovery will be slow and uncertain.*

This theory seemed plausible to start because therapies were speculative early on and prevention can scale faster than drug development. There was also public health precedent from past epidemics.

Why it lost

- Prevention does not solve the urgent survival problem of those already infected.
- The therapy pathway ultimately proved more tractable than early pessimists expected (especially once combination approaches matured). Prevention was not "wrong" it just answered a different version of the goal.

If the goal is defined as "create safe, effective therapies while preserving status quo scientific institutional process", the early institutional Theory of Change can be defended as cautious and rational.

But if the goal is defined as "minimize deaths under time pressure," the status quo seems miscalibrated, because it treats speed as secondary and uncertainty as disqualifying, in a setting where uncertainty was unavoidable.

Key Takeaway

The competing Theory of Change modified the system so that science could iterate fast enough to matter. The successful outcome, effective combination therapies and sharp mortality reductions, arrived only after institutional adaptation increased the ecosystem's learning rate.

Conclusion

Choosing the right Theory of Change may be the most important decision to save countless lives. Therefore, choosing the right Theory of Change, and pursuing the most effective Grand Strategy are heavy responsibilities. In the next section, we explore the dominant status quo in how people end up committing to specific Theories of Change.

2: How we Choose

Famous venture capitalist Marc Andreessen put it well: *"When a great team meets a lousy market, market wins. When a lousy team meets a great market, market wins."*

He then went on to explain: *"In a great market - a market with lots of real potential customers, the market pulls product out of the startup… Conversely, in a terrible market, you can have the best product in the world and an absolutely killer team, and it doesn't matter, you're going to fail."*

Choosing the wrong Theory of Change can similarly be catastrophic.

Ironically, in the same way many if not most founders spend little effort to choose the right market, many longevity doers spend little time or energy to choose a Theory of Change.

This may seem counter-intuitive. If choosing a Theory of Change is a critical first step to impact longevity, how could so few people engage with it directly?

Sarah Paine might explain this pattern as a failure of **Grand Strategy,** in the sense that any such strategy should have a coherent Theory of Change. The common failure mode is that actors substitute competent action for coherent causal reasoning. In Paine's account of Imperial Japan in WWII, Japan accumulated locally rational, well-executed actions; territorial expansion, naval buildup, decisive early battles, without ever articulating a credible theory for how it could win a long, industrial war against materially superior opponents. Strategic evaluation was avoided, perhaps because it would have revealed an uncomfortable truth: the strategic odds were poor.

Instead, leaders focused on questions that produced actionable plans. Paine might argue that longevity efforts exhibit the same failure mode. People choose actions that feel serious and legible (e.g. start or join a company, study or join a lab, volunteer for a non-profit) and avoid explicit Theory-of-Change evaluation because carefully examining whether one's favored actions can work is cognitively expensive, and potentially destabilizing. This can lead to impressive execution, wasted on a catastrophic trajectory.

Across domains the same pattern appears: when outcomes are distant, uncertain, and system-level, humans and institutions substitute action selection for causal reasoning.

This section attempts to explain strategic choice behaviors when it comes to longevity and propose a mental model for how people actually end up committing to specific Theories of Change.

Theories of Change and Why People Avoid Them

Imagine you have decided that extending healthy lifespans is your priority, and you now consider your next immediate step. Ostensibly, choosing a Theory of Change will determine your preferred Grand Strategy and therefore also the highest value action.

Instead, what seems natural to most is to consider concrete actions first, such as starting a company, joining a lab, investing, getting a degree, or volunteering with a non-profit.

If Theories of Change should be upstream of actions, why do most people avoid them?

Five barriers stand out:

1. Short-termism bias - people focus on immediate rather than long-term outcomes
2. Availability Heuristic - a tendency toward what's immediately accessible and actionable
3. Complexity Aversion / Cognitive Load Theory - thinking in abstractions requires significant intelligence, time investment, and has perceived unpredictable ROI
4. Temporal Discounting - People discount future value compared to present or closer value
5. Probability Neglect - People tend to round down low probabilities to zero

An evolutionary explanation reinforces these dynamics. Humans did not evolve to construct or optimize abstract, long-horizon causal models. Instead, historically, survival and reproduction depended on responding to immediate, legible signals: food, threats, status, and short-term social coordination.

This bias persists in modern institutional contexts. Individuals and organizations systematically favor actions with visible outputs, short feedback loops, and easily measurable progress, even when these proxies are poorly correlated with long-term outcomes. As a result, people and institutions drift toward Goodhart metrics and actions, optimizing what is measurable rather than what matters, while accumulating inertia. Explicit Theory of Change evaluation, which resists legibility and defers payoff, is therefore structurally disfavored despite its central importance.

Mimetics add fuel to the fire. Under uncertainty, individuals and organizations imitate the forms of other organizations (DiMaggio & Powell). In longevity, this manifests as repeated replication of familiar structures like companies, labs, and nonprofits, even when the underlying Theory of Change is unexamined.

Moreover, the mapping between actions and Theories of Change is often informal and uneven. In some cases, actions correspond tightly to a specific theory; in others, they relate loosely, overlap with multiple theories, or drift over time without re-evaluation. Because the theory is rarely articulated, inconsistencies, redundancies, and blind spots go unnoticed.

This creates a perverse bias towards availability and representation heuristics; individuals tend to overweight Theories of Change that map cleanly onto actions they are already positioned, trained, or incentivized to take. Scientists favor research-centric theories; entrepreneurs favor

venture-driven ones; advocates favor policy or cultural change. Rather than asking “Which Theory of Change is most likely to work?”, the implicit question becomes “Which Theory of Change justifies what I can already do?” Over time, this reinforces path dependence and fragments effort across locally coherent but globally suboptimal strategies.

In discussing “Emergent strategy” failure modes Henry Mintzberg argued that strategy often emerges implicitly from patterns of action rather than explicit planning. The dark side of this is that actors drift into strategies they cannot strongly articulate, defend, or revise. Mintzberg’s observations explain why rational actors who rarely thought about strategy often act under the impression they *do* have a coherent strategy.

Action feels productive; and feedback accrues to execution, not to the quality of the underlying causal model. Yet for long-horizon, system-level goals like extending healthy human lifespans, the highest-leverage mistakes occur upstream, at the level of theory, not execution.
High-impact, long horizon, non-standard approaches are thus often implicitly discarded rather than consciously evaluated.

And once actions begin, they generate careers, funding streams, reputations, and identities that make Theory-of-Change revision costly. This explains why drift persists even when evidence accumulates that the original theory is weak.

How People Choose Actions

Even without a Theory of Change, people rarely choose actions randomly or based on psychological biases alone. Instead, they rely on intuition guided by a loose, often unarticulated utilitarian calculus, one that estimates whether personal effort is likely to matter.

This can be roughly formalized as the Longevity Action Equation:

Action ≈ Age x Agency x All-Else-Equal Survival Probability

Where:

- **Age** reflects remaining expected lifespan and perceived urgency
- **Agency** reflects one’s perceived ability to influence outcomes (skills, capital, access, credibility)
- **All-Else-Equal Survival Probability** reflects a baseline belief about whether longevity breakthroughs will occur without one’s involvement

This equation is not a claim about what people should do. It is a descriptive model of how people implicitly reason when deciding whether and how to engage. For example, someone may never take action if they think aging is not modifiable. Similarly, someone who believes the solution to aging will almost certainly happen soon may choose to spend their time elsewhere. Let’s examine each of the factors more closely.

Age

Age shapes action by determining both perceived urgency and the size of the remaining time window in which action can plausibly matter. It does not simply increase motivation linearly; it changes which actions feel rational at all.

For many under forty, even pessimistic assessments of longevity progress are buffered by a sense of abundance. There appears to be decades to wait, to let others take risks first. Aging feels distant, and abstract, while the perceived costs of action - career risk, foregone options, financial instability - loom larger than uncertain future benefits. In this context, patience and deferral are locally rational, even for individuals who strongly endorse the goal of extending healthy lifespans.

As age increases, the calculus shifts. The remaining time horizon narrows, uncertainty compounds, and the cost of waiting rises sharply. Individuals in this range are more likely to pursue direct interventions, joining companies or labs, funding research, or building institutions, because marginal gains plausibly affect outcomes within their expected lifespan.

But this action-orientation to solve the core problem may be short lived. Beyond a further threshold, often when people are in their 50s or 60s, age and probability assessments jointly imply that breakthroughs are unlikely to arrive in time, and the action space contracts again. Long-horizon interventions lose relevance, and attention shifts toward strategies with potential immediate benefit. At this point, support for incremental innovation that may yield quick results could make sense, or more commonly, preservation-oriented actions such as cryonics as last-resort asymmetry: a bet that bypasses timelines altogether and requires only the refusal to accept irreversible loss.

Age, therefore, is not merely a multiplier on motivation. It reshapes the action set itself, determining whether and how waiting, building, or preservation dominate the decision landscape (see *Figure 1*).

Figure 1



Agency
Agency determines not *whether* action is desirable, but whether it feels possible. It reflects an individual's perceived ability to meaningfully influence outcomes, shaped by skills, capital, credibility, network access, and tolerance for personal disruption.

At low levels of agency, even strong interest does not translate into action. The individual may care deeply about longevity, but lacks the capacity to take concrete steps. In this context, inaction is not apathy; it gravitates to helplessness or resignation. Concern is expressed through consumption, discussion, or symbolic support rather than commitment.

Low agency is distinct from low capacity. Some people are in a life situation that substantially limits what they can do (e.g. given financial circumstance, family responsibilities etc). Yet even in that circumstance, a high agency person will find ways to do things within their scope; for example they can donate according to their ability, speak out online, connect with likeminded people and so forth.

At moderate levels of agency, action becomes conceivable but remains fragile. Individuals in this range often look to join existing initiatives rather than create new ones. However, this requires a rare alignment: the initiative must match their skills, values, risk tolerance, and compensation needs at the right moment. Because such alignment is uncommon, especially in frontier fields like longevity, many with meaningful agency still fail to act. Opportunity, not intent, becomes the binding constraint.

At high levels of agency, the opportunity constraint disappears. Individuals with minimal capital, autonomy, credibility, and risk tolerance are not dependent on existing structures to legitimize their action. They start new efforts, reshape existing ones, or assume leadership roles that redirect resources and strategy. In this regime, the decision is not whether to act, but *how* to deploy agency most effectively.

Agency therefore functions as a gating variable. Interest without agency produces concern. Moderate agency without opportunity produces frustration. Only when agency crosses a threshold does action reliably follow.

All-Else-Equal Survival Probability
There is a tight parallel between one's perceived chance of making it and the degree of risk one is willing to accept. When baseline odds are high, risk feels unnecessary; when they are moderate, calculated risk is tolerable; and when they collapse toward zero, risk aversion ceases to be rational.

Risk tolerance, in other words, is not a personality trait but a function of perceived survival odds.

Probability assessments shape whether action feels optional, urgent, or futile by determining how much marginal impact an individual believes their effort can plausibly have.

When the perceived probability of longevity progress is high, action often feels unnecessary and waiting becomes rational. Individuals in this regime defer, free-ride, or focus on local optimization, relying on time and collective effort to solve their problem.

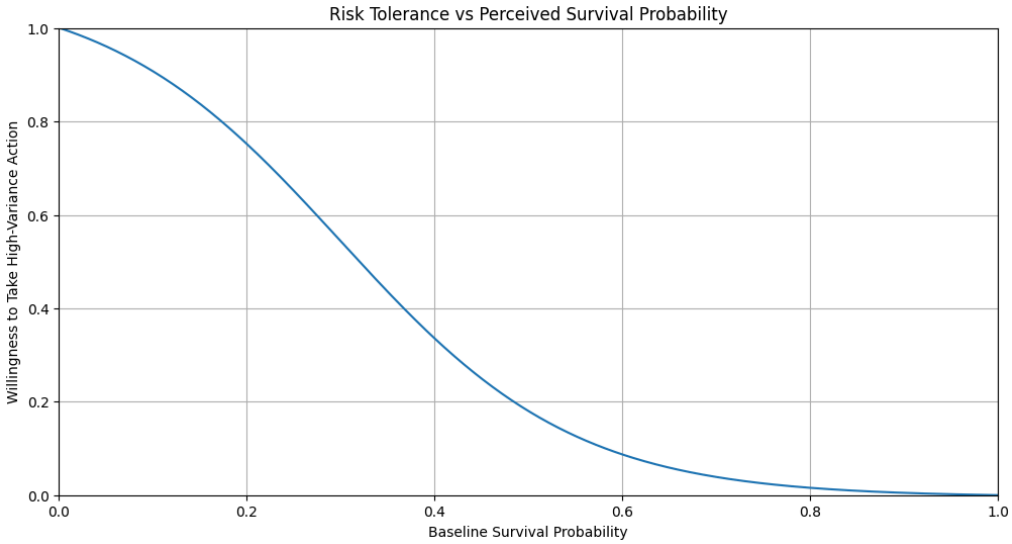
At moderate probability, action becomes selectively rational. There is enough uncertainty to justify engagement, but not high risk. Individuals in this range tend to favor incremental or legible actions; joining established efforts, funding proven approaches, or pursuing interventions with visible near-term impact. Accordingly, respective actions are filtered through considerations of career, reputation, and opportunity cost.

When perceived probability is low, the calculus flips. Inaction now feels equivalent to failure, and high-variance actions become rational. Moreover, even marginal contributions may seem like they can change the odds. Individuals in this regime are more willing to pursue unconventional approaches, start new efforts, or back unproven ideas, because conservative strategies no longer plausibly change the outcome.

Finally, when perceived probability approaches zero, the action space narrows again. At this point, even high-variance efforts within existing structures appear insufficient. Rational actors are thus compelled to pursue regime-level interventions or disengage from conventional action altogether, turning instead to preservation strategies such as cryonics that bypass the need for near-term success.

Probability, then, does not linearly increase or decrease action. It sorts individuals into qualitatively different action modes, each internally rational. The failure arises not from miscalculation, but from unexamined beliefs about where one sits on this probability curve (see *Figure 2*).

Figure 2



Five Illustrative Examples of the Longevity Action Equation in Practice

I: Inaction

Consider a 30-year-old with no direct involvement in aging research. Suppose they estimate a 40% chance that meaningful longevity breakthroughs will occur regardless of their actions. Under this model, **several forces converge toward inaction**:

1. **Marginal Impact Dilution**
   Their personal contribution feels vanishingly small relative to the scale of the problem. Highly capable researchers already exist; progress appears to be underway; the counterfactual impact of "one more person caring" feels negligible.

2. **Deferred Urgency**
   At 30, aging feels distant. The perceived benefits of action accrue decades in the future, while the costs, time, career risk, opportunity cost, are immediate.

3. **Agency Misalignment**
   If the individual lacks domain expertise, capital, or institutional access, their perceived agency is low. The equation penalizes action accordingly.

4. **Narratives of Inevitability**
   Some popular discourse emphasizes exponential technological progress, reinforcing the belief that the current trajectory is unlikely to remain linear, but rather benefit from rapidly compounding returns. Under this regime, breakthroughs feel "inevitable," making action feel optional.

From this vantage point, *doing nothing* is not irrational. It is a locally coherent decision under uncertainty.

II: Start or Join a Company

At first glance, starting a company may seem like an odd way to respond to the problem of aging. The odds are long, the timelines are extended, and the contribution of any single company to population-level outcomes appears uncertain. Yet in practice, company formation is one of the most common actions people choose when they decide to "do something" about longevity.

This choice becomes legible once we examine how individuals implicitly evaluate their options.

Imagine somebody who may or may not believe themselves able to contribute to fundamental research, who is 20-45 years old and who has a survival probability that is low, but not too low.

Importantly the perceived survival probability needs to be low enough that inaction feels equivalent to acceptance of failure, but high enough that the **expected marginal increase in survival probability** attributable to one's actions is meaningful.

For example, a small shift in low baseline odds, moving from 20% to 25%, represents a substantial relative improvement. The same absolute gain at higher baseline odds does not carry the same existential weight. As a result, personal contribution is perceived to be more valuable when baseline expectations are more pessimistic.

Starting a company is one of the few actions that plausibly offers such asymmetric upside. It is a high-variance bet with the potential to meaningfully shift one's personal survival odds, even if the median outcome is failure. When incremental or conservative actions are unlikely to move the needle, **risk-seeking behavior becomes rational**.

There is a further coupling between perceived survival odds and risk tolerance. As the chance of making it declines, the willingness to pursue low-probability, high-impact interventions increases. When the default outcome is death, downside risk loses salience. Founding a company, backing an unproven technology, or pursuing an unconventional approach becomes a defensible response to an otherwise bleak base rate.

Importantly, this logic does not require overconfidence. Nor does it require belief in inevitable success. It requires only the belief that *waiting is worse than acting*, and that certain actions preserve a nontrivial chance of changing the trajectory.

However, this same logic also explains why company formation is often chosen even when it is not the highest-leverage intervention at the system level. Starting a company maximizes personal agency and narrative coherence, but agency is not the same as leverage. Without an explicit Theory of Change, founders may correctly sense that bold action is required, yet still misidentify where such action most effectively shifts global probabilities.

That said, the type of company one chooses to start typically implies one of two Theories of Change, which attract founders with very different Longevity Action Equation profiles.

1. I will solve the problem of aging myself
2. I will catalyze the flywheel to flood the category with money and attention

The first type of longevity company is rarer, and implicitly built around the theory: "I will solve the problem of aging myself." Founders drawn to this path tend to have high perceived agency in a narrow domain, deep technical expertise, a strong prior, or a specific intervention they believe is bottlenecked on execution rather than coordination. Their Longevity Action Equation typically features relatively high agency, moderate-to-low baseline probability, and a belief that marginal technical progress compounds directly into survival odds. The implied Theory of Change is direct and solution-centric: identify the right mechanism, build the right product or therapy, and success follows.

This profile naturally attracts founders who believe that aging is primarily a scientific or engineering problem, and that the limiting factor is the absence of sufficiently focused, competent teams attacking it. When this theory is correct, such companies can be transformative. When it is wrong, when the true bottleneck lies in capital allocation, regulation, coordination, or epistemics, even exceptional execution risks producing local progress without system-level impact.

The second type of longevity company is built around a catalytic Theory of Change: “I will trigger a flywheel that floods the field with capital, talent, and attention.” Founders drawn to this approach typically believe that the baseline chance of making it is very low absent large-scale reallocation, and that direct attempts to “solve aging” are structurally blocked by regulation, incentives, and institutional conservatism.

Because aging is not recognized as a disease, these companies cannot legally pursue it as an indication. Instead, they target *stepping-stone indications*, recognized diseases or conditions whose treatment is expected, at best, to deliver incremental longevity benefits. The implicit strategy is that regulatory approval for such an intervention will enable widespread off-label use, label expansion, or cultural reframing—potentially analogous to drugs like Ozempic, thereby legitimizing longevity as a category and catalyzing a surge of investment and research activity.

This Theory of Change is nonlinear and indirect: success is not primarily defined by the magnitude of the therapy’s longevity effect, but by its ability to unlock a broader flywheel. However, it is also highly fragile. It depends on a long chain of contingencies: clinical success, regulatory approval, permissive off-label prescribing norms, cultural uptake, sustained investor enthusiasm, and eventual translation of incremental progress into genuinely transformative advances.

At each step, scope can narrow and timelines can stretch, sometimes to the point where the original ambition, meaningfully extending healthy human lifespan—no longer meaningfully governs decision-making. Some founders explicitly acknowledge a fallback logic: even if the flywheel fails to materialize, personal financial success would allow them to later fund moonshot research directly. This further weakens the coupling between company success and longevity impact.

More fundamentally, the strategy assumes that incrementalism is a viable path to a problem that may require step-function breakthroughs. If aging is bottlenecked not on capital or attention, but on fundamentally new paradigms, then catalysis via stepping-stone therapies risks producing motion without trajectory.

III: Become a Scientist

Choosing to become a scientist is one of the most common and intuitively serious responses to the problem of aging. If aging is fundamentally a biological problem, then working directly on its mechanisms appears to be the most coherent way to influence the outcome. This path seems

straightforward: acquire the necessary training, join or start a lab, generate new knowledge, and contribute to breakthroughs.

Typically, those who choose this path are 20-45 years old who have a low, but not too low survival probability.

Yet, like starting a company, this action embeds the following implicit Theory of Change:

“Aging is bottlenecked on scientific discovery, and concentrated research effort is the limiting factor.”

Under the Longevity Action Equation, individuals drawn to this path often share a particular profile: Moderate to high agency and confidence in intellectual domains (technical aptitude, academic access, tolerance for delayed gratification), moderate baseline survival probability, enough optimism to believe science can solve the problem, but not enough to justify waiting.

That said, scientists often have higher risk aversion relative to founders, especially since much of academia aligns with societally accepted pathways for contributing to large-scale knowledge production.

However, the scientific path contains its own fragilities.

First, it assumes that aging is bottlenecked on discovery rather than translation, regulation, capital allocation, or institutional coordination. If the limiting factor lies downstream of knowledge generation, additional papers do not necessarily shift outcomes.

Second, it assumes that given existing overall funding and context, directly doing science has a shot to meaningfully impact progress. This may be a siren song since in fields with deep combinatorial complexity, marginal discoveries may accumulate without producing step-function breakthroughs. Moreover, if solving aging requires new paradigms or vast changes in funding, policy or other attention vectors, the marginal utility of another lab may be lower than it appears.

Third, academic and scientific institutions operate under incentive structures that may diverge from longevity’s ultimate goal. Publication, grant cycles, career stability, and prestige metrics shape research agendas. These incentives can bias effort toward tractable subproblems rather than the most leverage-heavy questions. As a result, one may become highly productive within science while only weakly influencing survival probabilities.

This creates an important asymmetry: Becoming a scientist feels maximally serious because it engages the core domain of the problem. Yet **seriousness of domain does not guarantee leverage over the binding constraint.**

The central question is whether an individual’s potential impact within existing scientific structures meaningfully changes the probability distribution of outcomes. Without explicit evaluation of that Theory of Change, becoming a scientist can be a high-leverage commitment or a locally coherent but globally insufficient response.

IV. Meta-Actions => When the Odds Feel Very Low (5%-10%)

When a moderate to high agency individual concludes that their chance of making it under the default trajectory is very low, the goal shifts from picking the *right* solution to **creating conditions where many solutions can be tried in parallel**. From this perspective, the decision to start a single company in longevity is often rational, urgent, and misaligned all at once. At this point, action selection favors moves that increase surface area, optionality, and variance, rather than precision.

Several such actions recur across domains for a reason. Starting a moonshot fund, an institute, or an accelerator are three recurring responses, each premised on the belief that the core uncertainty is not which approach will work, but whether enough approaches can be tried under the right conditions.

These actions, however, are largely inaccessible. A moonshot fund of meaningful size, often on the order of $100M or more, typically requires either substantial personal capital or the credibility, network, and fundraising skill to aggregate it, placing it largely in reach of experienced investors or successful founders. Institutes demand even greater resource intensity and patience: large, multi-year capital commitments, tolerance for unclear milestones, and the ability to attract top talent without the promise of near-term commercialization, making them feasible primarily for philanthropists, governments, or unusually mission-aligned patrons.

Accelerators sit lower on the capital spectrum but higher on the coordination spectrum, requiring deep ecosystem knowledge, founder trust, and sustained operational effort to consistently surface and support credible teams. What unites these approaches is that they are not attempts to be right about a specific solution; they are attempts to reshape the opportunity landscape itself. When time is scarce and uncertainty is high, such meta-actions become some of the few moves that plausibly change the distribution of outcomes rather than merely participate in it.

V. Regime Change: When the Odds Are Diminishingly Small (1% or less)

Only when the perceived odds of making it collapse toward zero does the action set shift again.

At moderate pessimism, people still look for leverage within existing structures; companies, funds, institutes, accelerators. These are ways of increasing variance without directly challenging how resources are allocated. They are aggressive, but still legible.

When the odds feel *diminishingly small*, even this is no longer enough. At that point, the problem is no longer execution, talent, or ideas. It is allocation itself. If the world continues to invest roughly the attention, capital and policy affordances that it does today, the outcome is effectively fixed.

In this regime, rational actors stop trying to add effort at the margins and instead attempt to **re-route the flow of resources wholesale**; capital, talent, attention, legitimacy. The goal is not to produce better projects, but to change what kinds of projects are possible at all.

This is the point at which people pursue overtly political, institutional, or cultural interventions: building movements, reshaping incentives, founding new funding primitives, or attempting to redefine what society treats as urgent, fundable, or legitimate.

These moves are rare not because they require dramatic change, but because they only make sense once all lower-order strategies are exhausted. They are attempts at revolution, undertaken only when the alternative seems indistinguishable from failure.

3: Scoring Rubric

If choosing the best Theory of Change is critical, we need effective methods to evaluate and compare them.

Each Theory of Change may have multi-step and multi-agent causal chains, varied incentive structures, deep scientific context, and diverse dependencies and feedback loops. Some may be overlapping, others may be contradictory or even mutually exclusive.

This complexity calls for a systematic, robust and legible way to assess, compare, and integrate competing Theories of Change.

A useful assessment rubric should achieve three goals:

1. Identify a minimal set of scorable factors to judge a Theory of Change.
2. Define scoring so each factor can be rated on a legible scale.
3. Combine factors and scoring so multiple Theories of Change can easily be compared, surfacing key issues and driving decision-making.

The core principle is that a Theory of Change must both correctly diagnose the binding constraint, and be mechanically capable of acting on it.

As the foundation, nine operational scoring factors were chosen based on whether:

1. Each could be defined in simple terms that a rater can quickly understand and score.
2. A specific factor, if rated at the lowest level, would represent a major deficit for that Theory of Change. This means all factors are plausibly required components of a valid Theory of Change.
3. In combination, the factors enable high confidence decisions within a comprehensive context.

The tenth factor, Rationale and Key Assumptions, was chosen as a primary modifying factor, since a failed rationale will decisively lead to failure, and is the most salient part of a Theory in that it reflects the quality of the explanation of why a certain approach is expected to work.

Nine Scoring Factors for Theories of Change

(on a scale of 1-5)

1. Definition of Success

Definition: *What outcome, for whom, by when, counts as "achieved"?:*

1: Outcome is vague, success criteria unclear, Optimizes a weak proxy.
3: Outcome defined but still fuzzy on timeframe/measurement, Likely bottleneck, but could be displaced quickly.
5: Outcome is operational, measurable, time-bounded, and scoped, clearly targets the binding constraint; would achieve the core goal or very likely unlock desired progress, e.g. increase maximum lifespan for those born in 1970 by 10 years within the next 10 years.

2. Plausibility & Core dependent steps

Definition: *To what degree does the description of the theory's intermediate steps increase or reduce our confidence in its success.*

1: Many steps, unclear necessity, multiple "handwave transitions"
3: Steps mostly clear; a couple jumps remain
5: A few clearly defined, necessary steps; minimal handwaving; dependencies explicit

3. Actionability

Definition: *To what degree are there clear steps that can be taken with existing capabilities?*

1: Even if you believe it, you don't know what to do next
3: Some actions are available
5: Clear next actions are mapped to key hinge steps

4. Uncertainty structure

Definition: *How well do we understand the uncertainty related to this theory?*

1: Mostly unknown unknowns; fog; unbounded uncertainty
3: Major unknowns identified; partial paths to resolve them
5: Uncertainty is legible; dominant unknowns are identifiable and resolvable

5. Timeline

Definition: *How attractive is the timeline suggested by this theory?*

1: You'll know you're wrong only after 10+ years; metrics gameable. Hard to see this demonstrating major success in less than 25 years
3: Some leading indicators; mixed Goodhart risk[4]. Evidence of success expected in 15-20 years.

[4] See Manheim & Garrabrant (2018) for a taxonomy of Goodhart-type failures that arise when proxy metrics are optimized.

5: Conceivable to have dramatic success in less than 10 years. Fast, cheap learning loops; clear leading indicators; Goodhart risk managed.

6. Measurable indicators (leading + lagging):
Definition: *To what degree will we know whether the theory is working? How fast will our feedback cycle be?*

1: Unfalsifiable; “always right” narrative
3: Some updates possible, but evidence ambiguous
5: Clear near-term evidence that would meaningfully update confidence

7. Probability of Success
Definition: *If the Theory of Change works, how likely are we to achieve the goal? A useful score should help compare plausibility and even guess the percentage likelihood of success.*

1: Comparables point against success. No experts view success as possible. Mechanisms rely on handwaving transitions. Many dependent steps with compounding fragility. Requires major new capabilities. No clear leading indicators. Feedback cycles are long. Many core uncertainties, some of which are unbounded. Failure would not be surprising. Resourcing is highly unlikely. Vulnerable to psychological biases.
3: Past success comparables exist, but are partial. Some experts see plausibility. Core mechanisms are articulated. Dependency chain is clear; a few hinge steps dominate risk. Some key required capabilities exist. Beliefs can update over 2-4 years. Resourcing is conceivable. Some vulnerability to psychological biases.
5: Clearly successful past comparables. High expert congruence on feasibility. Mechanisms are well-specified and empirically grounded. Few dependent steps with low fragility. Capabilities largely exist. Leading indicators already present or emerging. Beliefs can update in 1-2 years. Few core uncertainties, and they are resolvable. Low vulnerability to psychological biases.

A 9-feature checklist to help evaluate probabilities

1. Past comparables? Conflict / Neutral / Support
2. Expert congruence on feasibility? Low / Medium / High
3. Mechanistic completeness? Low / Medium / High
4. Dependency depth / Fragility? Low / Medium / High
5. Capability overhang / gap? Low / Medium / High
6. Leading indicator present? No / Partial / Yes
7. Update-sensitivity? Low / Medium / High
8. Number of core uncertainties? Low / Medium / High
9. Portfolio substitutability? Low / Medium / High

8. Partial-Success Yield
Definition: *If this Theory of Change doesn’t fully work, do we still win in meaningful ways?*

1: Failure yields minimal reusable assets or learning. No generalization or dual use.
3: Mixed / Partial positive spillovers, but not transformative (knowledge, tools, talent). Learning doesn't compound. Mechanistic pathways may generalize.
5: Strong / High partial-success yield with durable value. Assets created compound over time (data, institutions, platforms, legitimacy, regulatory harness, manufacturing or delivery capabilities). Commercialization funds iteration even if effects are modest

9. Robustness to Likely Failure Modes
Definition: *How resistant is this theory to core risks that could torpedo it?*

1: One or two likely failure modes would severely or irreversibly compromise the Theory of Change. Many severe failure modes are mapped, and more likely exist.
3: Failure modes are well-identified, some are severe but others are recoverable. Partial or complete failure in one area does not necessarily collapse the entire ToC. Mitigations exist but are incomplete or untested.
5: The most likely failure modes are limited in severity, localized to steps that have partial-yield or redundant features, and/or quickly detectable.

Rationale and Key Assumptions as a Key Modifying Factor
A strategy can be mechanically elegant, measurable, and actionable, yet still fail decisively if its core explanation of why progress is stalled is wrong.

The Rationale encodes the theory's causal diagnosis: what must change, and why this approach will change it. If that diagnosis is flawed, execution strength only compounds effort in the wrong direction. By weighting all structural factors through Rationale, the model reflects the reality that correctness of constraint identification is upstream of execution quality, and that strategic success depends first on being aimed at the right problem.

Calculating a Total Score
Total Strategic Score = Rationale × (Geometric Mean of the 9 Scoring Factors)

Why This Model Is Superior to Simple Averaging
1. It Reflects Reality: Strategies fail primarily because the diagnosis of the bottleneck was wrong, or the mechanism could not reliably act on it. Multiplication captures this interdependence.
2. It Penalizes Fragility: If a theory scores well on execution but poorly on Rationale (e.g., misidentifies the constraint), the total score collapses. Likewise, if rationale is strong but mechanics are weak, the geometric mean exposes that fragility.
3. It Prevents Narrative Over-Optimization: High rhetorical coherence cannot compensate for long, fragile causal chains, 1/5 timeline scores or severe failure modes. The model forces strategic humility.

4: Sample Assessments

Every Theory of Change starts with a goal. Ours is to: "*Live in great health for as long as one wants".*

That said, when a goal seems far, Theories of Change often focus on intermediate goals like:

1. "Amplify attention to drive resourcing and policy"
2. "Unshackle innovation by changing how we innovate"

To start, let's describe some common categories of Theories of Change in longevity.

1. Mobilize the Masses

Here, culture is seen as the constraint. As long as society treats aging as acceptable, resource allocation and policies will be marginal. Even worse, the AIs we develop may be death-accepting and make death-oriented choices. Therefore, mass culture must change.

Grand Strategy:
Step 1: Create a compelling worldview as a sacred norm
Step 2: Embed it in identity, status hierarchies, and community
Step 3: Get the worldview widely adopted
Step 4: Generate (political) pressure
Step 5: Redirect capital, talent, policy
Step 6: Normalize aggressive life extending innovation, including in AI
Step 7: Achieve the goal.

Examples

Start a religion (e.g. Don't Die): Religions are identity-defining ideologies with a successful history of massive societal reallocation of capital and policies (e.g. Christianity, Mormonism).

The key constraint is seen as insufficient civilizational commitment, given the monumental technical innovation required; for commitment to happen, healthy life must be defined as a moral imperative.

Religions have historically reallocated massive capital, altered laws, reshaped norms, and sustained multi-generational coordination. If longevity became a meaning-level commitment rather than a lifestyle optimization, the scale of sacrifice, coordination, and patience could suffice.

Inform and change minds with media (e.g. Jellyfish DAO): Inspirations here include the Green Movement and the movement for LGBTQ rights. One of the core assumptions is that most people are simply not aware they can do anything about aging. Mass movements have repeatedly demonstrated the power to re-rank societal priorities by shifting norms. When enough people see aging as preventable and unacceptable, public pressure shifts elite incentives. Successful media narratives, coalition-building, and visible activism can feed this process, ultimately redirecting funding, policy, and institutional attention.

2. A Key Scientific Breakthrough will Unlock a Tsunami of Progress

Here, proof is the constraint. Without a decisive, undeniable breakthrough showing aging is scientifically tractable (whether in a lab or company), capital, policy, and talent allocation remain cautious and incremental.

Grand Strategy:
Step 1: Achieve a landmark scientific aging breakthrough
Step 2: Shift consensus
Step 3: Large-scale capital inflows and regulatory adaptation follow
Step 4: Accelerate follow-on research and commercialization
Step 5: Achieve the goal.

In one version, the breakthrough is sufficiently transformative on its own to initiate the cascade, much like penicillin, the transistor, or mRNA vaccines reshaped their fields. In another version, the breakthrough functions as a bridge: it may be limited but proves tractability, de-risks the category, and enables the next generation of tools and therapies. In both cases, the core belief is that once the barrier of "is this even possible?" falls, a massive wave of support and progress follows.

Examples

Double a mouse lifespan (e.g. LEV Foundation): The core belief is that mainstream resistance will fall once expert consensus shifts, and that a dramatic, reproducible extension of maximum lifespan in mice using a combination of damage-repair interventions is the minimum viable way to shift it.

Once credible experts publicly agree that aging can be dramatically reversed in mammals, influencers, philanthropists, venture capital, and governments will follow. The breakthrough is designed to collapse uncertainty and legitimize urgency.

Yamanaka Factors: The discovery of cellular reprogramming in 2006 in the Yamanaka lab showed that differentiated adult cells could be reset to a stem-like youthful state.

Shortly afterwards, the transient application of factors, and some evidence that the de-differentiation and rejuvenation aspects could be decoupled, sparked a cascade of capital (e.g. multibillion-dollar commitments to Altos Labs, NewLimit, Retro Biosciences), talent migration, and institutional legitimacy. It remains to be seen whether it becomes aging's "penicillin moment" or remains an early enabling technology; this may depend on its translation into safe organism-level rejuvenation.

3. Activate the Capitalist Flywheel

Here the idea is that a profitable wedge that demonstrates market demand, regulatory navigability, and scalable business models will create a fast spinning flywheel of attention and investment that will bring scientific breakthroughs.

The assumption here is that like semiconductors, personal computing, or GLP-1 drugs, early commercial wins create a positive feedback loop: profits fund iteration, competition accelerates improvement, valuations attract talent, and policymakers adapt to accommodate growth.

Grand Strategy:
Step 1: Commercialize a sufficiently adjacent intervention (even if incremental)
Step 2: Achieve regulatory approval and real-world adoption
Step 3: Generate great investor returns
Step 4: Legitimize longevity as a durable market category
Step 5: Attract top-tier capital, talent, and corporate entrants
Step 6: Expand R&D budgets and regulatory engagement
Step 7: Increase innovation velocity
Step 8: Achieve the goal

Examples

Catalyze the Market (e.g., XPRIZE Healthspan):
This approach aims to catalyze a market by creating a capitalist motive for teams to jump into the category. The logic is that a funding and attention spark is needed, which will compress innovation timelines and demonstrate commercial promise, by encouraging unconventional, diverse approaches.

Further, large moonshot incentive prizes can attract as much as 35x the prize size, and generate media attention that reframes aging as a competitive, solvable engineering challenge.

Achieve Meaningful Progress while Engineering a Commercial Wedge (e.g., Retro Biosciences, New Limit, Altos Labs):
Here, a sufficiently compelling therapy with clear clinical endpoints can unlock massive capital and category legitimacy, without need for an “Aging” indication. The goal is not incremental supplementation or lifestyle optimization, but a serious biotech play aimed at producing a breakthrough-level therapy within an investable timeframe.

Leverage aging biology commercial success with incremental assets (e.g., BioAge):
In this framing, aging biology and longitudinal human datasets can be mined to identify drug targets for age-related diseases with existing regulatory pathways.,

Their role is to prove that aging biology can be translated into profitable therapeutics, even if they do not increase lifespan. Once capital markets internalize that insight, the flywheel strengthens: companies who show some proof themselves can take bolder bets, as more startups form, more capital flows, and more ambitious age-targeting strategies become fundable.

Prove the Category in a Permissive Market (e.g., Loyal):
This approach assumes the fastest way to activate the longevity flywheel is to win in a **closely related but more navigable market first**. Instead of targeting human aging directly, begin with companion animals, where regulatory pathways are more flexible, time-to-outcome is shorter,

and consumer willingness to pay is high.

Dogs provide faster feedback loops, clearer endpoints, and a large addressable market of motivated owners. If aging-targeted therapies can generate significant revenue and regulatory precedent in animals, the perception of risk in human longevity biotech declines. The flywheel activates not because aging is solved, but because aging is shown to be **commercially tractable**.

4. Restructure How Innovation Happens

Here, the constraints are inefficient pipelines, conservative regulatory frameworks, misaligned incentives, and institutional inertia. We already spend hundreds of billions of dollars per year on biomedical R&D, yet progress on aging remains incremental.

Some versions focus on regulatory restructuring, arguing that agencies are the primary bottleneck and that reform, parallel-track pathways, or alternative jurisdictions are needed (e.g., Prospera). In others it's the way science is funded and done. Change the rules and process of innovation, and science will move faster.

Grand Strategy:
Step 1: Redesign how research is conducted and therapies are tested
Step 2: Reform regulatory pathways in existing or alternative jurisdictions
Step 3: Accelerate iteration and reduce time-to-human data
Step 4: Enable higher-variance experimentation
Step 5: Surface effective longevity interventions faster
Step 6: Compound innovation.
Step 7: Achieve the goal

Examples

Build Parallel Innovation Jurisdictions (e.g., Infinita):
The idea is to escape regulatory and jurisdictional inertia which slow iteration and suppress high-variance experimentation. Rather than waiting for reform, the strategy is to create alternative physical and legal environments where biomedical innovation can move faster.

If innovation-friendly zones prove effective, they can outcompete or catalyze reform in existing systems, increasing the global rate of innovation.

Decentralize and Rewire Scientific Funding (e.g., DeSci DAOs):
This version assumes the bottleneck lies in how science is funded, owned, and coordinated. Traditional grant systems are slow, risk-averse, and dominated by institutional gatekeepers. Intellectual property structures often misalign incentives between researchers, funders, and patients. The solution is to redesign the financial and governance infrastructure of science itself.

This will fund more research and rewire the capital formation and governance layer of biomedical science. With permissionless, agile and aligned funding, the rate of meaningful discovery increases even without spending increases.

5. Build Scientific Moonshots to Solve the Problem Directly

Here, what's seen as missing is focused attention by the right teams on discrete technical paths. Existing resourcing and policy are sufficient or close to it.

In one version, AI is the moonshot: the combinatorial complexity of biology exceeds human cognition, so accelerating advanced AI systems becomes the best path. In others, the bet is on elite human insight into biology: a small number of highly focused individuals or teams can discover what diffuse systems have not.

In both cases: the assumption is that the ingredients for finding a scientific solution exist and only attention and reorganization are needed. The specific scientific approaches differ, and may include approaches as diverse as body replacement, cryopreservation, epigenetic reprogramming and more.

Grand Strategy:
Step 1: Concentrate exceptional talent, technology and resources to solve aging
Step 2: Pursue paradigm-shifting interventions
Step 3: Achieve a breakthrough that meaningfully reverses or halts aging

Examples

Focus entirely on AI (e.g. Isomorphic Labs): A sufficiently advanced AI system could navigate the combinatorial search space of biology. The binding constraint is cognitive capacity, and accelerating AI development indirectly accelerates longevity. Given the tailwinds for AI investment overall, this is also a way to access resources on a potentially faster timeline.

Do the science and find the solution myself (e.g. Life Biosciences, Tomorrow Biostasis): Aging is a tractable scientific problem, and progress is bottlenecked by the absence of sufficiently capable individuals pursuing the right hypotheses with enough focus and persistence. The correct response is to personally become one of those individuals and discover the solution directly.

A core idea here is that exceptional individuals or small-team scientific insight can outperform large, diffuse systems, and the primary failure to date has been insufficiently concentrated talent working on the most promising ideas.

6. Surgically Drive Society's Attention by Finding Key Leverage Points

In this framing, leverage beats scale. Society does not move because "everyone" is convinced. It moves when key decision-makers re-rank priorities.

The goal is therefore not to mobilize millions, but to engage and impact the nodes that control capital allocation, regulation, narrative framing, and institutional direction. A small shift in elite perception can redirect trillions of dollars and reshape entire sectors.

Grand Strategy:
Step 1: Identify key political, financial, regulatory, and narrative nodes
Step 2: Build targeted relationships and strategic alignment
Step 3: Coordinate activity
Step 4: Reallocate budgets, policies, and incentives toward longevity
Step 5: Trigger downstream capital flows and institutional adaptation
Step 6: Accelerate global longevity development

Examples

Traditional Lobbying (e.g., A4LI):
Regulatory definitions and government priorities are the primary bottleneck. Aging is not officially recognized as a treatable indication in most jurisdictions, limiting funding, trial design, and drug approval pathways. Also, much basic science is still missing, and that can only be funded by public capital, since the commercial motive is too weak.

In this model, progress depends on incremental but compounding policy adjustments. Success is measured in regulatory language shifts, budget allocations, and formal recognition that unlocks new experimentation. Forging alliances with key industry players and coordinating their efforts with the government is key.

Influence Networks (e.g., Vitalism):
This version assumes the most powerful lever is not merely formal policy, but elite coordination.

Influence networks are coordinated webs of high-agency individuals, founders, companies, investors, scientists, philanthropists, policymakers, and align them around a shared strategic narrative. Lobbying and government engagement is a necessary but insufficient factor.

Past precedent (Koch Brothers, Soros Network) shows that hundreds of resource-rich, high-leverage stakeholders can move large-scale resource and attention reallocation without relying on mass adoption.

Comparative Scoring

The scores that follow are not meant to be definitive rankings, nor claim objectivity. They are simply an illustrative walkthrough of how the rubric can be applied in practice, using my own sensibilities about what constitutes a strong binding-constraint diagnosis, credible mechanics, and fast, falsifiable learning loops.

Reasonable people will weigh uncertainties, timelines, and failure modes differently, and different priors about what is actually bottlenecking longevity will produce different scores.

Treat these examples as a template: apply the same factors to the Theories of Change you personally find most compelling, make your assumptions explicit, and produce your own scores. Remember, the value of the exercise is less the exact numbers than the clarity it forces about *why* you believe a given strategy will work and *what evidence would change your mind*.

I: Mobilize the masses via Religion

(e.g. Don't Die)

1. Definition of Success: 2/5

The outcome is clear at a high level: make healthy indefinite lifespan a moral imperative such that it drives attention and resourcing. However, it is not time-bounded and does not specify measurable thresholds (e.g., % population adoption, funding shifts, regulatory change).

For 5 points, it would need something like: "Achieve 60% polling support in treating aging as a disease, measurable increases in AI's appreciation of life as sacred, $50B/year reallocation, and regulatory recognition within 5 years."

2. Plausibility (Core dependent steps): 2/5

The causal chain is long and fragile.

Identity-level ideology → Mass adoption → Political gravity → Capital and attention reallocation → Innovation acceleration → Lifespan extension.

Each individual step is plausible in isolation, and historical movements (religions, civil rights, environmentalism) demonstrate that large-scale norm shifts are possible. However, there are multiple handwave transitions, especially between cultural adoption and accelerated biomedical innovation. The dependency chain is long and compounding fragility is high.

To reach a 5, the hinge transitions (e.g. how moral adoption reliably converts to favorable life extending outcomes) would need clearer empirical anchors and fewer dependent steps.

Actionability: 3/5

There are actionable first steps: narrative building, media amplification, community formation, ritual and identity construction, influencer engagement, and coalition building. These are concrete and executable.

It's possible to start a movement, but scaling it to civilizational adoption relies on high-variance steps.

To reach a 5, clearer pathways are needed from early traction to durable mass adoption.

4. Rationale and key assumptions: 4/5

The key assumptions are explicit and historically grounded: Cultural legitimacy drives political and financial allocation. Aging is under-resourced because it lacks sacred or moral status. AI and future institutions may inherit prevailing cultural norms.

To reach a 5, assumptions would need tighter decomposition, ranked hinge steps, and clearer evidence that cultural shift can drive deep biomedical acceleration.

5. Uncertainty Structure: 2/5
Major uncertainties are identifiable but large; cultural systems are nonlinear and difficult to forecast.

1. What level of adoption is sufficient?
2. Does moral urgency translate into durable R&D acceleration?
3. Does movement identity increase epistemic distortion?
4. How does the actual development of AI intersect with Don't Die

To reach a 5, key uncertainties would need clearer modeling and faster resolution mechanisms.

6. Timeline: 2/5
Historical comparables typically unfolded over 30-100 years. While early traction can be measured, success may take decades to manifest. Moreover, metrics for cultural impact are susceptible to Goodhart's Law. Learning loops are slow and noisy.

Even if AI is going to radically change the world within 5 years, it's not clear that a cultural-change strategy could move faster-than historical precedent.

To reach a 5, there would need to be faster, clearer signals that show long-term change.

7. Measurable Indicators (Leading + Lagging): 3/5
Some indicators exist.

Growth in committed community size; Media coverage volume and tone; Polling shifts on aging as disease; Funding and policy language changes; Influencer alignment. However, causality remains ambiguous. It is possible to construct narratives that the theory is "working" even when structural acceleration is minimal.

To reach a 5, near-term indicators would need stronger predictive power for downstream biomedical progress.

8. Probability of Success: 2/5
Moral movements have historically driven massive shifts. However, few clear precedents exist where a religion-like movement directly accelerated complex frontier biomedical innovation.

The mechanism chain is long, fragile, and highly dependent on coordination quality. Failure would not be surprising; success would be extraordinary.

To reach a 5, there would need to be clearer historical analogues or clearing chain of operationalized ideas showing direct translation from identity movement to sustained deep-tech acceleration.

9. Partial-Success Yield: 4/5
This is a place where Don't Die shines. Even partial success creates durable assets: Moving the Overton Window; Increased cultural legitimacy; Attention spillovers; Potential funding increases.

It does not reach 5 because partial success may increase noise or polarization without proportionate biomedical acceleration.

10. Robustness to Likely Failure Modes: 2/5
Several severe failure modes exist.

Drift into pseudoscientific posture; Key-Man dependency; Polarization framing longevity as ideological; Institutional distrust increasing regulatory resistance.

Many of these failure modes are severe and could undo the entire effort; available mitigations are limited.

To reach a 5, failure modes would need to be localized, detectable early, and recoverable without collapsing the movement's legitimacy.

Conclusion

Under this framework, Don't Die presents a bold and conceptually coherent attempt to shift the cultural foundations underlying capital, policy, and institutional priorities. Its strength lies in the clarity of its moral thesis and the potential durability of its partial-success effects: even modest traction could expand legitimacy, move the Overton window, and increase attention toward longevity. However, its causal chain is long and fragile, its feedback loops are slow, and its core uncertainties are large and difficult to resolve quickly. While it carries extraordinary upside if cultural legitimacy is truly the binding constraint, it remains a high-variance strategy whose success depends on multiple compounding transitions that are hard to tightly control.

Don't Die yields a geometric mean of 2.36 and a Rationale score of 4.

Final Score: 4 × 2.36 = 9.46

This reflects a strategy with a relatively strong binding-constraint thesis, that cultural legitimacy is upstream of capital and policy allocation, but weak mechanics. The causal chain is long, uncertainties are large, timelines are slow, and robustness is limited. Although no structural 1/5 penalties apply, its operational fragility materially suppresses the final score. Don't Die therefore emerges as a high-upside, high-variance strategy whose success depends heavily on whether cultural legitimacy is in fact the dominant constraint.

II: Activate the Capitalist Flywheel (A Mix of Progress and a Commercial Wedge)

(e.g., Retro Biosciences, New Limit, Altos Labs)

Definition of Success: 3/5

The success condition is: develop a clinical intervention targeting aging biology (can be a stepping-stone indication) and show that longevity biotech is investable, approvable, and produces commercial returns.

It does not specify what level of revenue, market cap, or capital inflow constitutes flywheel activation.

To reach 5 it would need measurable triggers such as first FDA approval for an aging-linked mechanism, $X billion in follow-on capital within Y years, and entry of Z top-tier biotech firms into the space.

Plausibility (Core dependent steps): 2/5

The causal chain is clear and historically grounded: Developing a credible therapy that generates commercial returns will attract capital and talent, accelerating category growth. This mirrors how GLP-1 drugs, mRNA platforms, and oncology biotech scaled.

That said, without dramatic change in how we innovate, each “round” of incremental innovation will take a decade to reach people, even if an increasing number of programs exist. The dominant risk is thus incrementalism, whereby a capital flywheel will deliver innovations too slowly to save anybody. This is highly likely if one believes we are too far down the technology tree and need more fundamental science.

It would reach a 5 if the translation from early wedge to dramatic systemic acceleration were more convincing.

Actionability: 3/5

The path is operational and legible: Raise venture capital to build and demonstrate success via biotechs targeting tractable mechanisms. Execution of building the biotech is difficult, but not ambiguous.

However, the relationship between success at getting a drug to market and activating the capitalist flywheel is vague and mostly out of the company’s direct control, making it fragile.

It would reach a 5 if the expected impact and path to impact on the expected flywheel of attention, capital and policy were clearly articulated.

Rationale and Key Assumptions: 2/5

Assumptions are explicit: the binding constraint is market validation, not culture; capital flows follow demonstrated returns; once longevity becomes an investable category, competition and iteration accelerate progress. These assumptions are historically grounded and testable.

However, the model assumes incremental wins will somehow either compound or unlock more meaningful efforts that achieve deep age modification rather than stall at profitable but shallow outcomes.

The rationale supporting the identification of the binding constraint is weak.

It would reach a 5 with stronger evidence that market-driven incrementalism reliably produces paradigm-level breakthroughs.

Uncertainty Structure: 3/5
Major uncertainties are identifiable: Will early therapies produce strong enough effect sizes? Will regulatory pathways expand to aging itself? Will capital continue flowing after first exits? Will incentives bias toward incrementalism over deeper repair?

Can a strategy that doesn't explicitly focus on the goal of ending aging, but rather on a variety of proxy goals, even work?

It would reach a 5 with clearer predictive models of capital cascade dynamics in longevity biotech.

Timeline: 1/5
Clinical biotech timelines are often measured in decades, and venture cycles, regulatory approvals, and exits produce feedback loops within 5-10 years. When each step is incremental and activating a flywheel is the core bet, these cycles are simply too slow.

A dramatically different path to iteration would have to be true for it to reach 5 points.

Measurable Indicators (Leading + Lagging): 4/5
Indicators are clean and market-based: venture capital inflows; valuations and exits; regulatory milestones; clinical endpoints; entry of major pharma players. These are concrete and difficult to fake. Market signals update quickly.

It would reach a 5 if leading indicators reliably predicted systemic aging acceleration rather than isolated company success.

Probability of Success: 3/5
There are strong historical precedents where commercial validation transformed entire sectors (e.g., semiconductors, oncology biotech, GLP-1 drugs). The mechanism is robust within capitalist systems.

However, it assumes aging biology will yield sufficiently strong, monetizable effects to sustain a long-term flywheel. Failure would not be shocking, but success is structurally plausible.

To reach a 5, there would need to be better evidence that such efforts will not get stuck in a sick-care, incrementalist bog.

Partial-Success Yield: 3/5
Even partial success creates durable infrastructure: platform technologies; clinical trial networks; regulatory precedent; talent pipelines; manufacturing capacity; data assets. These assets compound and remain reusable even if individual therapies are modest. That said, given the sick-care focus required by existing regulatory context, it's questionable how much utility the assets will be for other aging work.

It does not reach a 5 because incremental commercialization may plateau without solving core aging mechanisms.

Robustness to Likely Failure Modes: 2/5
Failure modes include modest clinical effects failing to generalize; investor enthusiasm collapsing after early disappointments; regulatory friction persisting; incentives biasing toward short-term profitability over deep repair.
Mission drift as a result of founder loss of control.

There is deep vulnerability to the vicissitudes and aperture of commercial interest.

It would reach a 5 if incentives for long-horizon repair were structurally embedded.

Conclusion

The Capitalist Flywheel strategy is operationally legible and grounded in historical examples where commercial validation unlocked entire industries. It offers clean market-based indicators, concrete execution steps, and the potential to build durable infrastructure even in partial success. However, its central vulnerability lies in the assumption that market validation is the primary bottleneck to solving aging. If deeper scientific or translational constraints dominate, commercial compounding may optimize for incremental sick-care rather than paradigm-level repair. As framed, this approach provides a credible path to sustained biotech growth, but not necessarily to solving aging itself.

The Capitalist Flywheel earns a geometric mean of 2.36 and a Rationale score of 2.

Final Score: 2 × 2.51 = 5.01

Operationally, the strategy is legible and measurable, with strong market-based indicators and historical analogues. However, the structural fragility imposed by long biotech timelines and a weak binding-constraint diagnosis sharply depress its score under the multiplicative model. If market validation is not the true bottleneck, compounding commercial activity risks optimizing for incrementalism rather than solving aging. The scoring framework therefore treats this as a structurally coherent but strategically misaligned approach.

III: Build Scientific Moonshots to Solve the Problem Directly

(e.g. Life Biosciences)

Definition of Success: 5/5

Success is defined as directly discovering and validating an intervention that meaningfully halts, reverses, or repairs aging in humans.

Plausibility (Core dependent steps): 1/5

The causal chain is comparatively short: assemble a highly focused team → pursue the right mechanistic hypotheses → generate decisive experimental proof → translate into a therapy.

This avoids many sociological steps. However, it concentrates risk into a small number of scientific bets. The plausibility depends heavily on whether aging has “keyhole” mechanisms that a small team can crack faster than the broader field. Given very few teams are working on this, the odds seem extremely slim.

It would reach a 5 if the approach were anchored to a clearly identified, highly promising mechanistic wedge with strong prior evidence.

Actionability: 4/5

The steps are clear and executable: recruit exceptional scientists and engineers; pick a mechanistic thesis; build lab and experimental infrastructure; run tight iteration loops; publish and replicate; design translational plan. For a high-agency actor, this is directly actionable. It does not reach 5 because the most important hinge step—choosing the correct mechanism and executing breakthrough science—cannot be reliably proceduralized.

Rationale and Key Assumptions: 3/5

The core assumptions are: aging is fundamentally a solvable scientific problem; progress is bottlenecked by insufficiently concentrated talent and focus; a small team can outperform diffuse institutions, even with a relatively tiny amount of capital.

These assumptions are plausible and partially supported by historical cases where small groups delivered breakthroughs. However, they can be fragile if the true bottleneck is translation, regulation, or fundamental science (i.e. we are too low on the technology tree). The assumptions also fail if the efforts can’t clear a minimal resourcing hurdle to be tested.

It would reach a 5 with clearer decomposition of the scientific thesis, ranked uncertainties, and explicit reasoning for why concentration beats ecosystem-scale exploration.

Uncertainty Structure: 3/5

The key uncertainties are legible: is the chosen mechanistic thesis correct; will interventions show strong effect sizes; will results replicate; will safety be manageable; will translation stall. These are bounded scientific uncertainties rather than cultural fog, but they remain substantial.

It would reach a 5 if the approach had clear near-term experiments that sharply update the central thesis.

Timeline: 5/5
If the thesis is correct, progress can be rapid. Small-team iteration is fast, and early mechanistic proof can appear within 1-3 years.

Measurable Indicators (Leading + Lagging): 4/5
Indicators are relatively clean: mechanistic validation; biomarker shifts; functional improvements in animal models; replication; movement into IND-enabling studies; human trial initiation. These are falsifiable and update confidence over time.

It does not reach 5 because biomarkers can be Goodharted and early model wins may not predict human outcomes.

Probability of Success: 1/5
Breakthroughs driven by small teams have historical precedent, and concentrated execution can outperform fragmented ecosystems. However, aging may require multi-factor solutions and long translation cycles. Success is plausible but far from assured; failure would not be surprising. This is a classic high-upside scientific moonshot with moderate probability conditional on selecting the right thesis.

Partial-Success Yield: 3/5
Even if “solving aging” is not achieved, partial success creates durable assets: new mechanistic knowledge; datasets; tools; assays; trained talent; experimental platforms; translational infrastructure; potential spinoffs targeting specific diseases.

Many intermediate outputs are scientifically valuable and can compound but few teams are working on this with so few resources that it’s hard to imagine a very substantial partial yield.

It does not reach 5 because partial mechanistic insight may not translate into meaningful human lifespan extension without additional ecosystem changes.

Robustness to Likely Failure Modes: 2/5
At a small scale of teams and initiatives, the primary risks are concentrated: a wrong mechanistic bet, irreproducibility, safety issues, or failure to translate can collapse each effort. When you have few of them, your entire strategy is fragile.

It would reach a 5 if it were structured as a a large scale, well funded effort with a diversified portfolio of independent mechanistic bets with redundancy and rapid kill criteria.

Conclusion

The Scientific Moonshot approach is maximally direct: it targets the core biological problem without relying on cultural or market intermediaries. It benefits from fast learning loops, clear experimental indicators, and a sharply defined definition of success. Yet this directness comes with concentrated fragility. The strategy hinges on selecting the correct mechanistic thesis and executing breakthrough science within a narrow band of uncertainty. If that thesis is wrong, the entire effort collapses. This makes the approach a classic high-upside, high-risk bet—capable of rapid transformative success, but deeply sensitive to misdiagnosis of the true scientific bottleneck.

Building Scientific Moonshots yields a geometric mean of 2.68 and a Rationale score of 3.

Final Score: 3 × 2.68 = 8.05

This approach benefits from clarity of success definition (5/5), fast iteration potential (Timeline 5/5), and strong measurability. However, the multiplicative model exposes its concentrated fragility: two structural 1/5 scores trigger heavy penalties, and the binding-constraint diagnosis remains weakly supported. If aging is combinatorial or ecosystem-bottlenecked rather than mechanistically tractable via a narrow wedge, the entire effort collapses. The model therefore classifies this as a high-directness but high-collapse-risk strategy.

IV: Build Influence Networks to Redirect Society's Attention

(e.g., Vitalism)

Definition of Success: 4/5

Success is defined as shifting the priorities and resource allocation decisions of a small number of high-leverage actors (funders, founders, policymakers, regulators, media gatekeepers) such that longevity becomes a top-tier, institutionally legitimate objective. The goal is not mass adoption, but elite re-ranking that triggers cascading downstream effects. However, success criteria are often under-specified: what level of capital reallocation, policy change, regulatory adaptation, or institutional commitment constitutes "activation"?

It would reach a 5 if it specified concrete thresholds such as $X deployed into targeted longevity strategies, Y policy or regulatory wins, Z major institutions committing to longevity programs, within a defined timeframe.

Plausibility (Core dependent steps): 4/5

The causal chain is relatively tight and historically common in how societies actually allocate resources: build a dense network of high-agency actors → align them around a coherent theory of change → coordinate capital, talent, and attention → influence institutions and policy → accelerate innovation and deployment. Compared to mass movements, the chain has fewer steps and relies on known dynamics of elite coordination. The main uncertainty is whether network influence can overcome deeper bottlenecks (scientific tractability, regulation, translation).

It would reach a 5 if the network had a clearer path from elite alignment to measurable biomedical acceleration.

Actionability: 4/5

The steps are concrete: identify leverage nodes; recruit high-credibility members; curate convenings; produce memetic narratives and decision tools; broker collaborations; coordinate funding; build institutional partnerships; shape policy agendas. For high-agency organizers, this is directly actionable. It does not reach 5 because some hinge outcomes (elite conversion, sustained alignment, institutional capture) remain inherently high variance.

Rationale and Key Assumptions: 4/5

The key assumptions are explicit and plausible: society moves when decision-makers move; resources are controlled by small groups; legitimacy and coordination are binding constraints; shifting elite priors can redirect capital and policy far faster than mass persuasion; once re-ranked, the broader ecosystem follows. These assumptions are historically grounded in how major national and technological priorities have been set. It would reach a 5 with stronger evidence that elite coordination in longevity reliably translates into faster therapeutic progress rather than just more activity.

Uncertainty Structure: 3/5

Uncertainties are identifiable: Which leverage nodes matter most? Can the network maintain alignment over time? Can it avoid fragmentation into competing sub-tribes? Will institutional gatekeepers respond? Will scientific bottlenecks dominate regardless? These are bounded uncertainties and can be tested via experiments in convening, fundraising, and policy. However, causal attribution remains noisy.

Timeline: 3/5
Feedback loops can be relatively fast: network formation, convening outcomes, capital commitments, partnerships, and policy engagement can be observed within months to a couple years. This is much faster than cultural transformation and can meaningfully compress coordination timelines. It does not reach 5 because biomedical outcomes still take longer and network effects may lag.

Measurable Indicators (Leading + Lagging): 3/5
Leading indicators are concrete: growth in high-agency membership; capital committed; partnerships formed; policy engagements; institutional programs launched; media influence achieved; repeated convening attendance and retention. Lagging indicators include funded projects reaching trials or regulatory milestones.

It does not reach 5 because network “heat” can be Goodharted and may not map cleanly to biomedical acceleration without careful metric design.

Probability of Success: 2/5
Elite networks have repeatedly shaped real-world outcomes, so the mechanism is plausible. However, longevity is unusually complex, and influence alone cannot substitute for tractable science and workable regulatory pathways. Success is conceivable and resourcing is plausible, but the conversion from attention to solved biology is not guaranteed.

Failure would not be surprising.

Partial-Success Yield: 3/5
Even partial success creates durable assets: a high-agency coalition; repeat convening infrastructure; shared language and strategy; funding pipelines; partnerships between scientists and capital; policy relationships; legitimacy for the category. These assets compound and can support multiple downstream theories of change.

It does not reach 5 because partial success may produce coordination without sufficient scientific breakthroughs.

Robustness to Likely Failure Modes: 2/5
Failure modes include: network becoming a social club without hard execution; fragmentation into factions; reputational risk; capture by low-quality actors; misallocation toward fashionable but low-leverage projects. Very few people resonate with this strategy, making it brittle.

These risks are real but partly mitigable through governance, rigorous epistemics, and clear prioritization.

It would reach a 5 if redundancy and portfolio logic reduced dependence on a small set of personalities or donors.

Conclusion

The Influence Networks strategy focuses on shifting priorities among high-agency actors who control capital, policy, and institutional legitimacy. Its causal chain is tighter than mass cultural mobilization and less scientifically concentrated than moonshots, offering relatively faster coordination feedback loops. Even partial success builds durable assets: aligned coalitions, shared language, funding pipelines, and institutional pathways. However, influence alone cannot overcome fundamental scientific or regulatory constraints. Its ultimate effectiveness depends on whether coordination and legitimacy among elite decision-makers are the true binding constraints. If they are, this approach can be catalytic; if not, it risks producing alignment without proportional biomedical acceleration.

Influence Networks receives a geometric mean of 3.11 and a Rationale score of 4.

Final Score: 4 × 3.11 = 12.44

This is the highest-scoring strategy under the multiplicative model. It combines a strong binding-constraint thesis, that coordination and legitimacy among high-agency actors are central bottlenecks, with solid operational coherence and no catastrophic fragility. While it cannot substitute for scientific tractability or regulatory feasibility, its tighter causal chain and faster coordination feedback loops strengthen its structural score. Under a model that prioritizes constraint correctness multiplied by mechanical strength, this approach emerges as the most strategically aligned.

Comparative Analysis: My Personal Takeaways

The following are only my conclusions, based on my biases and my scoring. They are meant to showcase the utility of the scoring model when it comes to understanding an individual's decisions.

Having scored 4 major Theories of Change, my choice to focus my energies on Vitalism Foundation becomes more legible:

I score Influence Networks (Vitalism) highest at 12.44, versus Don't Die at 9.46, the Scientific Moonshot at 8.05, and the Capitalist Flywheel at 5.01.

Read together, these results capture the core reason I chose to pursue the Vitalist strategy and why I see Don't Die as the closest next bet, despite its low geometric mean. My analysis is that the dominant failure mode in longevity has not been "nobody tried" but rather massive effort spread across strategies whose mechanics are either slow, fragile, or mis-aimed at the binding constraint.

Tens of billions of dollars have already been spent on biomedical aging research over decades, and in the past five years alone, many billions more have flowed into commercial longevity biotech, Altos Labs (~$3B), Retro Biosciences (~$1B+), NewLimit (hundreds of millions), Calico (billions over a decade), along with hundreds of venture-backed companies and academic programs globally. In parallel, cumulative NIH aging-related research alone has amounted to tens of billions.

In other words, the Scientific Breakthrough and Capitalist Flywheel theories are not untested hypotheses. They have absorbed enormous time and treasure and shown legitimate triumphs. While for the purpose of actually achieving conclusive breakthroughs, that funding has still been a pittance, it is not a pittance when it comes to the Theory of Change promising the "tsunami cascade" those theories predict.

In other words, compared to 30 years ago, it's obvious today that aging is tractable. So where's the tsunami?

Meanwhile, the moonshot approach remains fragile and the flywheel approach risks incrementalism under long translational timelines. Mass cultural mobilization can be directionally right but is long-horizon and hard to falsify; so far no such effort has shown a major success.

In contrast, the influence-network approach scores best not because it "solves biology," but because it targets a constraint that sits upstream and may be modifiable with few resources - elite coordination; it thus offers a tighter causal chain, faster coordination feedback loops, and meaningful partial-success yield (coalitions, shared strategy, funding pipelines, institutional pathways).

5: Conclusion

This chapter was written as an urgent call to action.

We are not Marcus Aurelius attempting to build a jet in the wrong millennium. We have CRISPR, reprogramming, protein folding breakthroughs, AI-guided drug discovery, longitudinal datasets, cloned mammals with normal lifespans, and lifespan extension in model organisms. The question is no longer whether aging is tractable in principle, it is whether we will act intelligently enough to make it tractable in time.

History and logic show that when the stakes are high and the goal is system-level, success or failure are defined at the level of Grand Strategy and Theory of Change.

Kennedy did not ask, “Is Apollo expensive?”, he asked, “Is there a lower-risk path to a decisive geopolitical win under radical uncertainty?”

HIV activists did not accept slow science as inevitable. They redefined the binding constraint as speed and forced institutional change.

The lesson is also that execution cannot rescue a flawed Theory of Change. Yet in longevity, we continue to behave as if action selection precedes causal reasoning. We start companies, join labs, fund incremental programs, and optimize legible outputs.

All while avoiding the uncomfortable question:

“What is the explanation for why a specific Grand Strategy will unlock the binding constraint?”

Though it requires abstract reasoning, this concern is not abstract; executing brilliantly under the wrong Theory of Change will likely lead to tragic failure.

Thankfully the antidote is simple; be explicit. Write your Theory of Change, score it, expose its weak links, and identify what would falsify it. Then commit and set measurable thresholds you can re-evaluate regularly.

Choose thoughtfully. Act deliberately. And re-evaluate relentlessly.

Lives depend on it.

6: Bibliography

Theory of Change, evaluation, and causal frameworks

- Anderson, A.A. (2005) *The Community Builder's Approach to Theory of Change: A Practical Guide to Theory Development*. New York: The Aspen Institute.

- Chen, H.-T. (1990) *Theory-Driven Evaluations*. Newbury Park, CA: Sage.

- Funnell, S.C. and Rogers, P.J. (2011) *Purposeful Program Theory: Effective Use of Theories of Change and Logic Models*. San Francisco, CA: Jossey-Bass.

- Kellogg Foundation (W.K.) (2004) *Logic Model Development Guide*. Battle Creek, MI: W.K. Kellogg Foundation.

- Patton, M.Q. (1997) *Utilization-Focused Evaluation*. 3rd edn. Thousand Oaks, CA: Sage.

- Weiss, C.H. (1995) 'Nothing as practical as good theory: exploring theory-based evaluation for comprehensive community initiatives for children and families', in Connell, J.P., Kubisch, A.C., Schorr, L.B. and Weiss, C.H. (eds) *New Approaches to Evaluating Community Initiatives*. Washington, DC: Aspen Institute.

Strategy, grand strategy, and organizational choice under uncertainty

- Clausewitz, C. von (1832/1984) *On War*. Edited and translated by M. Howard and P. Paret. Princeton, NJ: Princeton University Press.

- Freedman, L. (2013) *Strategy: A History*. Oxford: Oxford University Press.

- Liddell Hart, B.H. (1954/1991) *Strategy*. 2nd edn. New York: Meridian.

- Mintzberg, H. (1994) *The Rise and Fall of Strategic Planning*. New York: Free Press.

- Mintzberg, H. and Waters, J.A. (1985) 'Of strategies, deliberate and emergent', *Strategic Management Journal*, 6(3), pp. 257–272.

- Paine, S.C.M. (2017) *The Japanese Empire: Grand Strategy from the Meiji Restoration to the Pacific War*. Cambridge: Cambridge University Press.

- Rumelt, R.P. (2011) *Good Strategy/Bad Strategy: The Difference and Why It Matters*. New York: Crown Business.

Decision-making, bias, incentives, and metric failure modes

- Ainslie, G. (1975) ‘Specious reward: a behavioral theory of impulsiveness and impulse control’, *Psychological Bulletin*, 82(4), pp. 463–496.

- Campbell, D.T. (1979) ‘Assessing the impact of planned social change’, *Evaluation and Program Planning*, 2(1), pp. 67–90.

- DiMaggio, P.J. and Powell, W.W. (1983) ‘The iron cage revisited: institutional isomorphism and collective rationality in organizational fields’, *American Sociological Review*, 48(2), pp. 147–160.

- Goodhart, C.A.E. (1975) ‘Problems of monetary management: the U.K. experience’, in *Papers in Monetary Economics*. Sydney: Reserve Bank of Australia.

- Kahneman, D. and Tversky, A. (1979) ‘Prospect theory: an analysis of decision under risk’, *Econometrica*, 47(2), pp. 263–291.

- Laibson, D. (1997) ‘Golden eggs and hyperbolic discounting’, *Quarterly Journal of Economics*, 112(2), pp. 443–477.

- Sweller, J. (1988) ‘Cognitive load during problem solving: effects on learning’, *Cognitive Science*, 12(2), pp. 257–285.

- Tversky, A. and Kahneman, D. (1973) ‘Availability: a heuristic for judging frequency and probability’, *Cognitive Psychology*, 5(2), pp. 207–232.

- Tversky, A. and Kahneman, D. (1974) ‘Judgment under uncertainty: heuristics and biases’, *Science*, 185(4157), pp. 1124–1131.

- Sunstein, C.R. (2002) ‘Probability neglect: emotions, worst cases, and law’, *Yale Law Journal*, 112(1), pp. 61–107.

Apollo, Cold War space policy, and mobilization precedent

- Chaikin, A. (1994) *A Man on the Moon: The Voyages of the Apollo Astronauts*. New York: Penguin.

- Logsdon, J.M. (2010) *John F. Kennedy and the Race to the Moon*. New York: Palgrave Macmillan.

- Roosevelt, F.D. (1940) ‘The Arsenal of Democracy’ (Fireside Chat), 29 December. (Use an archival transcript edition appropriate to your publisher’s style.)

- Logsdon, J.M. (ed.) (series) *Exploring the Unknown: Selected Documents in the History of the U.S. Civil Space Program*. Washington, DC: NASA History Division.

HIV/AIDS, regulation, and "speed as a binding constraint"

- Epstein, S. (1996) *Impure Science: AIDS, Activism, and the Politics of Knowledge*. Berkeley, CA: University of California Press.

- Shilts, R. (1987) *And the Band Played On: Politics, People, and the AIDS Epidemic*. New York: St Martin's Press.

- U.S. Food and Drug Administration (2006) *FDA Centennial Edition: Milestones in U.S. Food and Drug Law History*. Silver Spring, MD: FDA.

Longevity biology: "aging is malleable" anchors (primary science)

- Kenyon, C., Chang, J., Gensch, E., Rudner, A. and Tabtiang, R. (1993) 'A C. elegans mutant that lives twice as long as wild type', *Nature*, 366, pp. 461–464.

- López-Otín, C., Blasco, M.A., Partridge, L., Serrano, M. and Kroemer, G. (2013) 'The hallmarks of aging', *Cell*, 153(6), pp. 1194–1217.

- McCay, C.M., Crowell, M.F. and Maynard, L.A. (1935) 'The effect of retarded growth upon the length of life span and upon the ultimate body size', *Journal of Nutrition*, 10(1), pp. 63–79.

- Ocampo, A., Reddy, P., Martinez-Redondo, P., et al. (2016) 'In vivo amelioration of age-associated hallmarks by partial reprogramming', *Cell*, 167(7), pp. 1719–1733.e12.

- Takahashi, K. and Yamanaka, S. (2006) 'Induction of pluripotent stem cells from mouse fibroblasts by defined factors', *Cell*, 126(4), pp. 663–676.

- Wilmut, I., Schnieke, A.E., McWhir, J., Kind, A.J. and Campbell, K.H.S. (1997) 'Viable offspring derived from fetal and adult mammalian cells', *Nature*, 385, pp. 810–813.

Spend / funding / scale-of-resourcing anchors (for your conclusion)

- AbbVie (2014) 'AbbVie and Calico announce R&D collaboration to discover, develop and commercialize novel therapies for age-related diseases'. Press release.

- EvaluatePharma (2024) *World Preview 2024, Outlook to 2028*. (Industry report; use for global pharma R&D trendlines.)

- Financial Times (2024) ‘Longevity biotech Retro Biosciences seeks to raise $1bn’. *Financial Times*.

- National Institute on Aging (2025) *Fiscal Year 2026 Congressional Budget Justification* (or equivalent budget/appropriations history document). Bethesda, MD: NIH/NIA.

- Axios (2022) ‘Longevity startup NewLimit raises $40M’. *Axios*.