LongGame Ventures - Investment Themes

We are interested in investing at the intersection of deep translational biology, artificial intelligence, and scalable preventative health infrastructure. Over the coming decade, we believe the defining opportunities in longevity will be driven by unprecedented data density, enabling a shift toward highly personalized precision medicine and targeted treatments. Our thesis focuses on three distinct pillars within this ecosystem: from Theme 1, where we back the precise delivery and control mechanisms needed to scale cellular reprogramming therapies from localized tissues to systemic clinical treatments; to Theme 2, where we fund the proprietary, closed-loop wet lab data engines required to train next-generation AI foundation models to decode aging biology; and finally to Theme 3, where we support the rigorous measurement layers and science-backed consumer solutions that utilize individual data to deliver true personalization and physiological optimization to the market.

Theme 1: Reprogramming, beyond the local

The Problem

Partial reprogramming has been a striking preclinical result for years without a way to know whether it does anything in humans. That changed in 2026, when the first partial reprogramming therapy was dosed in people: an OSK gene therapy for age-related optic neuropathies, delivered by intravitreal injection and switched on by systemic doxycycline. Whatever the early readout shows, the category has crossed from claim to clinical question, and the next several years will be spent answering it.

The design of that first trial also explains why the first uses are local. The eye is immune-privileged, the dose is contained, and the inducible switch gives external control over how long the factors stay on. Those are the conditions that let you test the biology in humans before taking on systemic risk. A clean safety and vision readout would validate the approach in that narrow setting. It would not tell you whether reprogramming generalizes to tissues that are not structurally protected.

Opportunity

The bar is functional. A program has to show it improves how a tissue works, not just that an aging clock moved.

What clears that bar at scale is a specific combination: delivery precise enough to reach the intended cells, control tight enough to switch reprogramming on and off, and a safety profile good enough to give systemically. Get those together and the constraint flips. A program is no longer limited to whichever tissue happens to be contained and privileged. It can go after organs by how vulnerable they are, treating the tissues where aging does the most damage first, rather than the ones that are simply easiest to reach. The eye is first because it is contained. The companies that create lasting value will be the ones whose delivery and control let them prioritize by need.

The engineering questions follow. How specifically can reprogramming be aimed so the right cells in the right tissue get the signal. How tightly can it be held so cells move backward in age without losing identity or drifting toward pluripotency. How durable the effect is, and what the dosing looks like. And how you reach less contained tissue without systemic toxicity.

Examples of Solutions We Want to See

- Precise delivery that reaches the intended cells in the intended tissue
- Tight external control over when reprogramming is active and for how long
- Safety profiles that allow dosing beyond contained, privileged tissue
- Smaller or safer factor sets that widen the safety margin
- Approaches that let a program prioritize vulnerable organs rather than only accessible ones

What We Would Pass On

- Programs resting on clock movement with no functional readout
- Delivery that only works in contained, privileged tissue with no path to controlled systemic use
- Approaches that buy efficacy with a safety margin too thin to dose where it matters

Theme 2: The data engine, not the model

The Problem

The constraint is not the amount of data. It is the kind. The data that gets labeled longevity-relevant is abundant: wearable biometrics, consumer biological-age tests, self-tracked supplements and lifestyle logs. Owning that measurement layer is a real position in the consumer, which we come to later. For discovery it is too shallow to carry a mechanism, so more of it does not move the problem. What is needed is depth: molecular detail at the resolution models can use, tied to validated readouts and generated on aging biology. That data mostly does not exist yet, and the deeper reason is the timescale. The endpoint we care about plays out over decades, which forces reliance on surrogate aging biomarkers whose validity is unestablished, so even a large dataset can leave a model optimizing against the wrong target. No public resource fixes this, the large biobanks included; they are built for disease association, not to resolve aging mechanisms at the depth and over the timescale this needs. The teams that succeed will not be the ones licensing data. They will be the ones generating proprietary, well-characterized data on the right biology.

Opportunity

The cost of generating that kind of data is falling fast. Single-cell profiling has gone from dollars to cents per cell, and functional screens from a few hundred perturbations to genome-wide in a few years. Automation, cheaper sequencing, and stronger foundation models all push in the same direction. But the therapeutic bar does not move: a longevity drug has to act systemically, be druggable, and be safe in healthy people. That bar is what makes the data requirement steep. Safety in healthy people cannot be established from public datasets built mostly on sick ones, so the data has to be generated, at the depth and on the biology the question demands.

Generating that data is expensive, and the fair question is what stops a far better-funded team from outspending an early stage company to produce the same thing. The answer is focus: a specific biology the larger players are not working on, and an integration depth across wet lab, computation, and model that a bigger, more diffuse effort will not replicate.

Examples of Solutions We Want to See

- Closed-loop discovery platforms where experiments, computation, and model updates run as one iterative system rather than separate steps
- Automated experimental systems for aging biology, in primary human cells, organoids, or model organisms, at scales that generate usable signal
- Longitudinal multi-omic platforms designed from the start to produce improving data
- Model-building efforts grounded in a real data pipeline, not the model alone

What We Would Pass On

- Teams licensing public datasets and calling it a moat
- Model-first pitches with no proprietary data generation underneath
- Platforms whose data does not improve with use
- Engines sold on the platform story with no validated target or asset to show they work

Theme 3: Consumer longevity, on solid science

The Problem

This category has grown faster than the evidence under most of it. Advanced blood panels are common, supplements multiply after every viral preprint, and much of what consumers receive is closer to marketing than to science. That gap is the opportunity. In a market this noisy, value accrues to the few companies built on real science, and most of the work is the discipline to tell which those are.

Opportunity

The durable opportunities sit in two jobs, both able to rest on real science: preventing problems before they surface, and helping people who are not yet sick improve and hold their health. Prevention and optimization. Most of what is worth backing is a version of one or both.

These play out across a growing set of areas, and two are live now. One is continuous monitoring. Diagnostics are moving from snapshots to real-time, and the next step is rapid, scalable multiplexing: one device reading many classes of biomarker at once, hormones, metabolites, nucleic acids, instead of a single analyte at a time. That serves both jobs, catching drift early for prevention and letting healthy people track and tune. The other is the population GLP-1 drugs created: a large group now managing their metabolic health, where the opening is live tracking and tools to hold the benefit after people come off, when most of the weight returns. These are examples, not the whole list.

What separates the companies worth backing is the same across every one of these areas: they own the measurement layer that other products build on, not the app on top of it that becomes a feature, and their interventions rest on real mechanism and human data, not language borrowed from a viral preprint.

Examples of Solutions We Want to See

- Measurement infrastructure that improves real-time molecular readouts from the body with ability to multiplex.
- Companies working the next clinical questions in metabolic health beyond weight loss
- Evidence-based consumer interventions for preventative health that are rooted in solid scientific backing.

What We Would Pass On

- Interventions whose case is a narrative wrapped around a viral preprint
- Hardware with no credible answer on retention or margin
- Formulations marketed on adjacent science they do not actually rest on

Apply here: https://longgame.vc/l3