


The New Operating System of Biotech

A Zentynel Position Paper
By Cristián Hernández, General Partner

MAY 2026



Executive Summary

Most conversations about artificial intelligence in biotechnology still focus on tools, productivity, automation, and the use of AI inside existing workflows. That framing is understandable, but it misses the deeper shift now beginning to reshape how biology is studied, engineered, financed, and scaled.

At Zentynel, we believe AI is not simply becoming a software layer on top of biology, but an infrastructure layer in the construction of biological knowledge. When the production of knowledge changes, the architecture of companies changes with it, and the logic of venture capital eventually changes as well.

This paper is written for founders, investors, and ecosystem builders who need a practical way to understand this transition. It is not a technical review of AI models, nor a general survey of every use case across life sciences. It is a position paper about how AI may change the way biotech companies learn, operate, and raise capital.

The distinction matters because life sciences contain several layers of AI transformation. AI can improve health system operations, personalize diagnosis and treatment, and accelerate the development of new biological solutions. Our primary focus here is the third layer, where AI begins to reshape discovery, company architecture, and venture underwriting.

Biotech has always advanced through uncertain loops of hypothesis, experimentation, failure, interpretation, and reinvestment, usually at high cost and low speed. If AI can compress part of those loops without pretending to eliminate clinical risk, the industry begins to behave under a different investment equation.

The first durable wave of value will not come from removing biology from biotech, because biology remains the judge and patients remain the final test. It will come from companies that learn faster, fail cheaper, operate with less friction, and convert scientific insight into higher quality decisions.

The best biotech companies of the next decade will not be the ones that say they use AI, but the ones that could not exist without it.



1. From Tools to a New Knowledge Infrastructure

The most important mistake in today's AI and biotech conversation is to reduce the shift to a question of productivity inside existing workflows. Productivity matters, but it is not the full story when AI systems begin to participate directly in how biological hypotheses are formed and tested.

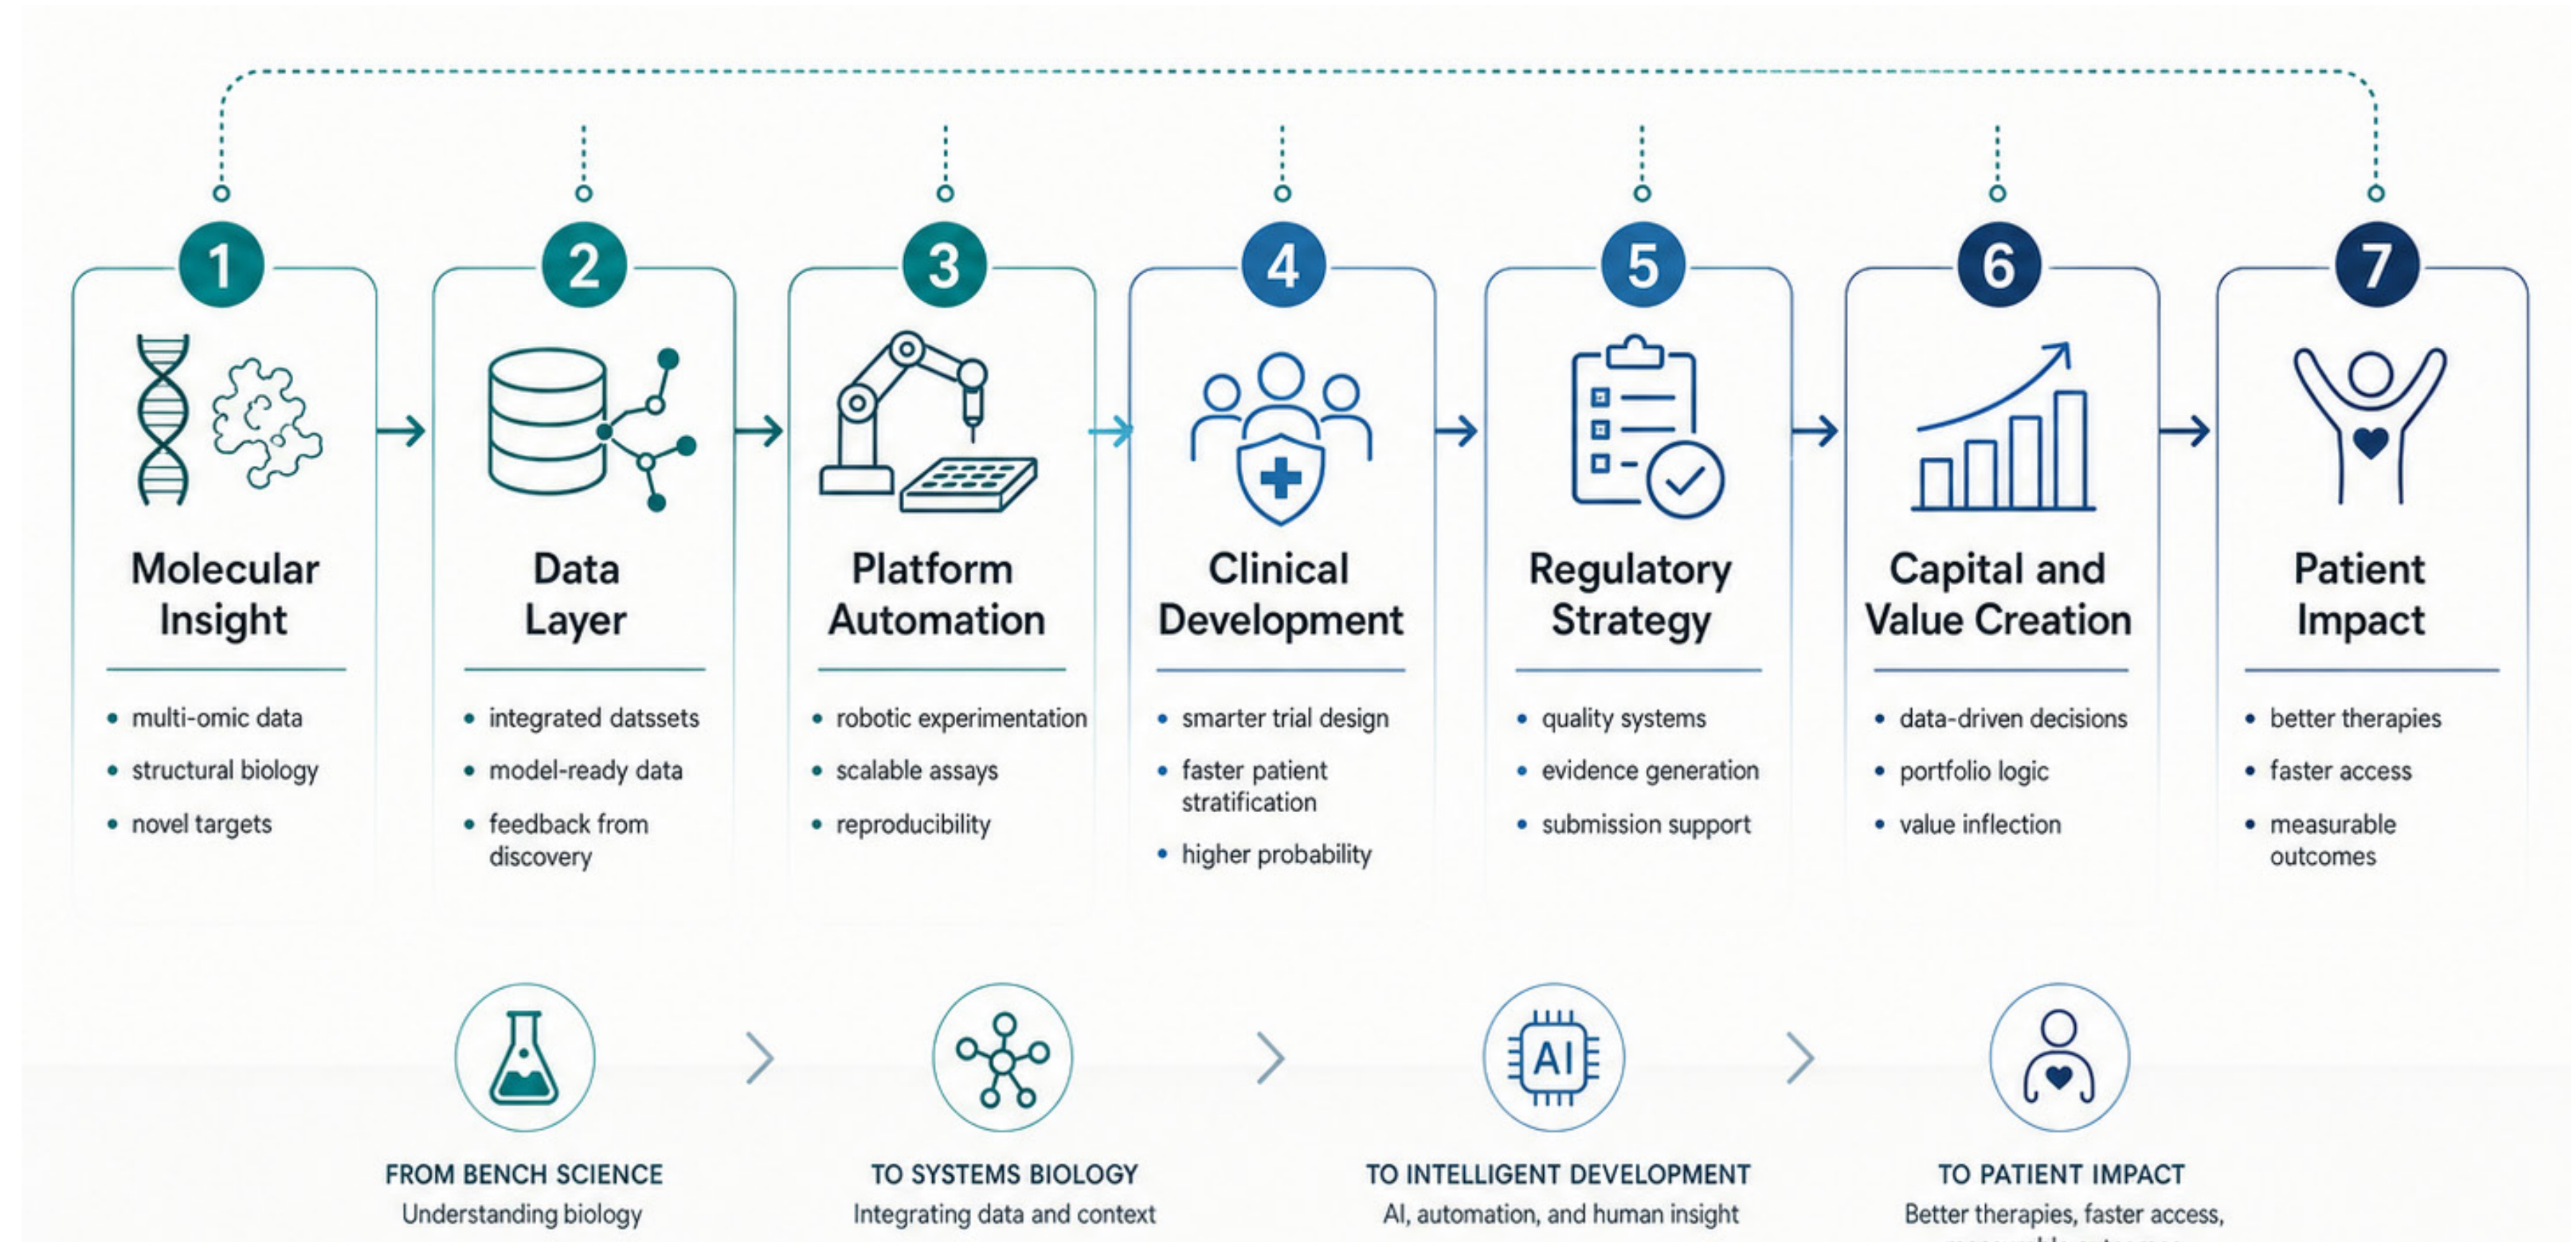
For decades, biotechnology was largely an exercise in observation under uncertainty, with expensive experiments and feedback cycles measured in months or years. Teams advanced by combining intuition, domain expertise, scarce data, and iterative experimentation, while capital funded the next layer of uncertainty reduction.

AI begins to alter this pattern by connecting discovery, data interpretation, automation, simulation, experimental prioritization, and decision support. The result is not a magical shortcut through biology, but the emergence of a different operating logic for how scientific companies can learn.

This is why the transition should not be understood only as a discovery acceleration story. It is also a story about learning systems, company design, and capital allocation under uncertainty.

When the speed and cost of learning change, venture capital must adjust how it evaluates teams, milestones, data assets, and capital requirements. That is the lens through which the rest of this paper should be read.

Figure 1. The emerging AI and biotech stack connects molecular biology, data integration, platform automation, clinical development, regulatory strategy, capital allocation, and patient impact into a more integrated operating system.



2. From Invisible Biology to Engineered Systems

For many years, AI in biotech appeared in conferences, pharma innovation units, venture decks, and scientific panels as an important but still abstract promise. The symbolic break came with AlphaFold, which helped move AI from the periphery of biotech strategy into the center of biological reasoning.

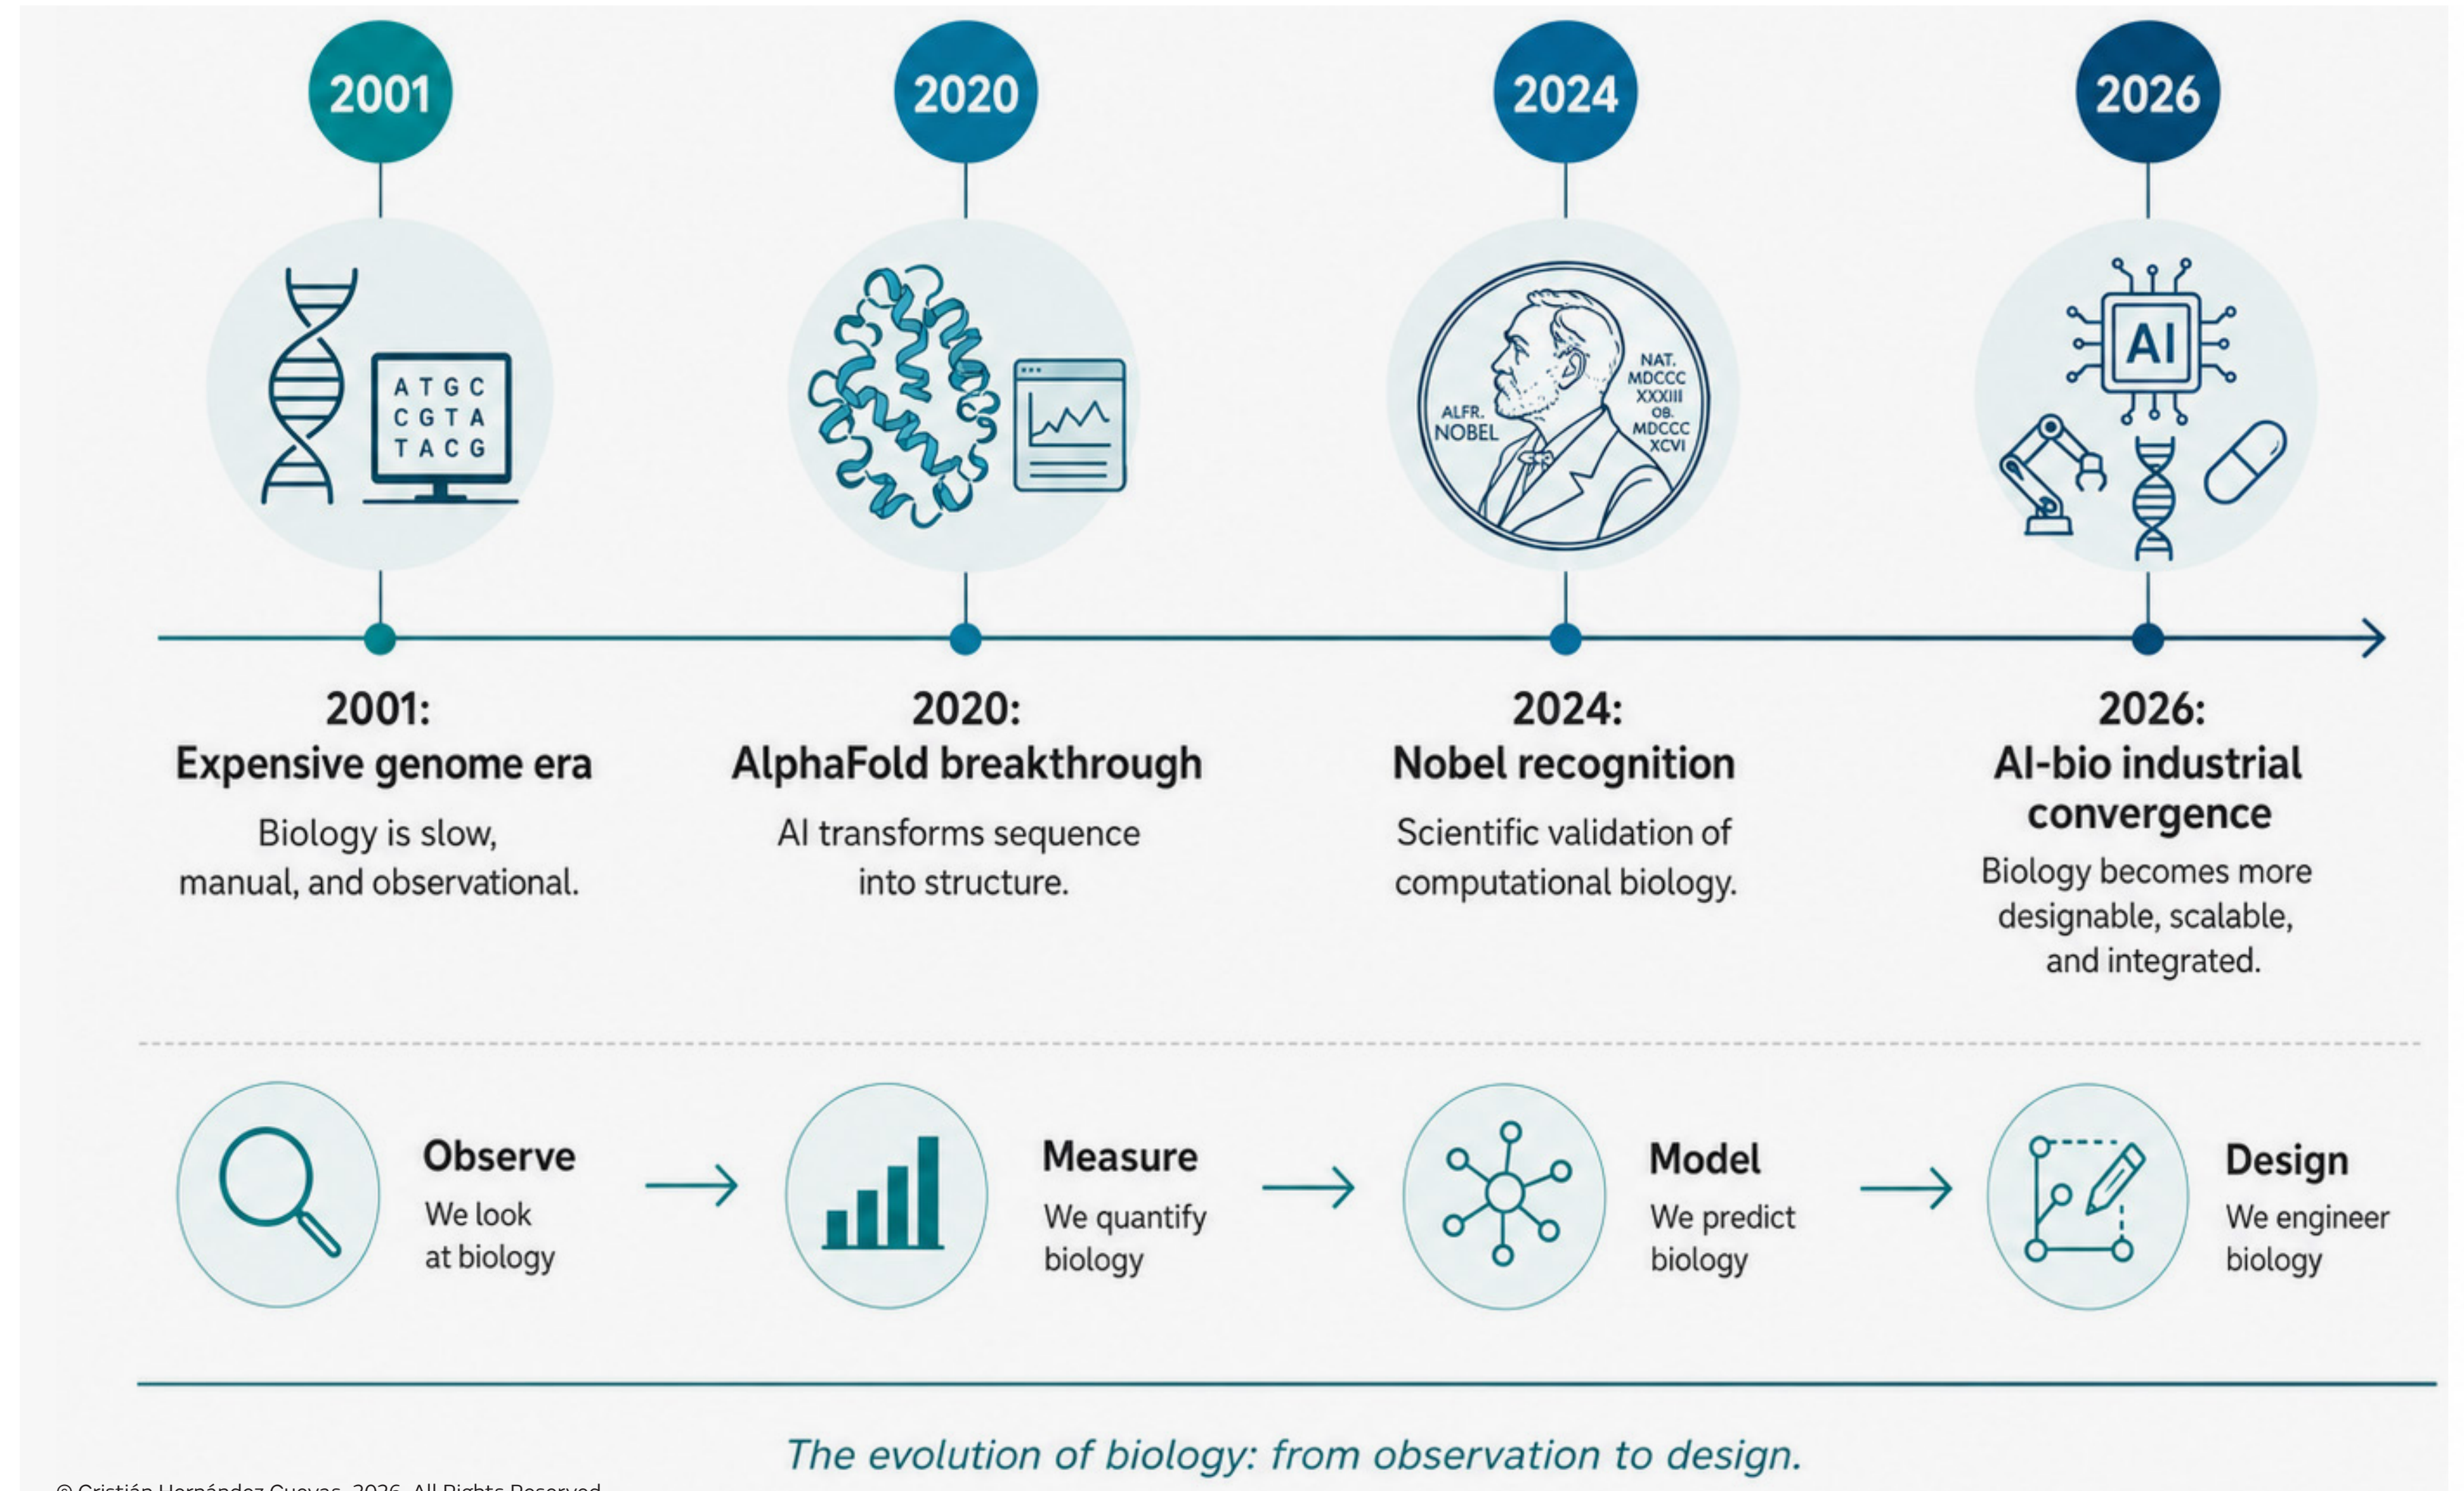
In 2024, the Nobel Prize in Chemistry recognized David Baker for computational protein design and Demis Hassabis and John Jumper for protein structure prediction. The Nobel Committee described the DeepMind work as solving a long standing problem in predicting complex protein structures from amino acid sequences.

That recognition did not mean that biology had been solved, but it changed the collective sense of what could be modeled, searched, designed, and iterated. This psychological shift matters because technological transitions often begin when a field changes its sense of what is possible.

Biology has always been partially invisible, because we observe phenomena, infer mechanisms, generate hypotheses, and then validate slowly through experiments. AlphaFold did not remove that uncertainty, but it demonstrated that parts of biological reality could be modeled computationally at extraordinary scale.

Once a field believes that biological systems can be modeled and designed with computational assistance, capital starts to move differently. The next section shows why that shift is already visible in the market.

Figure 2. The AI bio race has moved from the expensive genome era into an industrial convergence phase, where observation, measurement, modeling, and design increasingly operate as one connected trajectory.



© Cristián Hernández Cuevas, 2026. All Rights Reserved.



3. Capital Is Moving Before Complete Proof

The market is not waiting for AI enabled biotech to be fully proven, because it is already funding the possibility that it may become foundational. Isomorphic Labs announced a 600 million dollar external financing round in March 2025, led by Thrive Capital with participation from GV and Alphabet.

The direction accelerated further in May 2026, when Isomorphic announced a 2.1 billion dollar Series B round to scale its AI driven drug design engine. The speed of that capital formation illustrates how quickly the category is moving from scientific possibility to industrial infrastructure.

In January 2026, NVIDIA and Lilly announced a co innovation AI lab for drug discovery, with plans to invest up to one billion dollars over five years in talent, infrastructure, compute, and research. The collaboration is notable because it connects pharma expertise, accelerated computing, AI models, robotics, and experimental infrastructure.

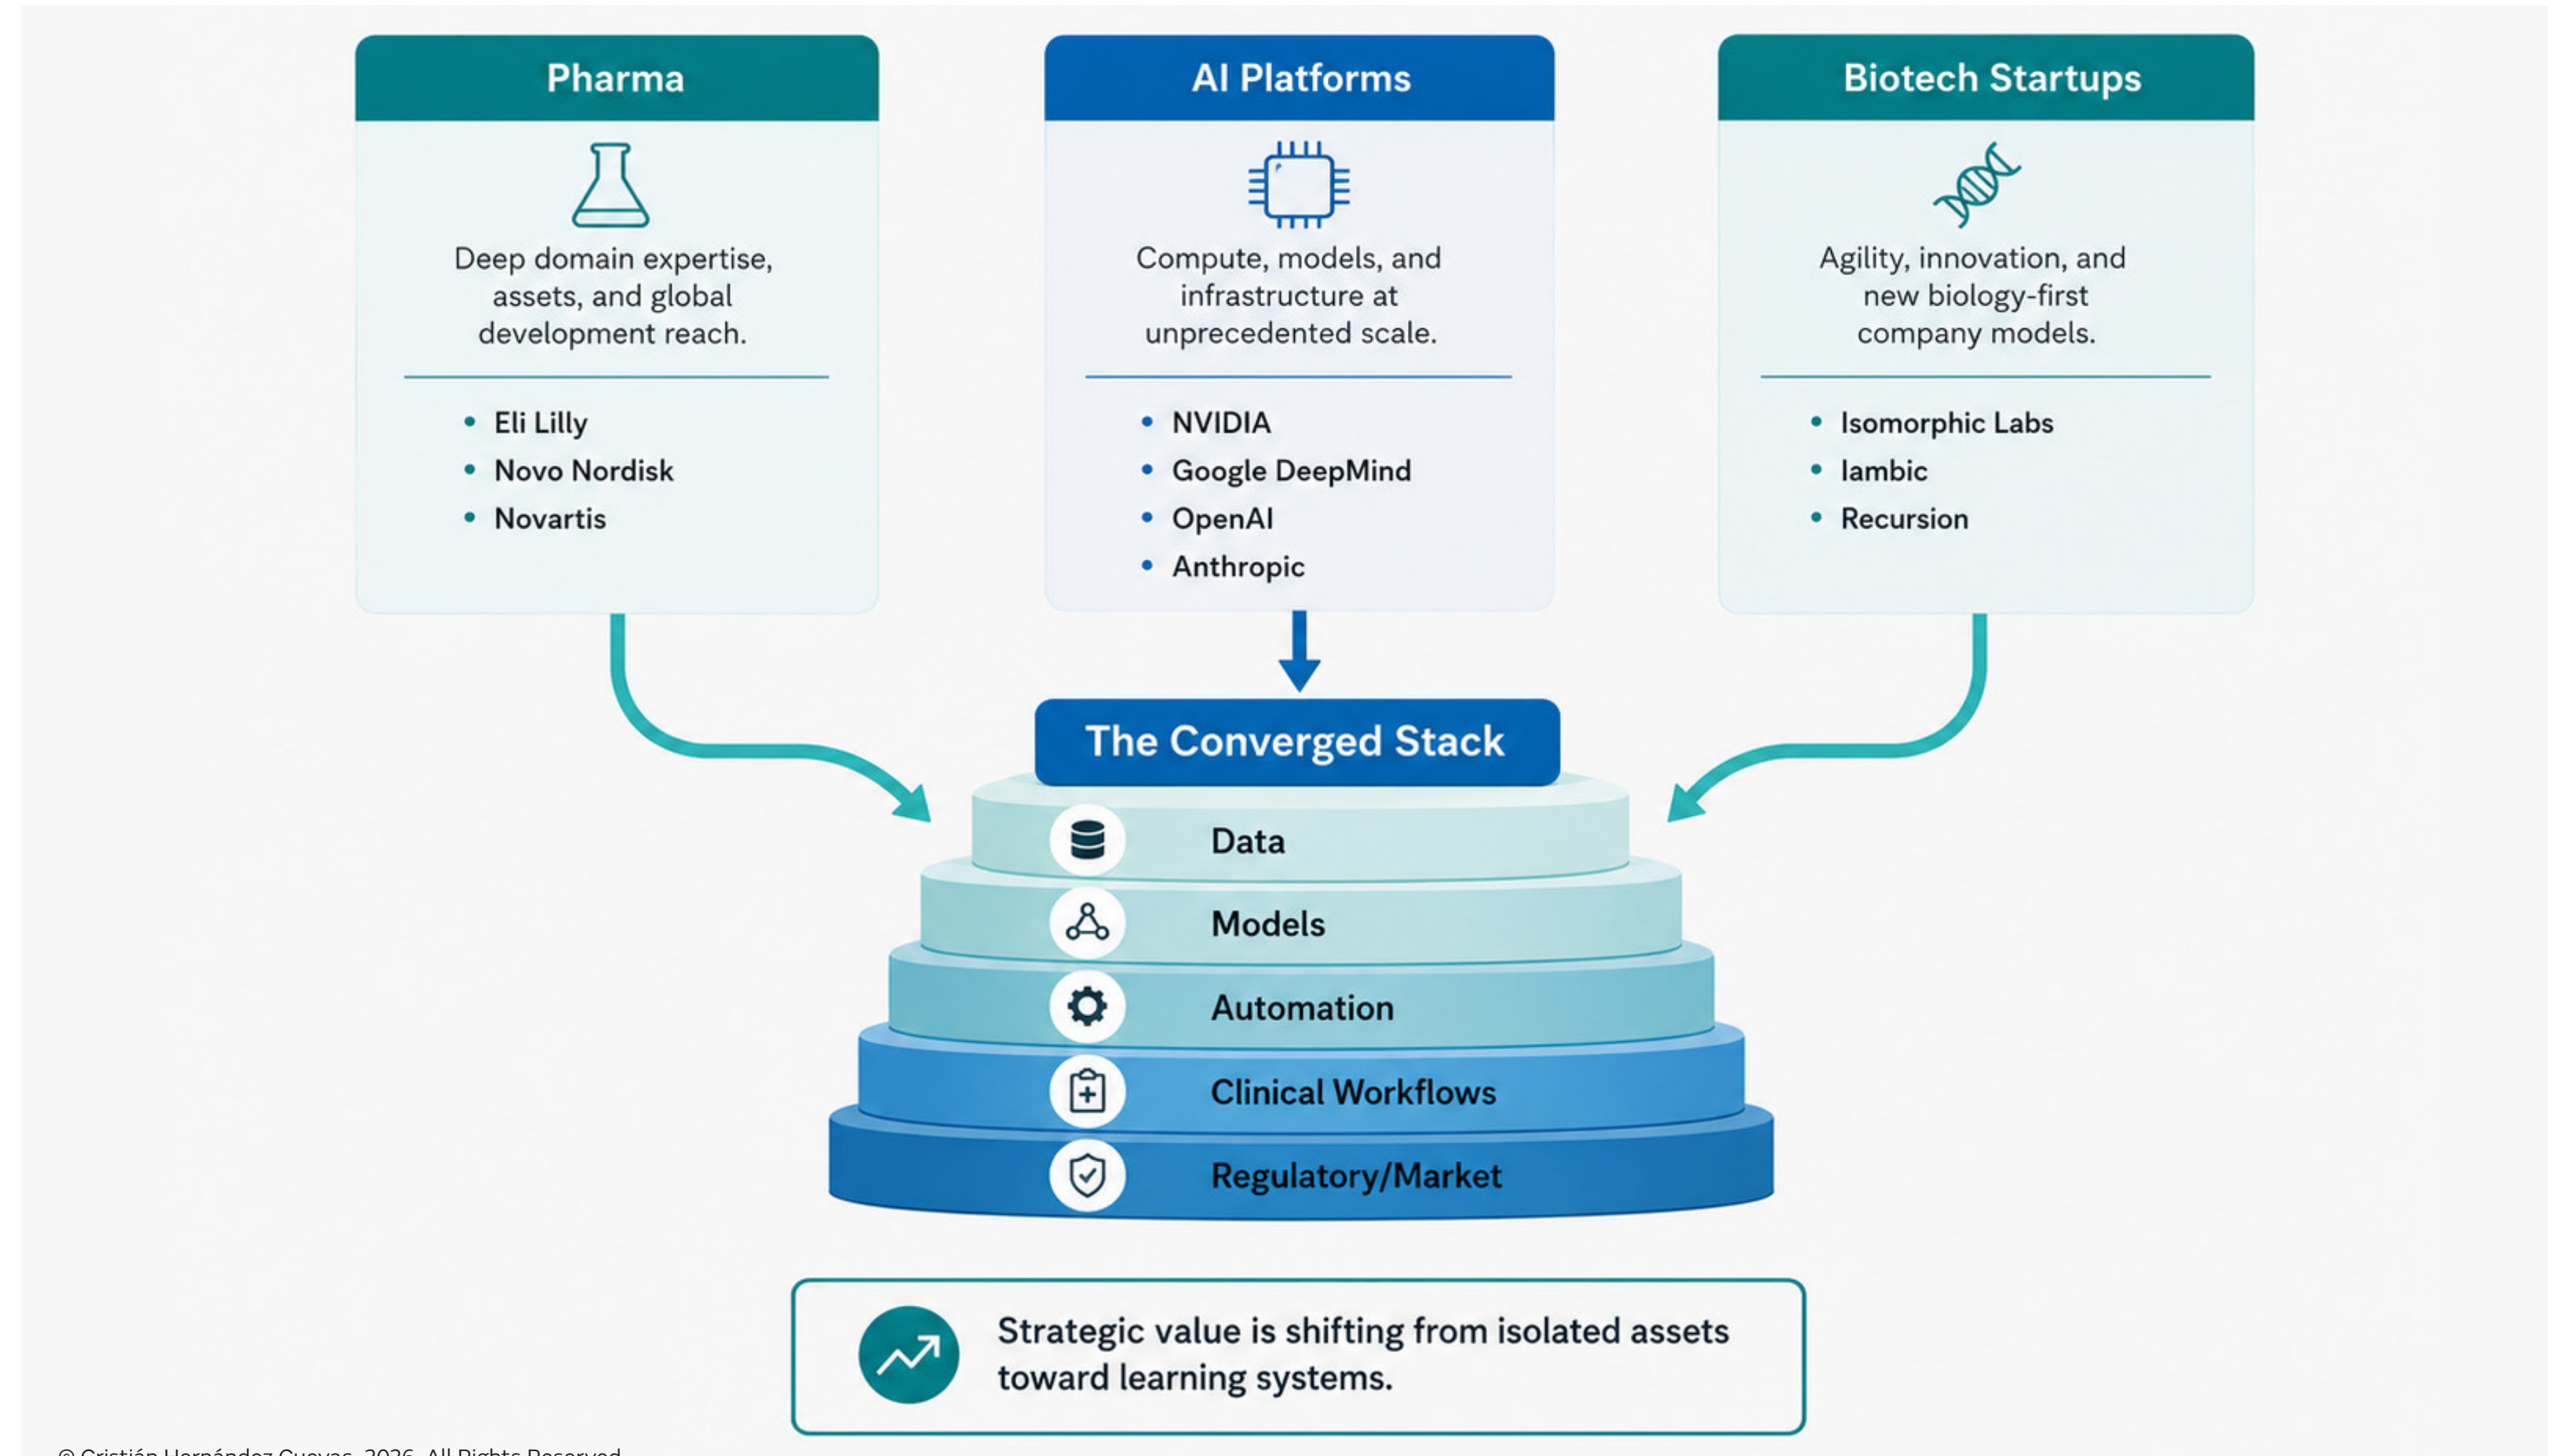
In February 2026, Iambic announced a multi year collaboration with Takeda using its AI driven discovery platform, with success based payments that could exceed 1.7 billion dollars. The agreement gives Takeda access to Iambic's models and wet lab capabilities to advance small molecule programs in oncology and inflammatory diseases.

These are not isolated headlines, but signals that pharma is increasingly buying access to new learning systems and not only to discrete molecules. The old question was who owns the best asset, while the emerging question is who owns the best system for generating, validating, and learning from assets.

This difference moves competitive advantage away from a single product and toward the data, workflow, and feedback architecture behind it. When an industry changes architecture, waiting for complete certainty can be rational in theory and competitively fatal in practice.

Capital is therefore funding optionality, infrastructure, data access, platform capacity, and a new operating system for drug discovery and development. This is not a bet that clinical biology has become easy, but a bet that biological learning may improve structurally.

Figure 3. Capital is moving toward the converged stack because strategic value is shifting from isolated assets toward data, models, infrastructure, validation workflows, and platform learning capacity.



© Cristián Hernández Cuevas, 2026. All Rights Reserved.



4. Biology Still Decides What Is Real

Every technological shift carries the same danger, because the narrative can run ahead of the evidence and create confidence before translation exists. In biotech, this danger is especially acute because biology has a way of correcting narrative excess through toxicity, heterogeneity, and clinical failure.

AI can accelerate hypothesis generation, compress search spaces, prioritize candidates, improve molecular design, and support experimental planning. It does not remove biological complexity, clinical execution risk, regulatory scrutiny, patient variability, or the need for a strong underlying thesis.

A model cannot invent information that the data does not contain, and much of the most valuable biological data remains fragmented or proprietary. This is the central tension of the current moment, because AI is accelerating discovery but has not yet compressed the full clinical arc.

Drug development remains long, expensive, and uncertain, with PhRMA estimating 10 to 15 years and 2.6 billion dollars to develop one new medicine. PhRMA also states that less than 12 percent of candidate medicines entering Phase I clinical trials are ultimately approved by the FDA.

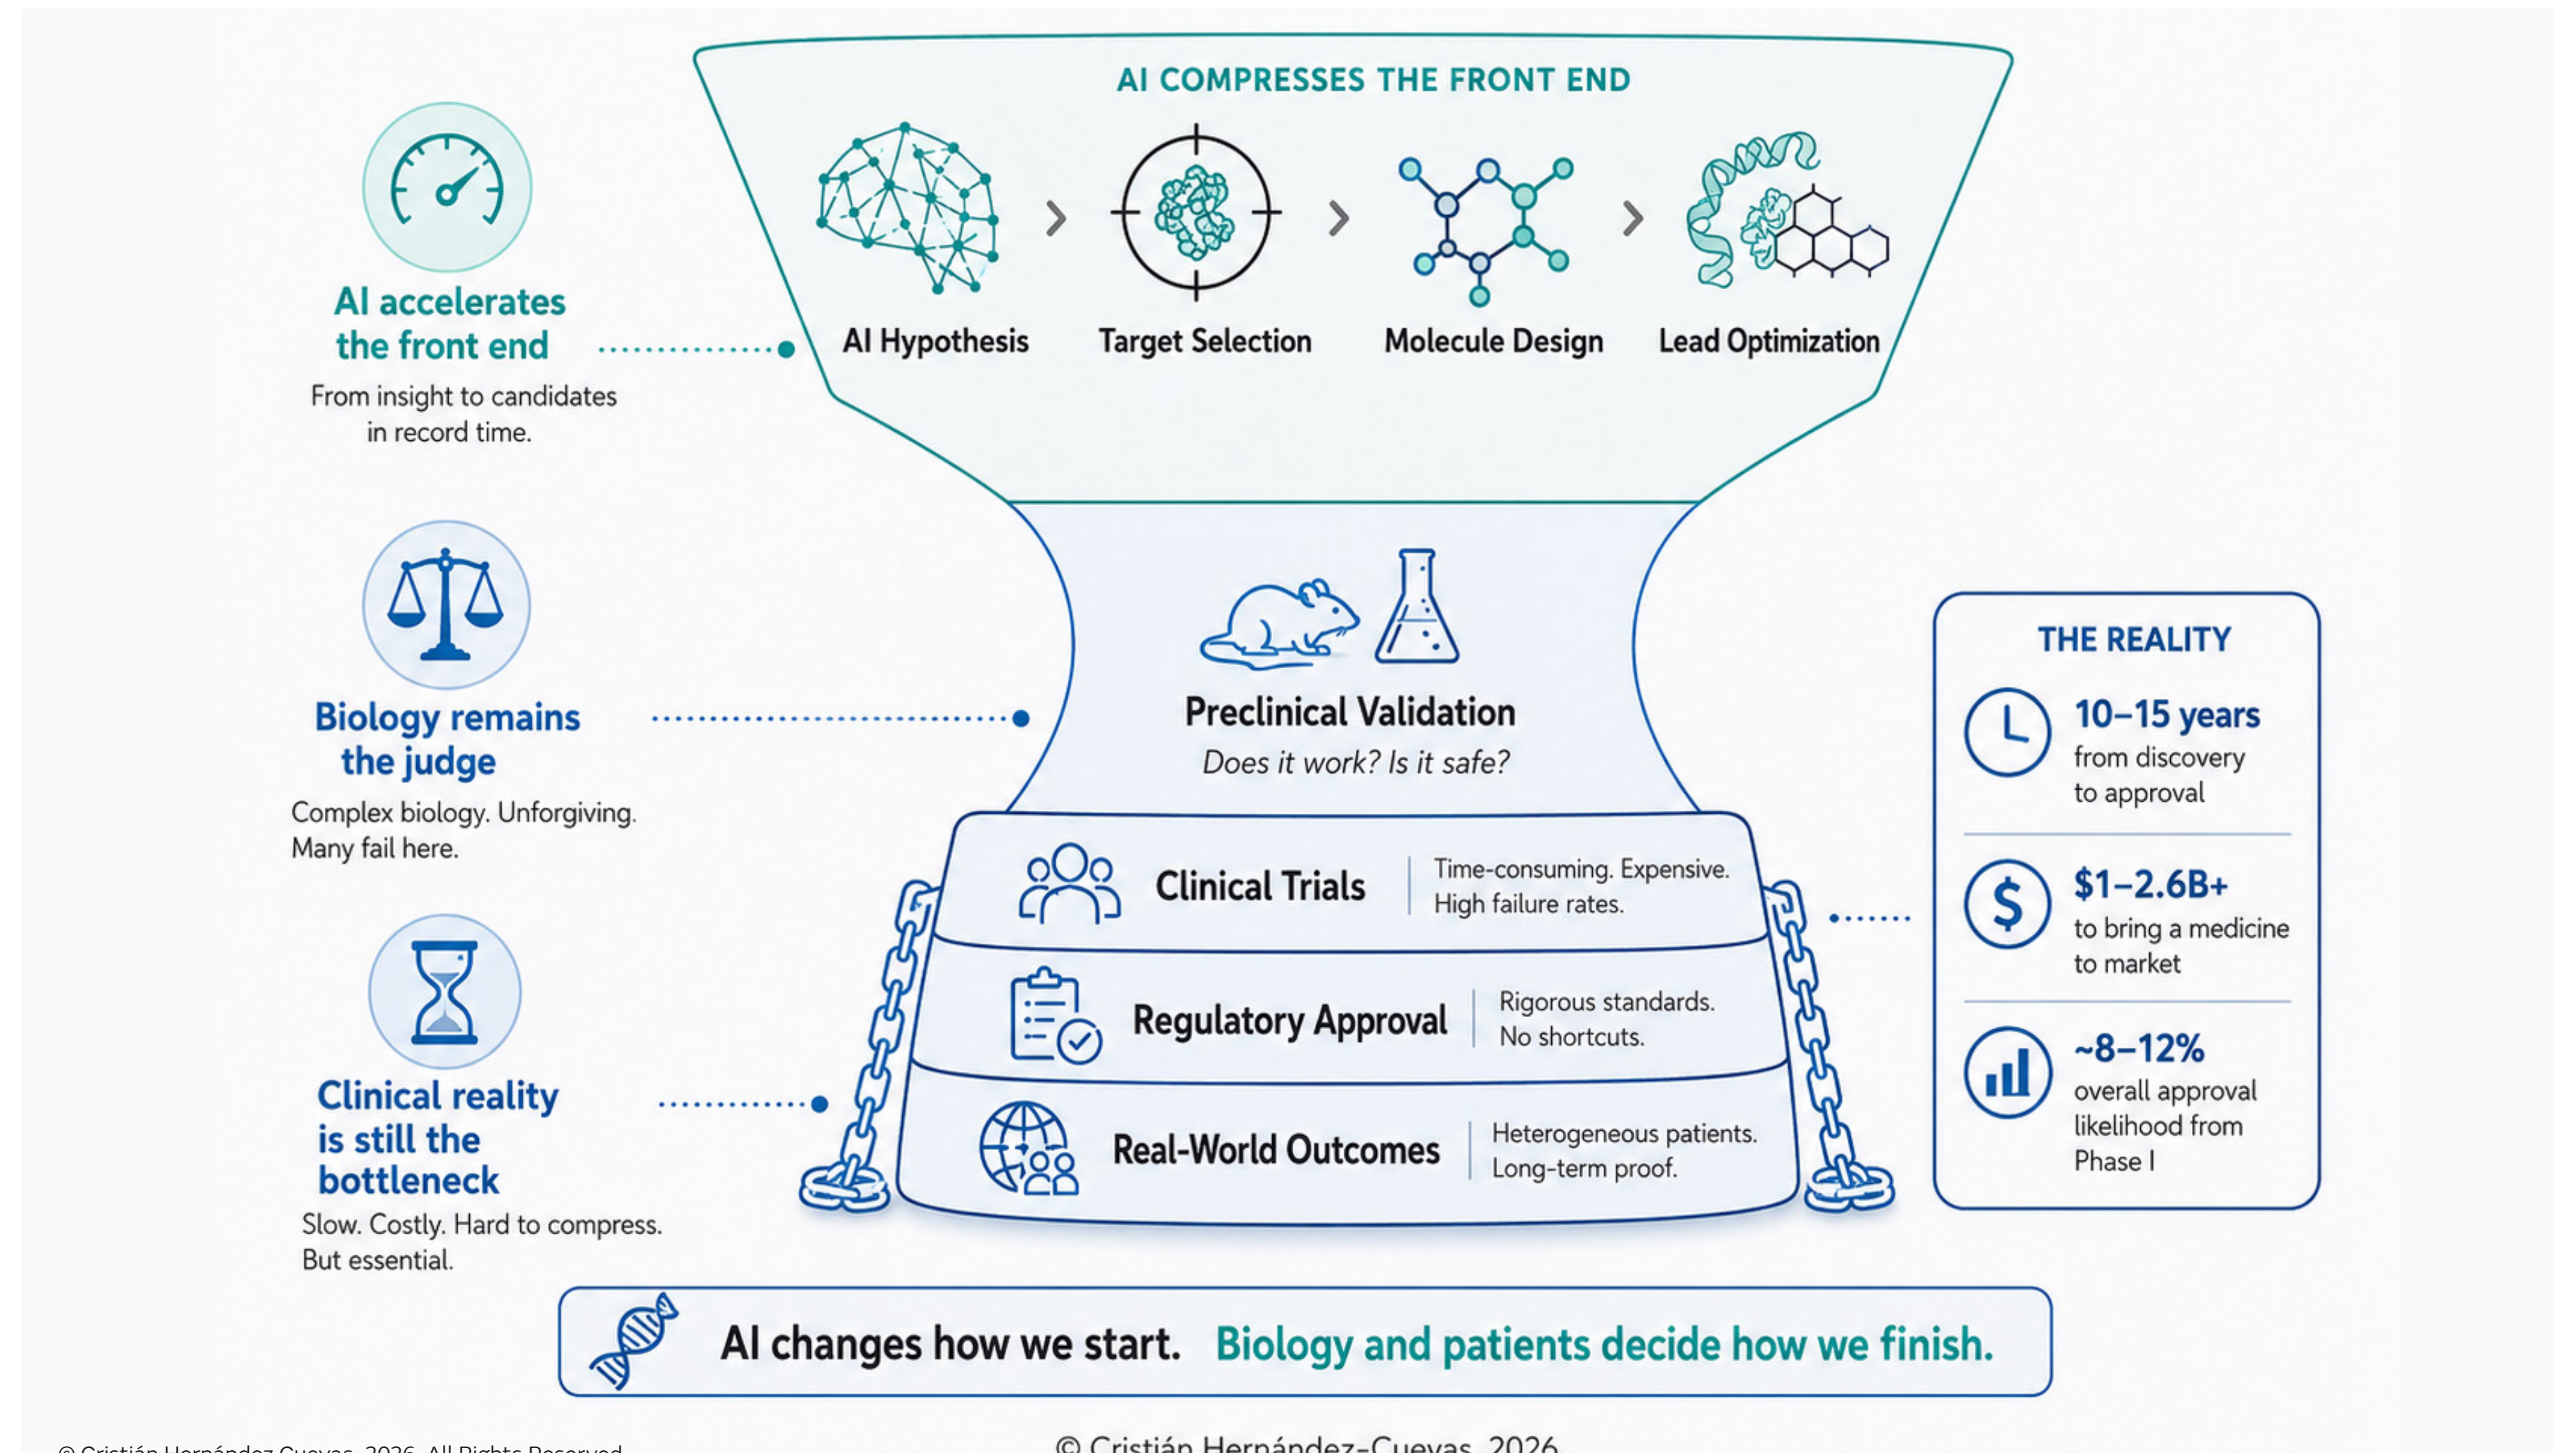
A BIO, Informa Pharma Intelligence, and QLS Advisors analysis found a 7.9 percent overall likelihood of approval from Phase I during 2011 to 2020. That data keeps clinical translation at the center of any serious biotech investment discussion.

The lesson is not that AI is overhyped, but that the standard must remain precise about where AI accelerates and where biology still governs. AI does

not remove biotech risk, yet it may change the cost, speed, and quality of biological learning, which is already enough to matter.

That framing leads to the core investment question. If biology still decides what is real, where can AI create durable value before clinical certainty arrives?

Figure 4. AI compresses the front end of biotech by improving hypothesis generation, target selection, design, and prioritization, while clinical trials, regulatory approval, and real world outcomes remain the decisive bottlenecks.



5. The Real Value Is the Economics of Learning

Biotech has always been a learning business, because every company spends capital to transform biological uncertainty into investable evidence. The question is not only whether a company has a promising scientific hypothesis, but how efficiently it can convert uncertainty into decision quality.

That conversion is expensive, because a weak candidate can absorb years of work and a financing round before its limitations become visible. A poorly designed experiment can create ambiguity rather than insight, leaving investors and teams with more activity than actual risk reduction.

AI does not eliminate this reality, but it can change the price of each learning cycle by improving hypothesis generation and experimental prioritization. If a company can interpret data faster and kill weak directions earlier, the capital model begins to shift before clinical risk disappears.

The value is not only in producing more success, because disciplined failure is itself a strategic advantage in a sector where learning is expensive. Capital efficiency in biotech is not merely about spending less; it is about learning more per dollar and reducing ambiguity with each decision.

This may be AI's first enduring impact in biotech, not by making the field easy but by making learning less wasteful and more cumulative. Once learning becomes cheaper and faster, the next question is whether the company itself has been designed to capture that advantage.



6. Company Architecture Is Now Part of the Science

Many companies still treat AI as a feature, adding it to the deck, hiring a data scientist, or integrating a model into existing workflows. That may improve productivity, but it does not by itself create an AI first company or a defensible scientific operating system.

AI first is not a product decision, because it is an organizational architecture that changes how work is assigned, executed, verified, and improved. The relevant question is no longer only how to make a person more efficient, but why a person is doing a repeatable task in the first place.

This architecture will matter most in companies where work depends on repeated cycles of prediction, testing, interpretation, and prioritization. Drug discovery platforms, synthetic biology companies, computational biology teams, diagnostic engines, clinical operations platforms, and translational research organizations are natural candidates.

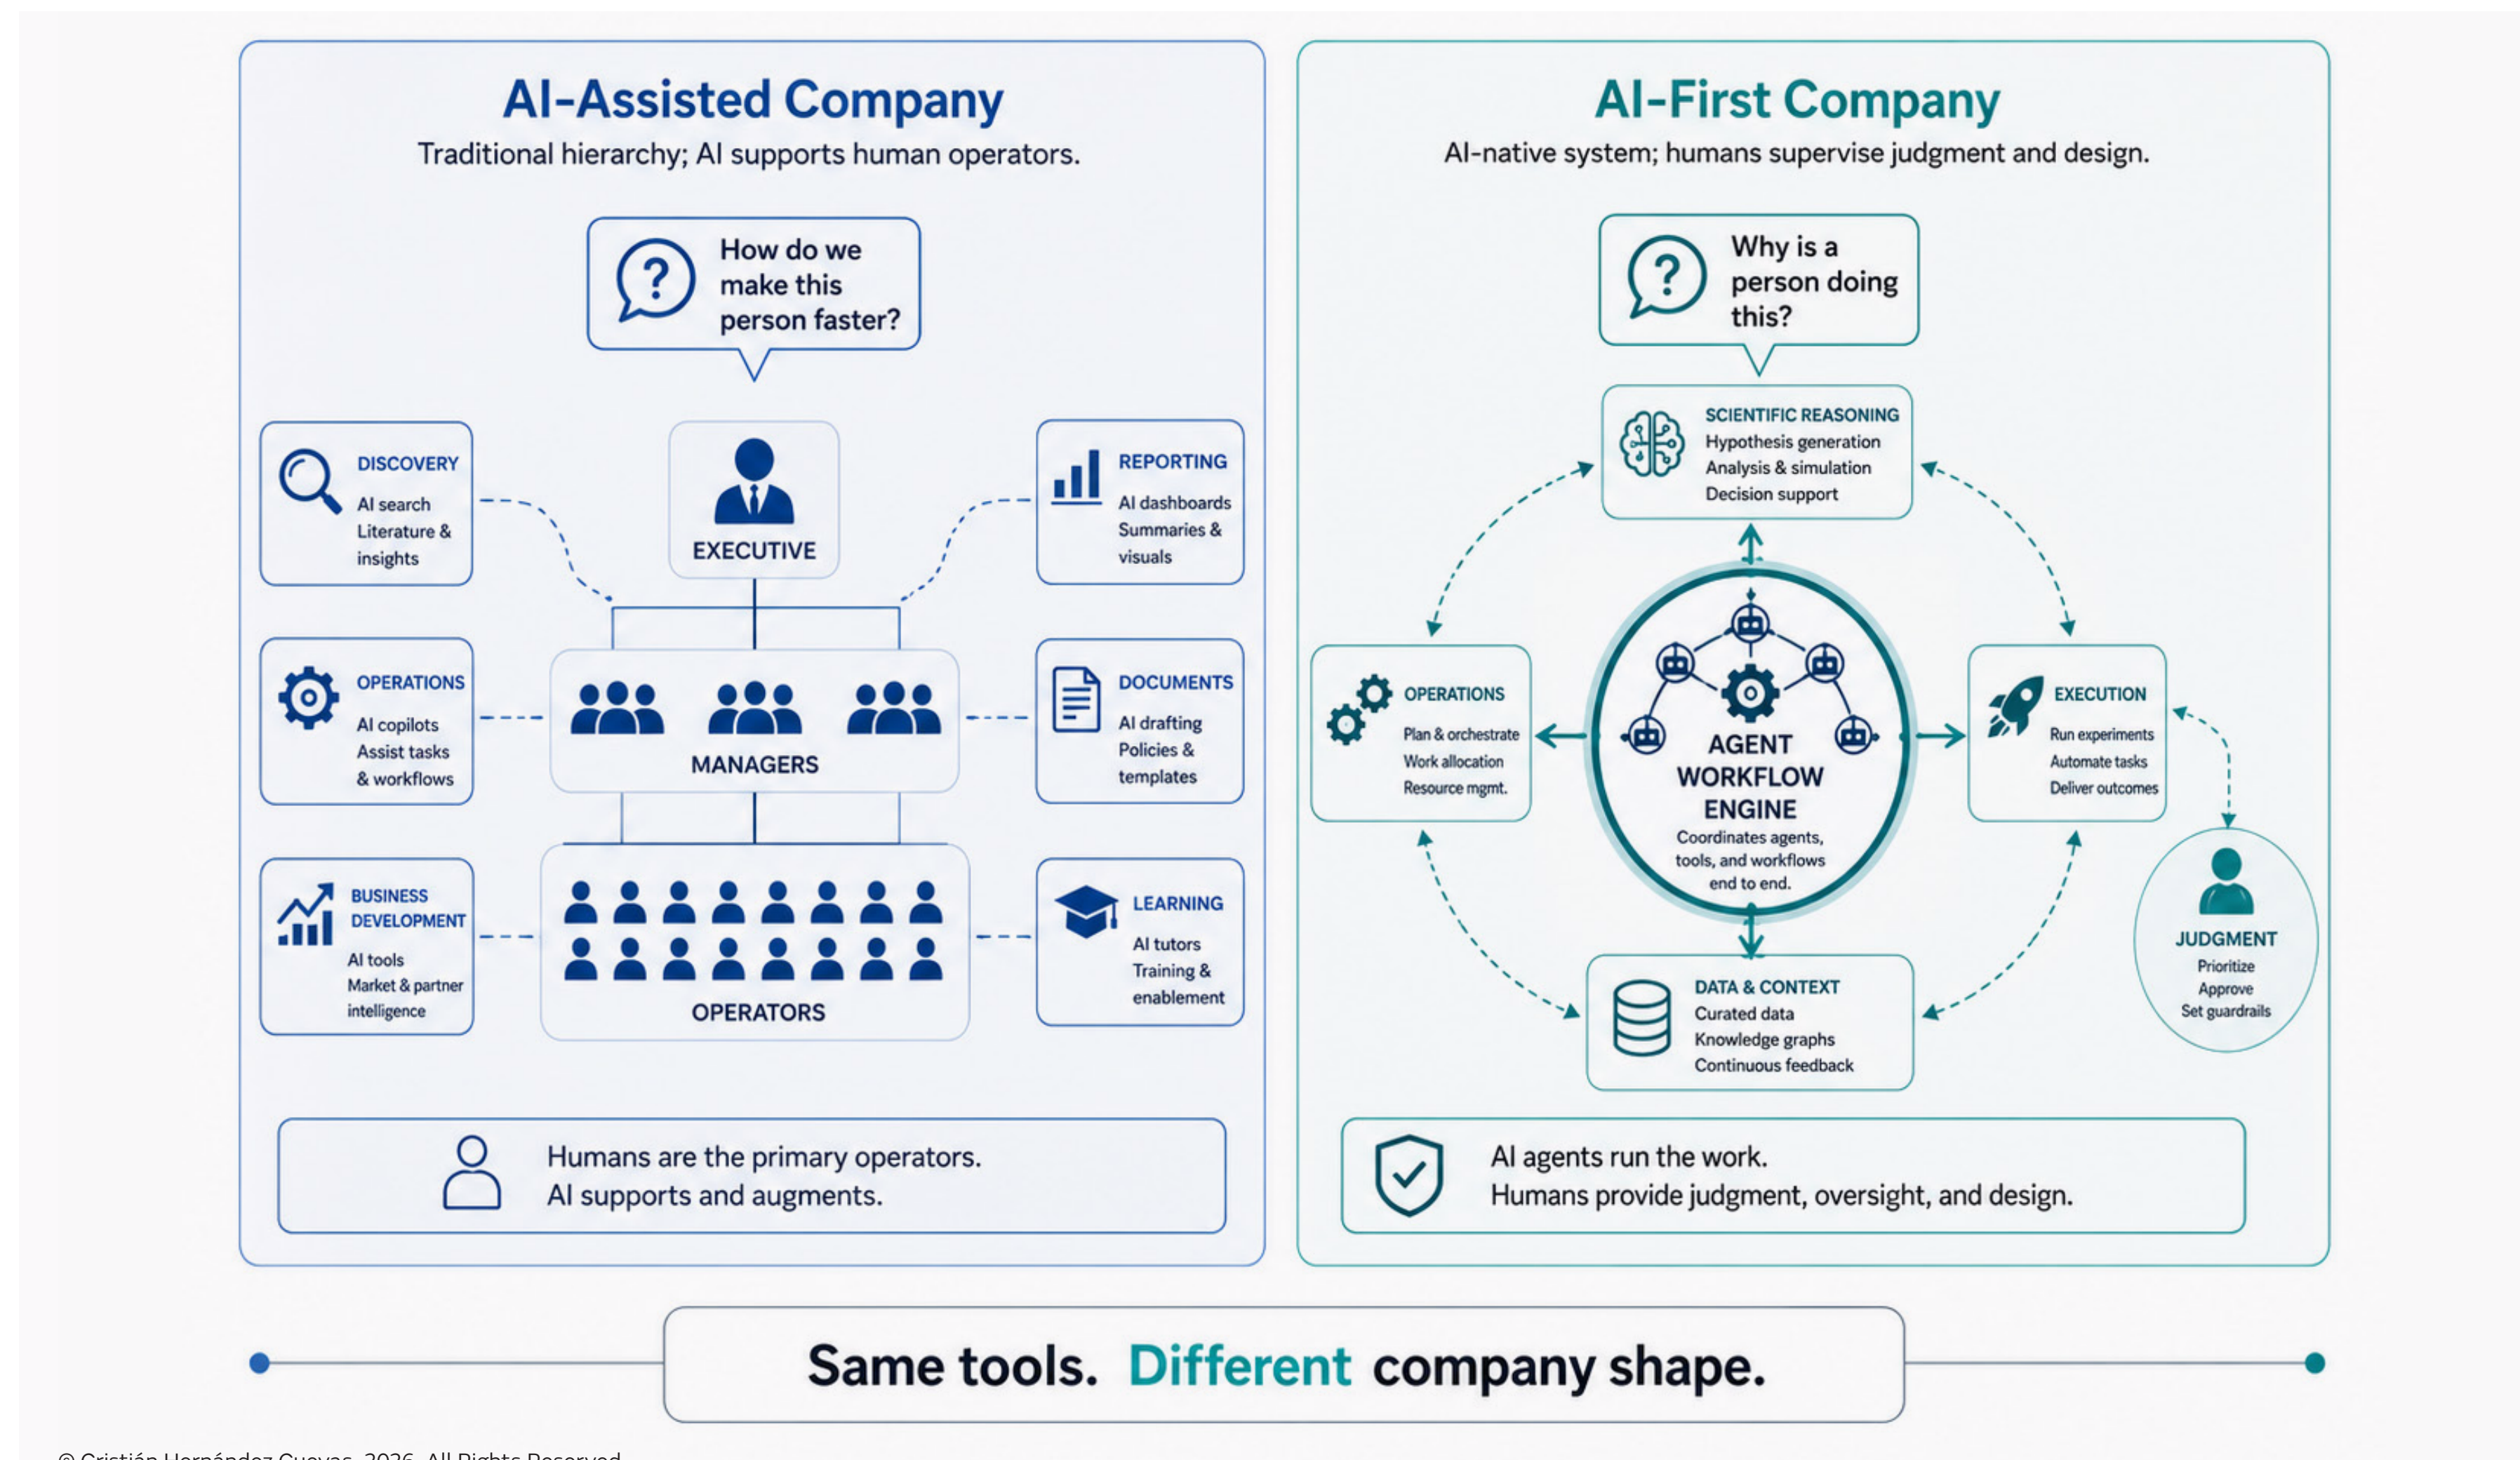
The same logic may also apply beyond therapeutics to areas such as agricultural biotech, One Health platforms, biomarker discovery, and regulatory intelligence. The common denominator is not the sector label, but whether the company can turn repeated learning cycles into a compounding operational advantage.

In traditional companies, humans remain the default operators and software assists them across fragmented parts of the organization. In AI first companies, automated systems increasingly become default operators, while humans move toward judgment, supervision, exceptions, and context.

The human role does not disappear, but it becomes more demanding because the premium moves from repetitive expertise to judgment under uncertainty. The strongest teams will define better questions, detect when models are wrong, and understand when speed creates risk rather than advantage.

This is why the AI transition is managerial, cultural, financial, and strategic, not merely technical or computational. Investors therefore need a sharper framework for distinguishing real architectural change from AI language added to an otherwise conventional company.

Figure 5. The operating question inverts when AI moves from assistant to default workflow operator, forcing companies to re-design roles, decision rights, feedback loops, and human judgment around a different architecture.



© Cristián Hernández Cuevas, 2026. All Rights Reserved.



7. Investors Will Underwrite Evidence, Not AI Language

For investors, the standard is changing because AI language is becoming easier to produce and therefore less useful as a signal of company quality. In biotech, we have always asked about the science, the team, the market, the regulatory path, and the capital required to reach the next milestone.

Those questions remain essential, but a new one must be added with much more discipline: how is this company actually built? The real question is not whether the deck contains a slide titled artificial intelligence, or whether the team references a model by name.

The real question is where AI is embedded in the operating system of the company, and which workflows are different because of it. Investors should ask which decisions are faster or better, which datasets become proprietary, and which feedback loops compound with use.

They should also ask which costs change, which bottlenecks remain stubbornly biological, and which claims can be validated through evidence. This is what investors will increasingly underwrite: evidence that the company learns faster, reduces experimental waste, and improves candidate selection.

Evidence should also show that the company can lower burn without lowering ambition, build data advantage over time, and improve decision quality. As AI language becomes cheaper, evidence becomes more valuable because it reveals whether the company has changed or only changed its vocabulary.

This evidence standard is relevant for founders as much as it is for investors, because it defines what must be prepared before a financing conversation becomes credible.



8. The Latin American Opening

Most global conversations about AI and biotech assume that value will concentrate in the same places it always has, including Boston and San Francisco. That may remain partly true, but we believe the AI transition creates a more interesting possibility because geography can matter differently.

Latin America has historically faced gaps in capital availability, later stage financing, global networks, specialized infrastructure, and exit visibility. Those constraints are real and we have lived them, yet they do not represent the full strategic picture of the region's biotech potential.

The region also has deep scientific talent, strong clinical and medical capabilities, biodiversity, agricultural complexity, and cost advantages. It also has increasingly sophisticated entrepreneurial ecosystems, even if the financing environment remains much more difficult than in the United States.

LAVCA reported that Latin America focused VC fundraising reached 548 million dollars across 52 funds in 2024, the lowest point in eight years. This does not mean the biotech capital gap has been solved, but it does show that the region is no longer starting from zero.

In the old biotech model, Latin America's advantages were often fragmented across universities, hospitals, clinical systems, biodiversity, and founder talent. AI may help connect these pieces by making distributed talent, differentiated datasets, efficient operations, and disciplined learning more valuable.

The opportunity is not to imitate Silicon Valley with less capital, because that would be the wrong ambition and the wrong strategic frame. The opportunity is to build a different kind of biotech company, combining global scientific ambition with Latin American operating advantages.

Efficient cost structures matter, as do differentiated biology, clinical execution, regulatory learning, agricultural complexity, and One Health contexts. AI does not automatically make Latin America competitive, but it may reduce the penalty of distance and reward teams that learn faster with fewer resources.

This is not only a company formation opportunity, because AI may also change how venture funds operate. Funds that can process scientific signals,

compare global landscapes, support capital readiness, and help founders translate complex science into investable evidence may become more valuable to the ecosystem.

For a region that has always had to build under constraints, disciplined learning is not only a survival skill but a potential strategic advantage. This is the field observation behind Zentynel's conviction that Latin America deserves a more serious role in the next biotech cycle.

Figure 6. The Zentynel AI Evidence Framework evaluates AI as measurable change in proprietary data, operating leverage, workflow impact, validation quality, and business outcomes, rather than as a label or marketing claim.

Value Signal	Proprietary Data	AI Role	Workflow Impact	Validation Evidence	Business Outcome
Reduced Experiments	Proprietary assays, omics, protocol history	Experiment design, active learning	30–70% fewer experiments	Retrospective benchmark, prospective validation	Lower R&D burn, faster iteration
Faster Candidate Selection	Multi-modal data, structure, biology, PK/PD	Ranking models, MOA prediction, prioritization	1–2x faster selection	Prospective candidates, hit-to-lead comparison	Shorter cycle time, earlier value inflection
Lower Cost-to-Serve	Automated pipeline, software analytics, shared data lake	Automation, workflow orchestration	20–50% lower cost per program	Unit economics, cost-per-decision metrics	Higher operating margin, more programs per team
Shorter Cycle Time	Integrated data + workflow history	Predictive timelines, bottleneck detection	25–60% faster cycle time	Before/after timelines, operational audit	Faster milestone completion
Better Clinical Triage	Patient data, biomarkers, outcomes	Patient stratification, endpoint prediction	1–2.5x better signal	Retrospective cohort, external validation	Higher PoS, lower capital wasted
Higher Closing Velocity	Decision logs, diligence history, market intelligence	Diligence copilots, signal synthesis	30–60% faster decision cycles	Decision quality review, investment outcomes	Better portfolio throughput



AI is a tool. Evidence is the signal. We underwrite what can be proven and repeated.

9. What Zentynel Is Learning From the Field

From our work across Latin American biotech company formation, venture investing, and cross border scientific commercialization, three lessons keep recurring. These lessons are not abstract conclusions, but patterns that emerge from working with founders, scientists, investors, and strategic partners across the region.

First, the most interesting companies are not always the ones with the most aggressive AI language, but the ones where AI changes the cost of learning. This leverage can appear in faster candidate generation, better experimental prioritization, automated patent mapping, improved clinical site selection, cleaner regulatory preparation, or better internal decision making.

The common thread is not the model itself, because the common thread is leverage embedded in the company's operating system. When AI is real, it changes how the company learns and how each cycle creates a better next cycle.

Second, the strongest founders are not using AI to avoid scientific depth, but to make scientific depth more operational and more testable. In weak companies, AI becomes a substitute for thinking, while in strong companies it becomes a forcing function for better questions.

It demands cleaner hypotheses, better structured data, more explicit assumptions, and feedback loops that expose whether the company is guessing or learning. That is the kind of discipline investors should reward.

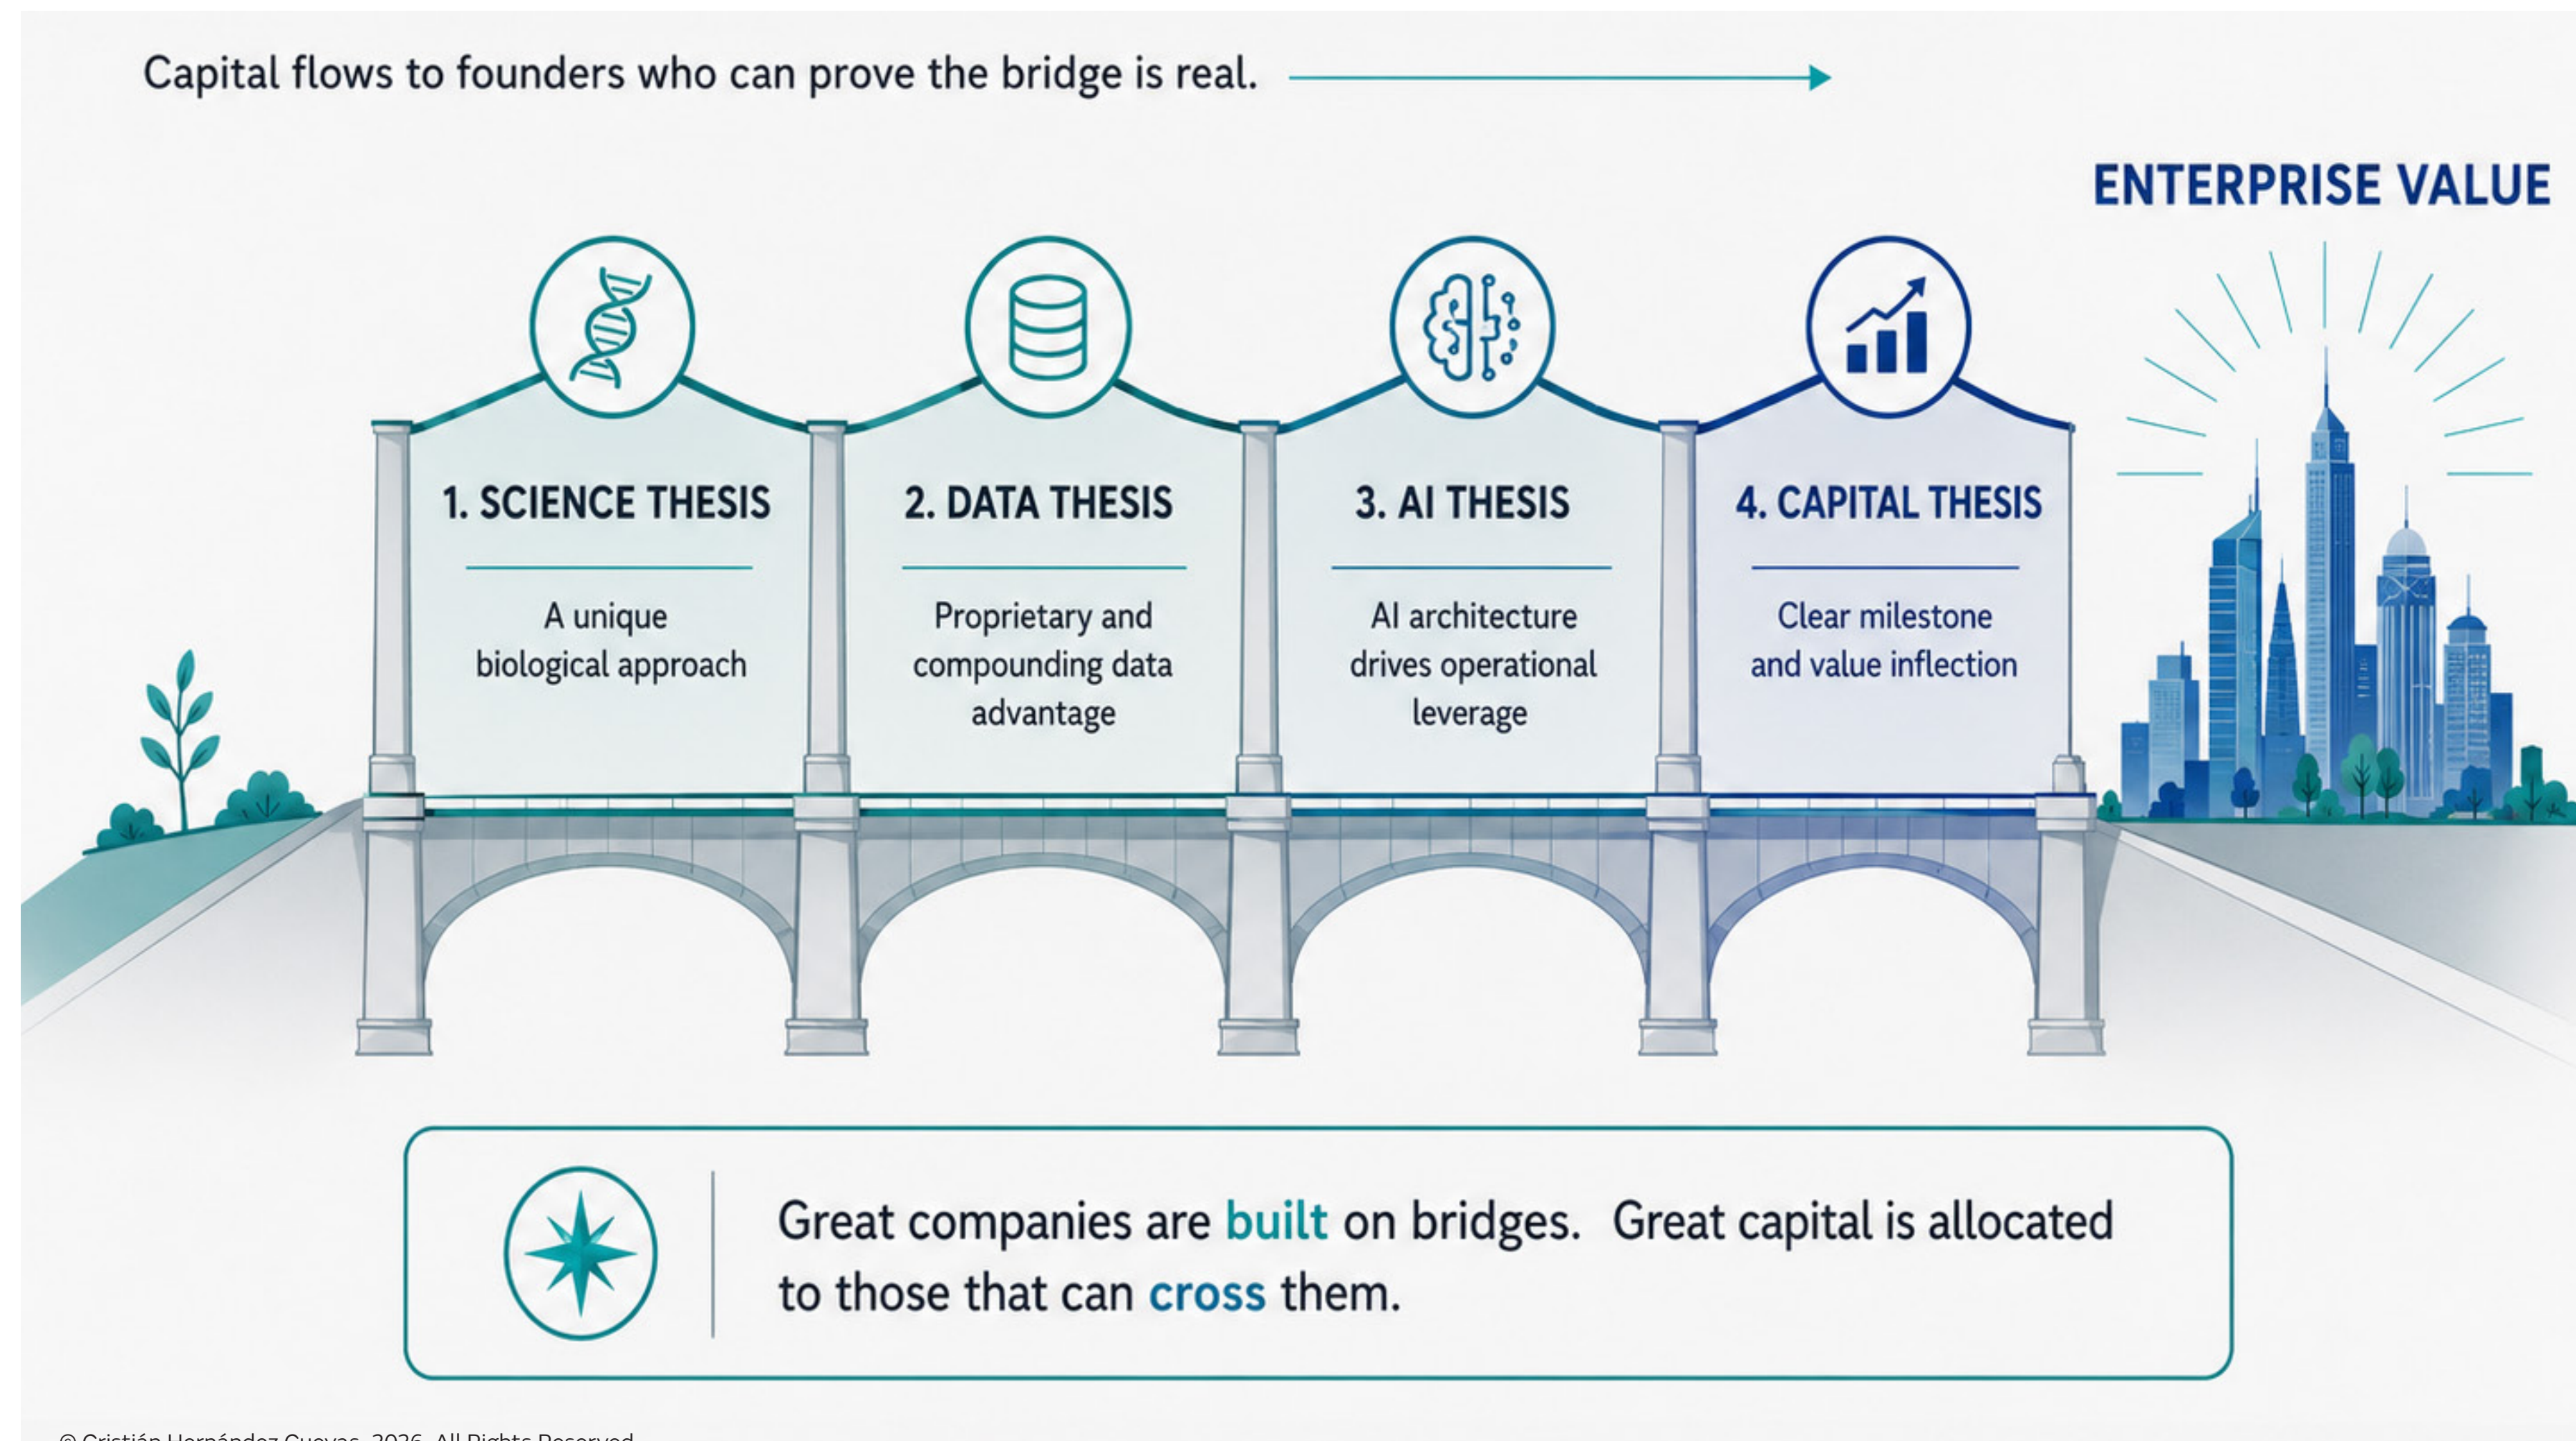
Third, capital efficiency is becoming a strategic language and not merely a financial constraint imposed by difficult fundraising conditions. Latin American founders have historically operated with smaller rounds, fewer specialized investors, and more

pressure to reach milestones with limited resources.

In the old biotech model, this scarcity was often a disadvantage because capital intensity could overwhelm even strong science and disciplined teams. In an AI enabled model, some of that discipline may become an advantage when paired with world class science, execution, and global ambition.

Constraint alone does not create excellence. But teams that already know how to learn efficiently may be better prepared than the market assumes, especially if AI compresses parts of the work that previously required large organizations and heavy capital intensity.

Figure 7. The capital bridge moves from science thesis to data thesis, AI thesis, and capital thesis, requiring founders to show how each financing round converts uncertainty into enterprise value.



10. What Founders Must Prepare to Raise Capital

For biotech founders, the implications are practical because raising capital now requires more than a compelling scientific story and a credible team. Founders need to explain how science, data, AI, operations, and capital strategy reinforce one another as a coherent learning system.

The first requirement is the scientific thesis, which should explain the problem, its timing, and the biological insight that makes the company credible. The second requirement is the data thesis, which should clarify what the company can access, generate, structure, or accumulate that others cannot.

The third requirement is the AI thesis, which should avoid model names and instead explain where AI creates real leverage inside the operation. That leverage may sit in discovery, design, experimental planning, clinical operations, regulatory preparation, manufacturing, market access, or decision making.

The fourth requirement is the capital thesis, which should define what milestone the round buys and why that milestone increases enterprise value. This last point is essential because fundraising is not merely selling a story, even if narrative remains important for alignment and conviction.

Raising capital is building a bridge, and the investor does not need to see the entire city if the next span is visible and can hold weight. The stronger the evidence behind that next span, the more credible the company becomes.



11. Do Not Prepare an AI Slide

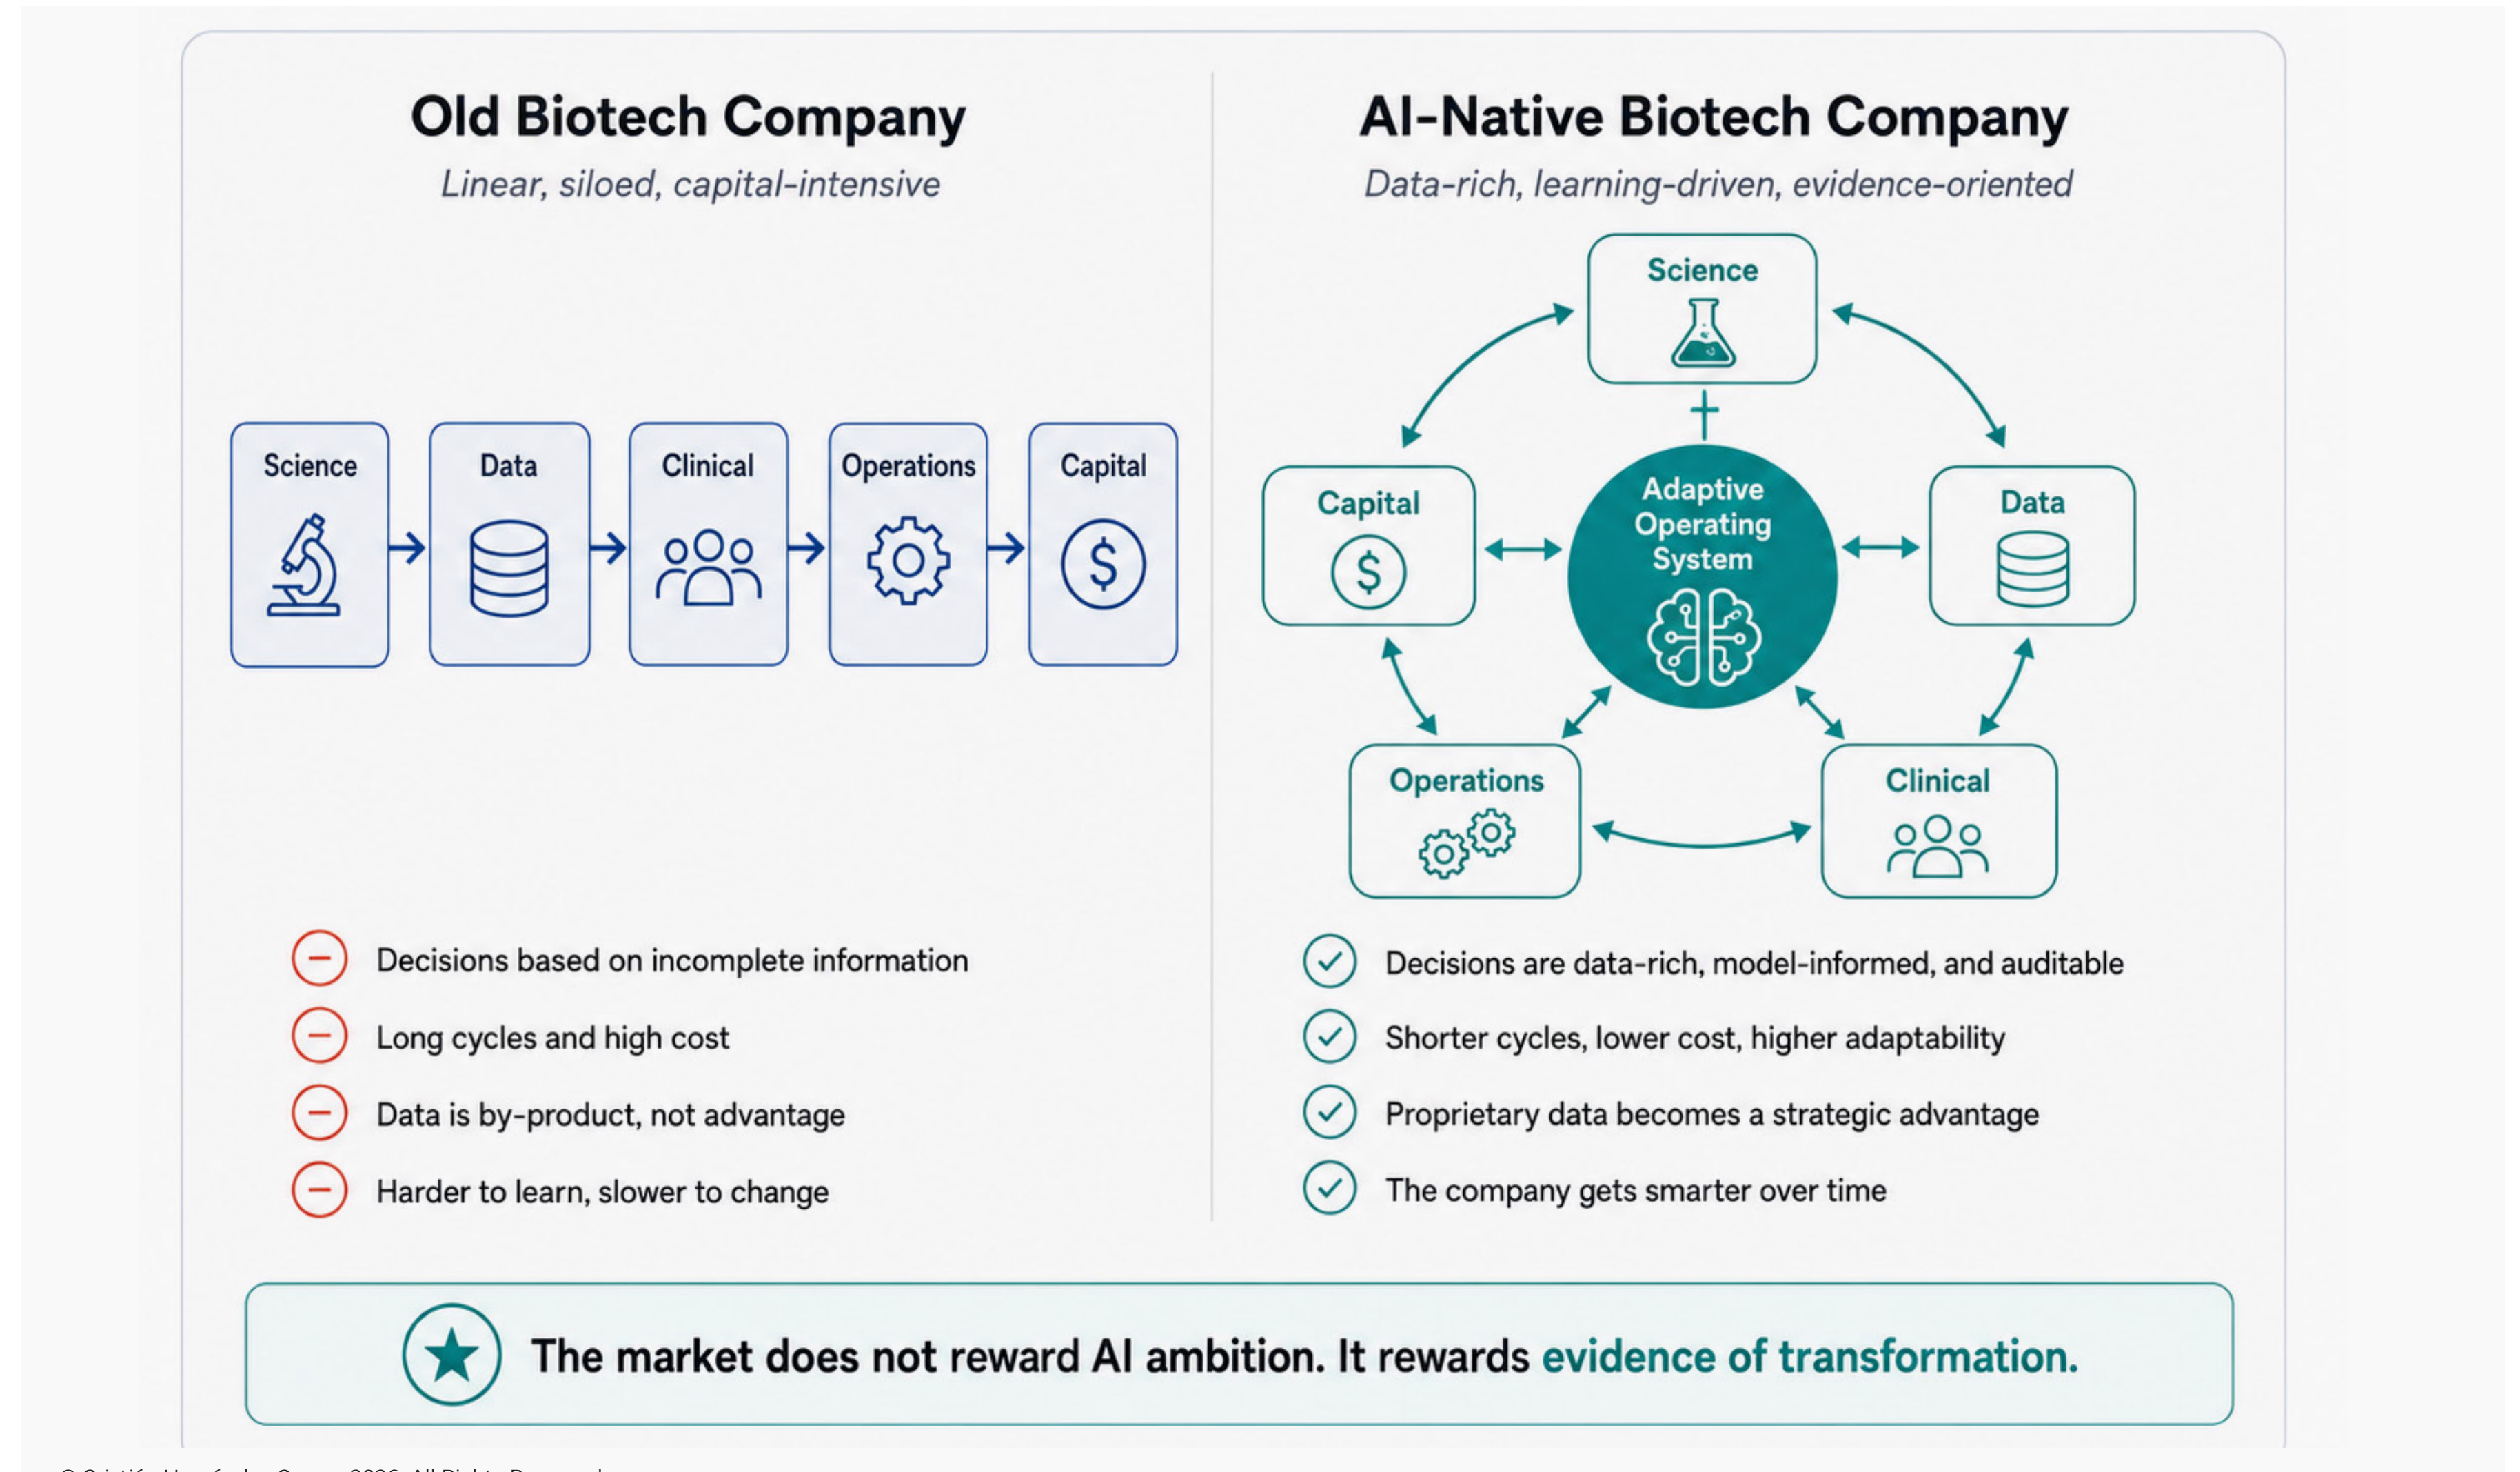
The practical message for founders is simple: do not prepare an AI slide, because the market will increasingly discount AI as a standalone claim. Prepare evidence that AI has changed how your company learns, decides, operates, and scales across the parts of the system that matter.

If AI changes burn rate, experimental velocity, decision quality, data advantage, or milestone efficiency, it becomes part of the investment case. If AI is only language in a deck, it will be treated as language in a deck and eventually filtered out by more disciplined investors.

The best companies will not say they use AI, because they will show something more uncomfortable and more powerful than a claim. They will show a company that could not exist without it, because AI has changed the structure of how the company learns and operates.

That is the new standard, and it is closer to AI enabled evidence than to AI ambition as a branding strategy. It also leads directly to the larger thesis of this paper: the coming decade will reward transformed architectures, not transformed vocabulary.

Figure 8. The shift from an old linear biotech company to an AI native biotech company is not cosmetic, because value emerges when AI changes decisions, data loops, operating leverage, and enterprise productivity.



11. Closing Thesis

The coming decade will not be divided simply between companies that use AI and companies that do not, because that distinction is too shallow. The real division will be between companies whose architecture was redesigned for an AI enabled world and companies built for a world that no longer exists.

In biotech, this matters more than in most sectors because the stakes are higher, uncertainty is greater, and the cost of poor learning is enormous. AI will not remove the risk of biology, but it will change who learns faster, fails better, and converts data into judgment with less friction.

It will also influence who can build globally relevant companies from regions that have historically been underestimated by the biotech capital markets. That is the opportunity Zentynel is watching closely, not because AI makes biotech easier but because it may make a new kind of biotech possible.

The company that matters will be built around biological reasoning, disciplined learning, capital efficiency, and global ambition from the beginning. It will not merely use AI, because AI will shape how the company learns, operates, and compounds advantage.

The next generation of biotech will not be defined by who uses AI, but by who turns biological uncertainty into a faster, cheaper, and more disciplined learning system. That is the standard we believe founders should prepare for, and it is the standard Zentynel intends to underwrite.

Acknowledgements

This paper was shaped by conversations with founders, scientists, investors, collaborators, and friends across the biotech ecosystem. Their questions, doubts, enthusiasm, and critical feedback helped clarify ideas that were still emerging through the work.

I am especially grateful to the people who continue to push Zentynel to think with more ambition, precision, and generosity about Latin America's role in the future of biotechnology. This paper is also a small act of biogratitude toward that community.



Cristián Hernández Cuevas

Zentynel General Partner

Cristián Hernández Cuevas is General Partner at Zentynel Frontier Investments, where he invests in and supports biotech, healthtech, and deep science companies across Latin America and global markets. His work sits at the intersection of science, venture capital, company formation, and cross border commercialization.

A Molecular Biotechnology Engineer from the Universidad de Chile, Cristián also earned a Master's in Bioscience Enterprise from the University of Cambridge. Over three decades, he has worked as a scientist, founder, operator, advisor, and investor, helping translate scientific insight into companies, partnerships, and investment strategies.

Through Zentynel, he focuses on bridging capital gaps, supporting founders, and positioning Latin America as a serious contributor to the global biotech landscape. His work emphasizes scientific ambition, disciplined learning, and the development of investment frameworks suited to emerging innovation ecosystems.



Selected Sources

The Nobel Prize in Chemistry 2024, Nobel Prize press release, October 2024.

Isomorphic Labs, announcement of 600 million dollar external investment round, March 2025.

Reuters, Isomorphic Labs 2.1 billion dollar Series B coverage, May 2026.

NVIDIA and Eli Lilly, co innovation AI lab announcement, January 2026.

lambic Therapeutics, collaboration announcement with Takeda, February 2026.

PhRMA, Biopharmaceuticals in Perspective and Research and Development materials.

BIO, Informa Pharma Intelligence, and QLS Advisors, Clinical Development Success Rates 2011 to 2020.

LAVCA, Trends in Tech 2025 report, February 2025.



www.zentyne.com