BioSpectrum
the business of Bio & Health Sciences
ASIA EDITION

The Model Wants a Lab Coat

November 5, 2026 | Opinion | By Ankit Kankar | ankit.kankar@mmactiv.com

Anthropic spent 2026 turning Claude from a clever assistant into a colleague, a contractor and now a drug developer in its own right. The harder question for 2027 is how much of life sciences it will actually shape, and who gets to decide.



Start with the smallest story of the year, because it may turn out to be the most important. Until this summer, Mithril Medicine was a biotech company with exactly one full-time employee. Its founder, Cody Tranbarger, was chasing treatments for two rare forms of muscular dystrophy. He had no research department in the usual sense. What he had was Claude, which he used to coordinate contracts, logistics and experiments across 11 outside laboratories. In under 18 months, one candidate went from idea to encouraging results in mice.

In July, Anthropic bought the company. Tranbarger now works inside Anthropic, and the firm says a human trial could begin as early as late 2028 if the candidates keep holding up.

Sit with that for a moment. The company that makes the model has acquired a drug programme that the model helped run. A software vendor now owns preclinical assets. That single transaction says more about the next phase of AI in life sciences than any benchmark, and it raises the question this essay tries to answer: in 2027, how much of our industry will Anthropic actually shape, and on whose terms?

To answer that honestly, we need to start not with Anthropic, but with the problem it says it wants to solve.

The problem underneath: why drug R&D keeps getting harder

For decades, pharma has lived with an uncomfortable pattern sometimes called Eroom's law, Moore's law spelled backwards: the number of new drugs approved per billion dollars of R&D has fallen steadily since the 1950s, even as every individual tool got better. The most widely cited estimate puts the cost of bringing one new medicine to market at around USD 2.6 billion once failures are included, over a journey of ten to fifteen years. Roughly nine in ten candidates that enter human trials never reach patients.

The crucial detail is where they fail. Chemistry is rarely the killer. Most expensive failures happen in Phase II, when a well-made molecule hits the right target and the patient simply does not get better. That is a biology problem: the wrong target,

the wrong patients, the wrong hypothesis about disease. It is also the most expensive place to be wrong, because by then the programme has consumed years and hundreds of millions of dollars.

This gives us a useful lens for judging any AI claim. An AI that designs molecules faster is valuable. An AI that helps us pick better targets and design better trials is transformative. The early evidence makes the distinction sharp. A 2024 analysis by BCG, published in Drug Discovery Today, found that AI-discovered molecules succeeded in Phase I 80% to 90% of the time, against a historical 40% to 65%. In Phase II, on a small sample, they succeeded at about 40%, roughly in line with the industry. AI, so far, is making safer, better-behaved molecules. It has not yet proven it can make biology more predictable.

AI has made drugs easier to make. It has not yet made biology easier to predict. That gap is where 2027 will be decided.

Keep that gap in mind. It is the yardstick against which Anthropic's strategy, and its rivals', should be measured.

Four identities in twelve months

Anthropic's path through 2026 is best read as a sequence, because each step changed what kind of company it is in our industry.

Anthropic's life sciences timeline

When	Move	What it signalled
Oct 2025	Claude for Life Sciences: connectors to Benchling, PubMed, 10x Genomics, Synapse.org and Wiley's Scholar Gateway	Toolmaker: plug into existing lab software
Feb 2026	HHMI and the Allen Institute become founding partners	Credibility with elite academic science
Apr 2026	Acquires Coefficient Bio (reported at about USD 400 million)	Buying in-house drug discovery talent
May 2026	Bristol Myers Squibb brings Claude across major business functions	Enterprise anchor in big pharma
Jun 2026	Claude Science workbench (60+ databases, compute, reviewer agents) and an in-house neglected-disease programme	Workflow owner, and first step into its own R&D
Sep 2026	Novo Nordisk collaboration; Bay Area wet lab confirmed; verification programme for researchers	From simulation to physical experiments
Oct 2026	Results from six drug development partners; Mithril Medicine acquisition disclosed	Participant: owns preclinical assets

Toolmaker, then workflow owner, then embedded partner, then participant. Each step pulls Anthropic closer to the molecule and further from the chat window.

Anatomy of the bet: three layers, one playbook

It helps to see Anthropic's life sciences effort as three stacked layers, each with a different business logic.

The reasoning layer. This is the model itself: Claude reading literature, writing analysis code, proposing hypotheses. It is horizontal, meaning the same model serves a bank and a biotech. Anthropic has been explicit that Claude Science runs on its general models rather than a special biology model. That is a deliberate bet that general reasoning, not narrow biological pattern-matching, is the scarce ingredient.

The workflow layer. Claude Science pulls databases, notebooks, compute and reviewer agents into one auditable environment. This is where Anthropic is copying its own most successful playbook. Claude Code won developers not because the model was uniquely clever, but because it sat inside the daily workflow and became hard to leave. Claude Science tries to do the same for the scientist's day. Workflow is where switching costs live, and where provenance, the record of who did what and why, gets captured.

The vertical layer. A wet lab, acquisitions, and internal drug programmes. This layer exists partly to learn: Anthropic says working inside real programmes is the fastest way to find out where Claude breaks. But it also changes the company's position. A vendor that owns molecules is, at least potentially, a competitor.

The Mithril insight: coordination is the real cost

The most interesting thing about Mithril is not the mouse data. It is the organisational shape. A virtual biotech that outsources work to contract labs is nothing new. What is new is one person running eleven labs at once, with the model handling the contracts, logistics and experimental bookkeeping that normally require a programme management team.

Economists will recognise the idea. Ronald Coase argued nearly a century ago that firms exist because coordinating work through the market is expensive. When coordination gets cheaper, firms get smaller. If AI collapses the cost of coordinating CROs, CDMOs and specialist labs, the minimum viable drug company shrinks from a department to a team, and sometimes to one founder. That is a structural shift, and it matters enormously for Asia, which supplies much of the outsourced work these tiny companies would coordinate.

What the early numbers say, and what they don't

Anthropic's October update was the first time it put numbers on the table rather than adjectives. Here is what each partner reported, and which part of the R&D problem it actually addresses.

Partner results disclosed on 27 October 2026

Partner	Stage	Reported result	Addresses the Phase II pro
Generate Biomedicines	Modality design	CAR-T component redesigned; one version about 70% shorter, roughly 10x more cancer-detecting receptors in lab tests	Indirectly: aims to extend CA solid tumours
Divergence Labs	Antibody design	About 35,000 antibody designs against IgE (food allergy) in roughly a month	No: speeds molecule genera
BigHat Biosciences	Lead optimisation	Claude plus lab automation; binding improved more than tenfold on two antibody fragments	No: improves molecule quali
Enveda	Manufacturing and regulatory documents	Found all 11 human-flagged issues plus 7 missed in 1,061 pages; 89% less review time; a three-month draft in four hours	No, but cuts time and cost a every stage

Partner	Stage	Reported result	Addresses the Phase II pro
The Column Group / Pleiades	Target selection	First Claude skill to assess drug targets, tested on 45 target-disease pairs	Yes: target choice drives Ph outcomes
Bristol Myers Squibb	Clinical planning	Claude evaluated early trial plans and estimated probability of success; expansion planned for 2027	Yes: better trial design and g go calls

Read the last column carefully. The headline-grabbing results, tens of thousands of antibodies and tenfold binding gains, are about making molecules. Valuable, but that is the part of R&D AI was already good at. The two partnerships that target the expensive failure point, target selection at Pleiades and trial planning at BMS, are the least visual and the least developed. They are also the ones that will matter most.

Then there is Enveda's documentation result, which I would argue is the most underrated number of the year. Molecule design gets the headlines, but documentation is where programmes quietly lose months at every stage, from IND filings to batch records. Every CDMO, regulatory affairs team and quality function in Asia understands that arithmetic instantly.

Now the caveats, which are real. Almost everything here is preclinical. The results are company-reported and not peer reviewed. Nothing has been tested in a human being. Anthropic's own head of life sciences, Eric Kauderer-Abrams, said in September that the final test of biology remains real lab work, and will be for a while. A tenfold improvement in a dish is a hypothesis about a medicine, not a medicine.

Anthropic is not alone: four theories of AI drug discovery

It would be a mistake to analyse Anthropic in isolation. 2026 was the year every major AI player declared its theory of how AI will change drug development, and they are strikingly different.

Competing approaches in 2026

Player	Theory	2026 moves	Owns molecules?
Anthropic	General reasoning model plus agentic workflow; learn by doing real programmes	Claude Science, wet lab, six partners, Coefficient Bio and Mithril acquisitions	Yes, small internal progra
OpenAI	Purpose-built science model plus agents that link findings to development decisions	Rosalind brand and GPT-Rosalind; Novo Nordisk partnership (Apr); Orchestra Bio acquisition (Oct)	Not so far
Isomorphic Labs (Alphabet)	The model company as a drug company	USD 2.1 billion raise (May); first human trials targeted for late 2026 after a delay; plans to license after early trials	Yes, core strategy
Nvidia with Eli Lilly	Compute and infrastructure; pharma owns the science	Up to USD 1 billion co-innovation lab over five years (Jan); Lilly TuneLab models for biotechs	Lilly does

Player	Theory	2026 moves	Owns molecules?
Insilico Medicine and China's AI-native biotechs	End-to-end AI pipeline company	Rentosertib enters Phase III in China (Jul), about 320 IPF patients at 47 centres	Yes, core strategy

Three observations follow.

First, big pharma is refusing to marry anyone. Novo Nordisk signed with OpenAI in April and with Anthropic in September. That is not indecision; it is strategy. Pharma companies are treating frontier models the way they treat CROs: as competing suppliers to be benchmarked, multi-sourced and swapped. If that holds, the durable advantage will not be the model. It will be how deeply each vendor embeds in workflow, data and decision records.

Second, the closest thing to proof is coming from Asia, not Silicon Valley. Insilico's rentosertib, designed with AI and now in a Phase III trial in China, is the most advanced test of whether AI-discovered drugs can clear the Phase II and III hurdles that matter. Its readout, whenever it comes, will move sentiment for the entire field, including Anthropic's customers.

Third, OpenAI and Anthropic are converging on the same insight. OpenAI bought Orchestra Bio to build agents that connect scientific findings to studies and decisions. Anthropic bought Mithril, whose founder used Claude to coordinate eleven labs. Both are betting that the next frontier is not smarter molecules but smarter orchestration. That is the Coase argument again, and it is now the consensus among the two leading model labs.

Where Anthropic is genuinely distinctive is in combining all three layers with an unusually public safety posture. That combination is both its selling point and its biggest constraint.

The safety paradox at the centre of the strategy

Anthropic is, by its own account, the AI developer most worried about biological misuse. It is also the one leaning hardest into biology. That tension went public in June.

When Anthropic released Claude Fable 5, its most capable general model, it wrapped biology and chemistry queries in safety classifiers so cautious that they rerouted ordinary questions to an older model. Reviewers found that even a question about what mitochondria do could trigger the fallback. Anthropic admitted it had chosen to be deliberately over-conservative. By September, the company said its updated safeguards intervened on benign biology requests 85% less often. Meanwhile, the unrestricted sibling, Mythos, remains reserved for vetted partners and approved biology researchers.

Commercially, this is a real cost: every falsely blocked question is a reason for a scientist to try a competitor. Strategically, it may be an asset. Pharma companies operate in the most regulated industry on earth, and a vendor that can document how it prevents misuse may find procurement and compliance conversations easier, not harder.

Verification becomes a credential

The deeper implication is that access to the most capable biology model is turning into a credential, in the way a GMP certificate or a BSL-3 facility is today. In 2027, being verified may decide which labs get the best tools. That raises questions our region should be asking now. How are labs in Hyderabad, Seoul, Suzhou or Singapore assessed? How long does verification take? Does it favour large institutions with compliance departments over the startups that drive so much of Asia's innovation? A tiered-access world is defensible on safety grounds, but it is not neutral in its effects.

There is also a reputational edge. When the wet lab became public in September, some commentators asked why a company warning about AI risk was running physical biology experiments. Anthropic says the lab focuses on fundamental biology under human supervision. That answer will be tested every time the lab's output becomes more ambitious.

The money underneath the mission

Anthropic filed publicly for its Nasdaq listing on 26 October. The prospectus shows first-half 2026 revenue of USD 16.3 billion, up from USD 1.2 billion a year earlier, and names life sciences as a frontier market, noting that global pharmaceutical spending is expected to exceed USD 2 trillion this year.

The same filing carries the sober counterweight. Among nearly 100 pages of risk factors, Anthropic warns that model errors could harm patients, and that regulators may decide some of its products fall under healthcare regulation, which could force it to restrict offerings or absorb significant compliance costs. It also shows concentration: two unnamed customers accounted for about 23% of first-half revenue. And it records how quickly policy can bend the curve: US export controls briefly forced Anthropic to suspend its two leading models in June, before access was restored on 1 July.

Public market investors will want life sciences to become a visible revenue line rather than a research showcase. That pressure cuts both ways. It means more investment, more partnerships and faster product cycles in 2027. It also means a company built around caution will be under quarterly pressure to move faster in the most safety-sensitive domain it touches. Watch whether the tone of Anthropic's life sciences announcements shifts from "what we learned" to "what we sold."

The regulator question

No amount of model capability matters if regulators will not accept its outputs. Here the picture is moving, but slowly. On 14 January 2026, the US FDA and the European Medicines Agency jointly published ten guiding principles of good AI practice in drug development, covering nonclinical, clinical, manufacturing and post-market work. Their themes are risk-based assessment, human-centred design, data governance and lifecycle management. They are principles, not binding guidance, but they signal the direction of travel.

Two implications stand out. First, the regulatory emphasis on traceability plays directly to the workflow layer. Claude Science's focus on auditable artifacts and provenance reads like an answer written in advance for exactly this scrutiny. Whoever records the reasoning behind an AI-assisted decision best will have an edge in regulated settings.

Second, Asia risks a patchwork. Japan's PMDA, India's CDSCO, Singapore's HSA, Korea's MFDS and China's NMPA will each need a view on AI-drafted dossiers, AI-designed molecules and AI-informed trial designs. If those views diverge, companies serving multiple Asian markets will face duplicated validation work, eroding the very efficiency AI promises. Harmonisation through ICH and regional regulatory dialogue is not a nice-to-have in 2027; it is a competitiveness issue.

Asia, market by market

Asia is not one story. The opportunity and the risk look different in each major market.

India: the services model meets the 89% number

India is already Claude's second-largest market, and Anthropic opened its Bengaluru office in February, its second in Asia after Tokyo. That footprint was built around enterprise, education and social impact, but the logic points straight at life sciences. India's strengths in contract research, contract manufacturing, generics and the global capability centres of multinational pharma are exactly where documentation-heavy work concentrates.

Here is the uncomfortable part. Much of India's life sciences services income is priced on effort: hours, FTEs and page counts. An 89% cut in document review time is wonderful news for a sponsor and an existential question for a business that bills by the hour. Indian CROs and CDMOs that reprice around outcomes, speed and quality, and build AI-native delivery, will win share. Those that defend the billable hour will be squeezed between clients who expect AI savings and AI-native competitors who can offer them.

China: a parallel ecosystem, moving fast

China is building its own AI-native drug industry, and it is not waiting for anyone. Chinese biotechs signed licensing deals worth about USD 93 billion with global partners in the first five months of 2026, according to HSBC figures reported by the South China Morning Post, and AI drug-design companies are part of that wave. Insilico's Phase III trial is running in China. This ecosystem will be shaped as much by geopolitics as by science, which means the question for 2027 is less

whether Chinese and Western AI drug discovery compete, and more whether their outputs can still flow across borders through licensing.

Japan and Korea: scale to negotiate as peers

Japanese and Korean pharma and biomanufacturing groups have the scale, data and balance sheets to negotiate with frontier AI companies as peers rather than as customers. Anthropic has had a Tokyo office since October 2025. The strategic opportunity is to co-develop rather than simply license: to make sure Asian clinical data, Asian patient populations and Asian manufacturing know-how are shaping these tools, not just consuming them.

Singapore and the smaller hubs: the trust layer

For Singapore and similar hubs, the opening is in trust infrastructure. If verified access becomes the gateway to the most capable biology models, there is room for a credible regional verification, biosecurity and data governance hub. That is a role that plays to Singapore's regulatory reputation and could keep Asian startups from waiting in a queue designed elsewhere.

Three scenarios for 2027

Forecasting a fast-moving company a year out is humbling, so here are three scenarios rather than one prediction, with my rough sense of their likelihood.

How 2027 could play out

Scenario	What happens	What would signal it
Base case: the plumbing year (most likely)	Claude becomes standard in analysis, documentation, regulatory drafting and design iteration across many large pharma firms. Clinical decision work expands cautiously under contract. Internal programmes progress quietly.	More enterprise deals; partner case studies time saved; no major clinical claims
Upside: the decision year	BMS-style trial assessment influences real go/no-go calls. Pleiades shows a small team can run many programmes. Regulators accept AI-drafted dossier sections with defined human sign-off. More one-person biotechs appear.	Named programmes advanced or stopped informed analysis; regulator statements on drafted submissions; Pleiades financing
Downside: the friction year	A safety incident, an over-broad safeguard, or a regulator classifying products as health software slows adoption. Partners balk at conflicts with Anthropic's own programmes. IPO pressure pushes faster commercialisation and erodes trust.	Partners pausing deals; new contract claus data walls; regulatory inquiries; tightening biology access

What each player should do now

Pharma leaders. Multi-source, as Novo has, and benchmark models against your own historical decisions, not vendor demos. Negotiate data walls and IP terms before deep integration, especially now that model makers are running their own programmes. Point AI at target selection and trial design, where failure is most expensive, not just at molecule generation.

CROs and CDMOs. Assume clients will expect AI savings to show up in your pricing within two years. Reprice around outcomes and speed. Make your systems easy for agents to work with, because the one-person biotech will choose suppliers its AI can coordinate efficiently.

Biotech founders. The Mithril model is now a template. A small team plus disciplined use of AI and outsourced labs can reach preclinical proof faster and cheaper than before. The scarce skill is no longer headcount; it is scientific judgement about which hypothesis deserves the shot.

Investors. Capital efficiency per programme may improve sharply, which changes portfolio maths. Ask founders how much of their programme management runs on AI, and how they validate it. Watch the Insilico Phase III readout as the field's bellwether.

Regulators and policymakers in Asia. Coordinate positions on AI-assisted submissions early, ideally regionally. Engage frontier labs on how verification and tiered access work for Asian institutions. Build the trust infrastructure that lets local companies access the best tools safely.

Signals to watch in 2027

1. **The BMS expansion.** Does Claude's trial assessment move into real early clinical development decisions, and does either company say so publicly?
2. Does a Claude-run, small-team biotech raise serious capital and advance multiple programmes?
3. **The first peer-reviewed results.** Do any of the October partner results appear in peer-reviewed journals, with independent replication?
4. **Verification at scale.** How many institutions, and how many outside the US, gain trusted biology access?
5. **Partner terms.** Do pharma contracts start including explicit walls between partner data and Anthropic's internal programmes?
6. Any Phase III signal from Insilico will reset expectations for AI-discovered drugs across the industry.

So, how much will Anthropic shape life sciences in 2027?

My honest assessment is this. Heavily in the plumbing: analysis, documentation, regulatory drafting and design iteration, where the gains are already measurable and the risks are manageable. Moderately in decisions: trial planning, target selection and portfolio choices, where adoption will be cautious and contractual. Barely, if at all, in approved medicines: by Anthropic's own timeline, its earliest internal candidate is not expected near the clinic until late 2028 at best.

In other words, 2027 will be the year Anthropic shapes how drugs are made far more than which drugs are made. That is not a small thing. If the Coase argument is right, the most important change hiding in this year's announcements is not any single molecule. It is that the minimum viable unit of a drug programme is shrinking, from a department to a team, and in Mithril's case to one determined founder with a model and a list of contract labs.

But the yardstick from the start of this essay still applies. Until AI makes biology more predictable, until it moves the Phase II needle, the revolution will be one of speed and cost rather than of success rates. Speed and cost are worth a great deal. They are not yet the same as more medicines for patients.

The question for boardrooms is no longer whether to use AI. It is which of your own capabilities you are willing to hand to a vendor that is now, quietly, also becoming a peer.

-