


When the Toolmaker Decides to Also Make the Drugs: Anthropic's Claude Science Launch, Its In-House Preclinical Bet on Neglected Diseases & What Tool Vs. Competitor Tension Means for Pharma AI

JUL 06, 2026



Share

 **Part I Podcast free on Spotify.**



Thoughts on Healthcare Markets & Technology



Part I: When the Toolmaker Decides to Also Make the Drugs: Anthropic's Claude Science Launch, Its In-House Preclinical Bet on Neglected Diseases & What Tool Vs. Competitor Tension Means for Pharma AI

Anthropic launched a science research platform AND announced its own drug programs on the same day. The toolmaker opened a mine next to its customers' mines...

▶ Listen now

15 minutes ago · Thoughts on Healthcare

 **Part II Podcast episode for paid subscribers only. Also available on Spotify.**



Thoughts on Healthcare Markets & Technology



Part II: When the Toolmaker Decides to Also Make the Drugs: Anthropic's Claude Science Launch, Its In-House Preclinical Bet on Neglected Diseases &

What Tool Vs. Competitor Tension Means for Pharma AI

Anthropic launched a science research platform AND announced its own drug programs on the same day. The toolmaker opened a mine next to its customers' mines...

▶ Listen now

12 minutes ago · Thoughts on Healthcare

To listen to paid episodes in Apple or Spotify, link settings on those platforms (instructions inside the Podcast).

Thoughts on Healthcare
Technology is a regular
publication. To receive
support my work, consider
free or paid subscription



Discover more from Thoughts on Healthcare Markets & Technology

Expert analysis of healthcare and life sciences markets, technology, investment, entrepreneurship policy, and AI — for investors, entrepreneurs,...

Over 4,000 subscribers

Enter your email...

Subscribe

By subscribing, you agree Substack's [Terms of Use](#), and acknowledge its [Information Collection Notice](#) and [Privacy Policy](#).

Already have an account? [Sign in](#)

Table of Contents

1. Why a model company suddenly cares about wet labs
2. What Claude Science actually is under the hood
3. The part where the toolmaker becomes the competitor
4. Neglected diseases, public benefit, and the convenient economics of doing good
5. The competitive board, from DeepMind to OpenAI to the techbio pile
6. Where the skeptics are right, which is the target problem
7. The money math and the IPO clock ticking in the background
8. What is actually worth watching

Abstract

On June 30, 2026, at an “AI for Science” event in San Francisco, Anthropic put a new product called Claude Science on stage next to its other agent flagships and, in the same breath, said it would start running its own preclinical drug programs aimed at neglected diseases. So the company that sells the shovels now wants to pan for gold too, standing next to the customers it just sold shovels to. This piece walks through what the product is, why the dual announcement is stranger than it sounds, how it

stacks against DeepMind, OpenAI, and the pure-play techbios, where the honest criticism lands, and what any of it means for a company staring down a profitable quarter and an IPO.

Thoughts on Healthcare Markets &
Technology is a reader-supported
publication. To receive new posts and
support my work, consider becoming a
free or paid subscriber.

Quick numbers worth holding onto:

1. Pharma spends roughly 150 to 200 billion dollars a year on R&D and has produced somewhere around 800 to 1,000 approved drugs in 120 years.
2. Claude Science integrates 60-plus scientific databases and functions and runs on a lab's own private infrastructure.
3. Isomorphic Labs raised a multibillion-dollar round in May 2026, among the largest ever in AI drug discovery, and has still dosed zero patients.
4. A Boston Consulting Group look at about two dozen AI-discovered molecules found Phase 1 safety success jumping to 80 to 90 percent versus a historical 50, while Phase 2 efficacy fell right back to the usual 40.
5. Anthropic is offering up to 30,000 dollars in credits across roughly 50 research projects, applications closing July 15.

Why a model company suddenly cares about wet labs

The setup is worth pausing on, because on paper it makes very little sense. Anthropic sells inference. Its whole business is fast, clean, scalable, and made of electricity. Wet-lab biology is the opposite of all four. It is slow, it is filthy, it does not scale, and it has a long and glorious history of humbling the smartest people in any given room. Pharma pours something like 150 to 200 billion dollars into R&D every year and has, across the entire modern history of the industry, gotten fewer than a thousand drugs to market. That is a brutal denominator.

So the natural first reaction to a frontier AI lab announcing it wants to develop its own medicines is a raised eyebrow and the question of whether anyone there has met a mouse.

Anthropic's answer is the feedback loop. The pitch, stripped of polish, is that a company cannot build a genuinely good tool for a craft it does not practice. The same logic that produced Claude Code, which came from an outfit that writes an enormous amount of software, is being pointed at science. Eric Kauderer-Abrams, who runs life sciences there, framed the drug programs and the product as inseparable, and his partnerships counterpart Jonah Cool made the blunt version of the argument, which is that you cannot credibly advance a scientific workbench and think about preclinical work if you are not doing any preclinical work yourself. It is dogfooding elevated to strategy. Whether that justifies standing up actual drug programs, as opposed to, say, hiring a few scientists to file bug reports, is the open question that hangs over the entire announcement.

The bigger tell is where Anthropic slotted this. Life sciences got positioned as the single largest opportunity for the company's stated mission, right alongside its coding and knowledge-work agents. That is not a side project framing. That is a company telling the market it thinks the richest vertical it can chase, both in mission terms and in dollar terms, is the one where the customers have the deepest pockets and the most fragmented tooling.

What Claude Science actually is under the hood

First, the boring but important clarification: Claude Science is not a new model. It is an environment built on top of the existing Claude models, the same way an IDE is not a compiler. Think of it as a research workbench that happens to have an agent living inside it. The core thesis, and it is a good one, is that scientific progress is not currently bottlenecked by model intelligence. It is bottlenecked by the absurd fragmentation of the tools a working scientist has to juggle. PubMed for literature, Jupyter and R for analysis, a cluster terminal for the heavy compute, some structure viewer for proteins, a pile of half-documented lab scripts, and a lot of copy-pasting in between. Claude Science tries to pull that mess into one place.

Under the hood it wires in more than 60 databases and built-in functions spanning genomics, single-cell work, proteomics, structural biology, and cheminformatics. It renders proteins, structures, and molecules natively rather than making a scientist screenshot them from somewhere else. It plugs into Basecamp Research's EDEN dataset, which is sequencing data from millions of microbe species, the sort of thing that can turn weeks of pathogen legwork into a single working session. It writes code, and then, crucially, it runs that code on the compute clusters a lot of researchers

depend on and quietly hate managing. Babysitting an HPC job is nobody's idea of a good afternoon.

Two design choices matter more than the feature list for this audience. First, Anthropic emphasizes that Claude Science can be deployed so that large or sensitive datasets stay on the lab's own infrastructure, which is the only way to make data governance and compliance tractable. In a field where those are load-bearing, that is not a nice-to-have, it is the difference between a pilot and a hard no from legal. Second, reproducibility is treated as a first-class feature rather than an afterthought. Every figure, every result, traces back to its underlying source code, the message history that produced it, and a plain-language account of what happened. There is also multi-agent orchestration, where a lead agent spins up specialized sub-agents, and session forking, so a researcher can branch an analysis and compare two approaches without torching the original.

That reproducibility point is the whole ballgame, and it deserves emphasis because it is the thing that could actually convert a skeptical bench scientist. The perennial complaint about large models in science has never been that they are dumb. It is that you cannot trust the figure, because you cannot see how it was made. A workbench that shows its work, down to the code, is attacking the real objection instead of the imaginary one. Early testers reportedly used it for single-cell RNA sequencing, CRISPR screen design, and protein structure prediction, and in one widely repeated anecdote a UCSF researcher caught viral contamination sitting in a dataset within minutes, contamination that had gone unnoticed. The product is in beta and available to paying Claude Pro, Max, Team, and Enterprise subscribers, with Novo Nordisk and the Allen Institute named among the early adopters.

The part where the toolmaker becomes the competitor

Here is the structurally weird bit. Anthropic wants to sell Claude Science to every pharma and biotech it can reach, and it also wants to run its own drug-discovery programs. That is selling shovels while quietly opening a mine next to your customers' mines. In most industries that arrangement gets you a very awkward sales call.

The official defense is that Anthropic will only chase neglected diseases that its pharmaceutical customers have no interest in, so nobody's toes get stepped on. It is a reasonable-sounding boundary, and it is also softer than it looks. Disease interest is not static, targets get reprioritized, and "the stuff you don't want" has a way of becoming "the stuff you want" the moment somebody demonstrates a mechanism. More to the point, the mere fact of running internal programs invites the question

every paranoid pharma data officer is already asking, which is whether the workflows and data flowing through the tool end up informing the vendor's own pipeline. The private-infrastructure design is partly an engineering answer to exactly that fear, a way of saying your data stays yours and never crosses into ours.

There is a useful contrast with how the DeepMind spinout Isomorphic Labs plays it. That company keeps its drug-design engine private and largely does not sell tooling at all. It builds an internal pipeline and licenses or partners the assets. Anthropic is attempting the harder both-and: monetize the horizontal tool and run programs off to the side as credibility and as live product research. The logic is coherent, dogfood the thing and you build it better, but coherence does not dissolve the conflict of interest, it just gives everyone a story to tell about it. For an analytical reader the honest framing is that the feedback loop is real and the conflict is also real, and Anthropic is betting the loop is worth more than the friction the conflict creates.

Neglected diseases, public benefit, and the convenient economics of doing good

Anthropic is a public benefit corporation, and it leaned on that structure to explain the disease choice. Normal drug economics punish small patient populations, so those conditions get orphaned, and a public benefit company can pick programs on the basis of patient benefit rather than net present value. Cool put it plainly, that these are areas ordinary development economics do not favor. All true, and also very convenient, and both things can be the case at once.

The technical case for pointing AI at rare disease is genuinely strong and worth stating clearly, because it is the part that is not just optics. A large share of rare disorders are monogenic, meaning a single damaged gene drives the whole thing. That makes the causal biology legible in a way that Alzheimer's, type 2 diabetes, and heart disease simply are not, since those are polygenic messes tangled up with multiple tissues and decades of environment. If the current generation of models has a fighting chance anywhere, it is in the corner of biology where the story runs closer to one gene, one mechanism, one target. Cleaner ground truth is exactly what these systems are hungry for.

Now the slightly cynical footnote, offered in the spirit of an audience that gets paid to be cynical. Neglected diseases are also precisely the terrain where Anthropic will not collide with paying customers, where a failed program is far less embarrassing, and where the mission halo shines brightest for the least commercial risk. Doing good and de-risking the optics happen to point in the same direction, which is a nice place to be standing. Reinforcing the point, the company has disclosed essentially nothing

concrete: no named targets, no budget, no timelines, no team size, just a lot of we-are-at-the-start-of-this and we-will-share-more. That is either appropriate humility or a marketing posture, and there is currently no way to tell which from the outside.

The competitive board, from DeepMind to OpenAI to the techbio pile

None of this is happening in a vacuum, and the board is crowded. For a decade the vanguard here was DeepMind, whose AlphaFold work earned Demis Hassabis and John Jumper a Nobel.

Its drug-focused spinout, Isomorphic Labs, raised a multibillion-dollar round in May 2026, among the largest rounds the AI drug space has seen, on top of earlier funding, and has papered deals with Lilly, Novartis, and Johnson & Johnson that could reach into the billions in milestones. It is also the cautionary tale sitting right there in the data. Isomorphic has disclosed almost nothing about its pipeline while raising billions of dollars, which is roughly the biotech equivalent of a magician charging billions for a ticket and refusing to show the trick, and its first-in-human timeline has already slipped toward the end of 2026 with no patients dosed at last public word.

Money and secrecy are not the same as a drug.

OpenAI is coming the other way, through partnerships, having lined up Novo Nordisk, Lilly, Moderna, Sanofi, and Amgen, and it too talks openly about autonomous researchers aimed at discovery. Then there is the pure-play pile, which is deep. Insilico Medicine signed a Lilly collaboration worth up to 2.75 billion, the largest AI-era pharma deal on record. Chai Discovery, valued around 1.3 billion, has its own Lilly pact. Recursion, Schrodinger, Eikon, Exscientia, BenevolentAI, AbCellera, Xaira, which launched with a billion dollars in 2024, and Manifold Bio, which signed with Roche, all populate the field. The macro backdrop is a genuine boom: AI-related biopharma dealmaking approached 10 billion dollars in 2024 alone, and the total value of AI partnerships in the space rose about 120 percent year over year into 2025.

Where Anthropic actually differs is instructive. It is not raising a dedicated pharma megaround. It is not licensing an asset pipeline. It is folding Claude Science into existing Claude subscriptions, which is an accessibility play rather than a bespoke enterprise-only motion, and it is treating its drug programs as a credibility signal attached to a horizontal software business. The enterprise engine is already visible

elsewhere in its book, including a May 2026 deal to roll Claude out to more than 30,000 Bristol Myers Squibb employees. In other words, the drugs are the garnish. The entree is selling the workbench to everyone who does science.

Where the skeptics are right, which is the target problem

The most useful criticism of this whole category is not the reflexive it-will-never-work take, which is lazy, but a sharper one about where AI has and has not moved the needle. The Boston Consulting Group analysis of roughly two dozen AI-discovered molecules is the number to internalize. Phase 1 safety success climbed to somewhere between 80 and 90 percent, well above the historical figure near 50, which sounds like a triumph until the next line: Phase 2 efficacy dropped right back to the industry's ordinary 40 or so. Read together, those two numbers tell a specific story. AI got very good at the tractable, data-rich, chemistry-and-physics problem of designing a molecule that behaves. It has barely touched the expensive, biology-limited problem of choosing a target that actually drives the disease.

That is the crux, and it is worth being blunt about it because the marketing tends to blur it. Molecule design was always the cheaper early stage. Target selection depends on understanding human biology that remains sparsely mapped, and no workbench, however slick, fixes a map that has not been drawn. Frank von Delft at Oxford made the plain version of the point, that these models have not come close to making experiments unnecessary. Jared Auclair at Northeastern added the operational worry, that general-purpose systems hallucinate and miss the fine print in regulatory guidance and assay design, and that in drug development those small errors compound into large, expensive ones. The wet lab remains the final arbiter, and it does not care how confident the model sounds.

There is a reason 2026 keeps getting described as the year AI drug discovery meets clinical reality. A meaningful cohort of AI-discovered candidates is finally reaching Phase 3, and the first real efficacy readouts are expected by year end. Those results, not another molecule-generation platform, are what will re-rate or deflate the entire thesis. Set against that backdrop, the honest read on Claude Science is that it is a serious productivity tool aimed at the parts of the workflow that were already the least broken. Compressing data wrangling, contamination checks, figure generation, and cluster babysitting is real value, and it frees scientists to spend more of their scarce judgment on the hard calls. That is a good thing. It is also not the same thing as curing disease, and anyone selling it as the latter should be met with a polite but firm show me.

The money math and the IPO clock ticking in the background

Strip away the mission language and there is a very legible financial story underneath. Anthropic is reportedly set to post its first profitable quarter and is pointed at an IPO later in 2026. Its compute bill is enormous, with reporting putting one arrangement at well over a billion dollars a month, and the broader token-consumption gold rush that has floated the whole sector is not guaranteed to stay hot forever. A company in that position needs durable, high-margin, defensible revenue, and it needs a story an underwriter can put in a deck. Pharma, sitting on far deeper pockets than the academic labs also represented at the launch, is about the best-looking vertical on the menu.

The strategic shape is clean once you see it. Build horizontal agent products, point them at the richest verticals, and let the vertical with the strongest mission narrative double as the marketing centerpiece. Running a handful of internal drug programs is, in that light, remarkably cheap advertising relative to the size of the enterprise contracts it might unlock, and it doubles as live R&D that makes the product better. The move that looks eccentric from a biology desk looks almost obvious from a corporate-strategy desk.

The risk cuts the same way it always does with this kind of hybrid. Taken seriously, in-house drug programs are a money pit that can swallow attention and capital and produce nothing for years, which is the base rate for the industry. Taken un-seriously, they are a stunt, and sophisticated pharma buyers, who have watched a parade of tech companies wander into healthcare and wander back out, will smell it immediately. Anthropic is threading between those two failure modes, and it has to keep threading them straight through an IPO roadshow, which is not a forgiving audience for ambiguity.

What is actually worth watching

A few things separate signal from theater from here. The first is disclosure. If Anthropic keeps the drug programs vague indefinitely, no named targets, no timelines, that is a strong hint the programs are primarily positioning. Specifics, on the other hand, would signal actual commitment and actual expense. The second is adoption friction. Whether pharma will run Claude Science on private infrastructure at real scale, or whether data-governance instincts cap it at cautious pilots, decides whether this is a business or a demo.

The third, and the one that matters for the whole category rather than just this company, is the Phase 3 wave landing by the end of 2026. If efficacy shows up across the AI-discovered cohort, from Isomorphic to Insilico and the rest, the entire space re-rates and Anthropic's timing looks prescient. If it does not, the cure-all narrative takes a real bruise and this launch reads as either brave or a couple of years early. The fourth is quieter but telling: whether reproducibility and auditability actually win over bench scientists, because that transparency is the wedge, and wedges either split the wood or they do not.

And then there is the IPO itself, which is the honest bottom line. If Claude Science matures into a line item that investors can underwrite, the story stops being another chatbot company and becomes something closer to scientific infrastructure, the plumbing under a large slice of research. That is the actual prize, and here is the punchline that makes the drug programs make sense: Anthropic can win that prize whether or not it ever ships a single pill. The medicines are the mission. The workbench is the business. Everyone at the launch event knew which one pays the compute bill

Thoughts on Healthcare Markets &
Technology is a reader-supported
publication. To receive new posts and
support my work, consider becoming a
free or paid subscriber.

[← Previous](#)

Discussion about this post

Comments

Restacks



Write a comment...