The Molecule Economy

When biology sets the price.

Jay Srinivasan
CEO, infinitusbio
November 2026

One question, priced everywhere

Every major capital decision in biopharma now turns on a single question: *will this molecule work, and in whom?*

A pharma company buying a biotech asks it. So does a royalty fund wiring a billion dollars before Phase 3, and a hedge fund sizing into a readout. A biotech in China negotiating an upfront payment asks it, and so does a trader quoting a contract on a trial outcome.

They are different buyers using different instruments. But they are all pricing the same thing: the future behavior of a molecule inside human biology.

We call the market that has formed around that question the Molecule Economy.

Value now moves through molecules, not companies

For most of the industry's history, capital bought companies and hoped the science inside them would pay off. In 2025 and 2026, capital started buying, licensing, financing and trading individual molecules directly.

- ***Acquisitions.*** *Biopharma deal value reached a record $170 billion in the first half of 2026. Most large deals were bets on one or two assets: an IL-13 antibody, a lung-cancer kinase inhibitor, a myeloma cell therapy.*
- ***Licensing from China.*** *Companies in China took eight of the ten largest global licensing deals in the first half of 2026. Western buyers are now shopping for molecules, not partners.*
- ***Royalties.*** *The royalty market reached a record $10 billion in 2025. In May 2026 Blackstone Life Sciences committed up to $1.3 billion against a single antibody that had not started Phase 3.*
- ***Recycling.*** *Specialist buyers now acquire biotechs within weeks of a failed trial, keep the cash and sell off the remaining molecules. One buyer had made seven such offers by August 2025.*
- ***Trading the outcome itself.*** *In July 2026 Kalshi, a federally regulated exchange, announced contracts on clinical-trial results and FDA decisions.*

The unit of value has shrunk from the company to the molecule. Whether that molecule works has become the most expensive question in the industry.

Everyone is pricing the same blind spot

About nine in ten drugs that enter clinical trials never reach approval. Roughly half of those failures come down to one cause: the drug did not work well enough in the patients who were tested (Sun et al., *Acta Pharmaceutica Sinica B*, 2022). Counting the failures, each approved drug now costs an estimated $2.67 billion (Deloitte, 2026).

The cost of being wrong lands on whoever holds the molecule on the day it reads out. In May 2025, a closely watched lung-cancer combination improved progression-free survival but missed on overall survival, and its developer's shares fell as much as 31% that day. In December 2025, a rare-disease company's lead drug failed its Phase 3 trial and its shares fell as much as 90% in a day.

The reverse is just as costly. A molecule that failed in the wrong indication, or in unselected patients, is often sold for little more than its cash, even when it might work somewhere else. Most failed drugs are not bad molecules. They failed in the wrong indication, in the wrong patients or, in some cases, with the wrong partner.

More data has not changed the odds

The industry's answer to uncertainty has been more of the same: more data, more expert calls, more models trained on past trials. Buy-side firms keep raising these budgets, and 94% expected to spend more on alternative data in 2026. The approval odds have barely moved, for two reasons.

1. ***Opinion converges.*** *Expert networks and key-opinion-leader calls are useful, but every serious fund calls the same few clinicians. The result is a well-informed consensus, which is exactly what the price already reflects.*
2. ***History repeats its biases.*** *Models trained on past trials learn which kinds of drugs used to succeed. They inherit survivorship, sponsor and publication bias, and they are weakest where the money is: new mechanisms, new combinations, rare diseases and molecules nobody has tested in a given indication.*

Meanwhile, most AI investment in TechBio still goes to drug discovery: finding and designing more molecules. Better-designed molecules have not made better drugs. The bottleneck is not finding molecules. It is knowing, before capital is committed, what a molecule will do in real patients.

Read the biology, not the history

There is another way to ask the question: start from the biology instead of the track record. Take mRNA data from a patient's tumor or tissue biopsy. Build a working model of that patient's cells. Add the drug's known targets. Then watch, mechanistically, what happens inside the cell.

That is what we built at infinitusbio, the intelligence company for the Molecule Economy.

Our macro-connectionist multi-agent AI simulates a whole human cell with more than 200,000 interacting agents. Limited inputs from public or proprietary data are enough to build digital models of patient cells and biopsies. It learns nothing from past trials, so every call is out of sample by construction. It is deterministic, so the same input gives the same answer. And every prediction traces back to the molecules and pathways behind it, so it can be audited.

The early evidence is encouraging, and we state it plainly:

- *Ranking 17 small-cell lung cancer cell lines by drug response with no training data, the engine matched wet-lab results with a Spearman correlation of 0.82.*
- *In all seven retrospective tests in cancer and inflammatory disease, it agreed with wet-lab outcomes.*
- *Prospectively, it predicted that a major oncology drug would fail; the sponsor later ended the trial after Phase 2. It also identified a new inflammatory bowel disease target that a client confirmed in the lab.*

These are predictions, not clinical proof. But they answer the questions every buyer in the Molecule Economy is pricing: which indication, which patients and which combination.

- *An acquirer can test a target's next indication.*
- *A royalty fund can check how wide a label might grow.*
- *A licensor in China can give buyers in the US and Europe an independent read to audit.*
- *A buyer of a failed asset can look for the subgroup where it works.*

Information on equal terms

A read on a molecule is only worth something to a market if everyone can get it on the same terms. The moment one party sees it first, it stops being market information and becomes leverage: a head start to trade or reposition, or a window to keep it from being published at all. Markets already expect companies to disclose to everyone at once. Research about their molecules should meet the same standard.

Our public reads follow three commitments.

1. **One clock.** Every public read is released to everyone who orders it at the same moment. No one gets an early look: not a fund, not a partner and not the drug's owner. The owner receives a copy at that same moment, and may publish a response, which we may post alongside the read.
2. **One price.** Each read carries a published price estimate before orders open. Every buyer pays the same final price, set within that estimate when orders close.
3. **One record.** Each read is sealed with a public, timestamped fingerprint before anyone sees it, so it cannot be edited or withdrawn. After the readout it is released free to everyone and scored, right or wrong. We score every call in public, matched or missed.

Two lines we will not cross.

- Companies can commission private reads using their own data; that work stays confidential, never enters a public read, and cannot change, delay or remove one. And no one, including a drug's owner, can pay us to keep a read off the record.
- Public reads use public data only. We take no positions and give no investment advice.

That record starts now. We are sealing calls on a first wave of trials in oncology and autoimmune disease that read out in 2027. Each will be released on the same clock and scored in public when its result arrives.

Discovery AI fills the network

Most of the money flowing into AI in TechBio is spent on finding new molecules. In AI drug R&D, discovery engines captured the majority of the $3.8 billion in equity funding raised in 2024. That spending is not our competition. It is our supply.

Every molecule a discovery platform designs enters the Molecule Economy as an asset that someone must decide to build, buy, sell or fund. Each of those decisions turns on the same question: will it work, and in whom? Discovery AI produces more candidates. It does not answer that question.

As we network orchestrate the Molecule Economy, connecting the owners, acquirers, licensors, funds and investment banks who price molecules through one shared read, every new molecule adds demand for that read and another node to the network. The more capital discovery AI attracts, the more assets need an independent read, and the more valuable the network becomes. We benefit from that investment without having to make it.

Read the molecule before you price it

Keytruda, the first PD-1 drug, was approved in 2014 for melanoma. By 2017, about 70% of its sales came from cancers added later, led by lung cancer. The best molecules often sit in plain sight, underpriced because nobody has read them against every disease where they might work.

The Molecule Economy will keep growing. As it does, infinitusbio will read each molecule against every indication where it might work, so more molecules reach their full therapeutic potential. More assets will cross borders, more capital will fund trials directly, and more of the outcome will be traded in the open. A market like that runs on information that is independent, sealed in advance and available to everyone at once. In it, the edge goes to those who understand what a molecule will do before the price reflects it, and who can show their reasoning.

We will be in San Francisco for JPM week in January 2027.

If your work depends on what a molecule can do, we would like to compare notes.

infinitusbio provides scientific research on drug efficacy. It is not investment advice. Predictions are not clinical evidence and do not replace clinical trials or regulatory review.

Sources

- BioWorld, record $170.11 billion in biopharma deal value in H1 2026 (July 17, 2026): bioworld.com
- IQVIA, largest biopharma acquisitions of H1 2026 (July 8, 2026): iqvia.com
- VCBeat, China's share of the top 10 licensing deals in H1 2026 (July 15, 2026): vcbeathealth.com
- Royalty Pharma, record $10 billion royalty funding market in 2025 (J.P. Morgan Healthcare Conference release, January 2026): barchart.com
- MedCity News, Blackstone Life Sciences' up to $1.3 billion financing of Apogee's zumilokibart (May 27, 2026): medcitynews.com
- BioSpace, Concentra's buyout offers so far in 2025 (August 13, 2025): biospace.com
- GV Wire, Kalshi and AppliedXL contracts on clinical-trial results and FDA decisions (July 16, 2026): gvwire.com
- The Motley Fool via Nasdaq, Summit Therapeutics' share move on the HARMONi readout (May 30, 2025): nasdaq.com
- Business Wire, Rezolute's sunRIZE Phase 3 failure and share move (December 11, 2025): businesswire.com
- Exabel buy-side alternative data survey, 2026 spending intentions: hedgeweek.com
- Sun D. et al., Why 90% of clinical drug development fails and how to improve it?, Acta Pharmaceutica Sinica B 12(7), 2022: doi.org/10.1016/j.apsb.2022.02.002
- Deloitte, Measuring the Return from Pharmaceutical Innovation, average cost per approved drug in 2025 (May 4, 2026): deloitte.com
- CB Insights, AI in drug R&D market map, 2024 funding and the share captured by discovery engines (May 23, 2025): cbinsights.com
- Merck, Keytruda's accelerated approval as the first FDA-approved anti-PD-1 therapy, in melanoma (September 4, 2014): merck.com
- BioPharma Dive, Keytruda sales by indication: about 40% lung cancer, 30% melanoma (May 2, 2017): biopharmadive.com