SEP 2, 2026

We Finally Got Our Flying Cars, They're Called Neolabs

Neolabs brought ambitious research back, now investors have to learn how to price it.



By Liang Wu, Jenny Xiao, and Jay Zhao

Leonis Capital · leoniscap.com/research/we-finally-got-our-flying-cars-they%E2%80%99re-called-neolabs/



For years, Silicon Valley's favorite self-critique was that "we wanted flying cars and got 140 characters instead." The complaint became a genre. Venture had abandoned ambition, the argument went, and settled into financing photo apps and dashboards while the hard problems went unfunded.

Then recently investors started doing the opposite. Thinking Machines raised $2 billion at a $12 billion valuation before launching a product. Project Prometheus raised $12 billion at a $41 billion valuation to pursue an "artificial general engineer" for the physical world. Jeff Dean's Discovery Loop has reportedly been in talks to raise $1 billion at roughly a $10 billion valuation, shortly after Dean and several longtime Google researchers left to start the company. Capital was once again

following researchers into problems whose technical outcome, product, and eventual market were still unresolved.

Since 2024, investors have put $94.3 billion into 119 AI research labs. Most are pre-revenue and some are pre-product, by design. Only 16 disclose any revenue at all. The market has started calling them “neolabs,” and together they represent one of the largest bets on research-before-revenue in venture capital history.



A neolab is a company organized around a research bet. It trains its own models in pursuit of a technical breakthrough, with any product downstream of the research rather than the reason for it. Early in the company’s life, the research itself is the product. Training a model is necessary but not sufficient. [1] The question is whether the research serves the product or the product, if one arrives, serves the research.

Our dataset applies that test to 119 labs spanning frontier models, physical AI and robotics, world models, AI for science, new architectures, formal reasoning, and recursive self-improvement.

The capital came back faster than the underwriting did. Venture finally got its flying cars, now it has to learn how to price them.

From Flying Cars to Neolabs

Venture's retreat from hard technology in the 2010s was less about ambition than incentives. Software had become an unusually good asset class for venture capital. Companies could reach product-market fit on a few million dollars, feedback arrived within months through usage and revenue, and each incremental dollar bought observable progress against observable demand. Marginal distribution costs were near zero, and successful companies could compound quickly. Deep tech offered the opposite profile: hundreds of millions of dollars (or more) was needed before commercialization, years before the central thesis could be tested, and technical risk stacked on top of manufacturing and regulatory risk, with repeated dilution along the way. **Faced with those two sets of economics, venture simply made the rational choice.** It funded the asset class where experiments were cheaper, feedback was faster, and progress was easier to measure.

Foundation models changed that tradeoff. Foundation model labs combine research risk with something closer to software-speed commercialization. When a lab produces a meaningful technical breakthrough, the path from research breakthrough to paying customers can be measured in months rather than decades. While the margin structure differs from software, commercialization still more closely resembles software than traditional deep tech. OpenAI and Anthropic showed that a research organization could make this transition at enormous scale. Foundation models are also absorbing capabilities once expected to sit in the application layer, strengthening the case that durable value may accrue closer to the model layer.

Most of the 119 labs in our dataset are still in the phase before commercialization. A neolab can raise hundreds of millions or billions on a thesis and founding team, then spend that capital on compute and researcher salaries while the commercial value of the work remains unclear. It may raise again at a higher valuation before shipping much. Early commercialization was never the point. Revenue, if it arrives, follows the research. The result is a long research cycle with sparse external readouts financed through staged private rounds.

The Neolab Underwriting Vacuum

Venture has funded a close analogue before in biotech. Neolab financing resembles biotech because investors fund years of research before the central technical thesis is fully resolved. Biotech reprices across rounds through a milestone-based system built around clinical evidence: preclinical results, Phase I, II, III, and approval, with endpoints and regulators enforcing the gates. SaaS reprices as product adoption turns into revenue, growth, and retention. Neolabs have no equivalent external sequence anchoring valuation to evidence.

A neolab reprices on a different set of signals. Founder pedigree can influence the formation price, a leading researcher joining or a major compute commitment can move it again. Private

evals are shared selectively with investors, and strategic partnerships are announced. Each signal can improve the probability of success, but none provides the comparability of retention or a clinical endpoint.

Technical milestones exist, but they are heterogeneous, mostly private, and rarely comparable across labs. Without a common external scoreboard for how much research risk has been removed, valuations can move substantially before an externally legible technical result arrives.

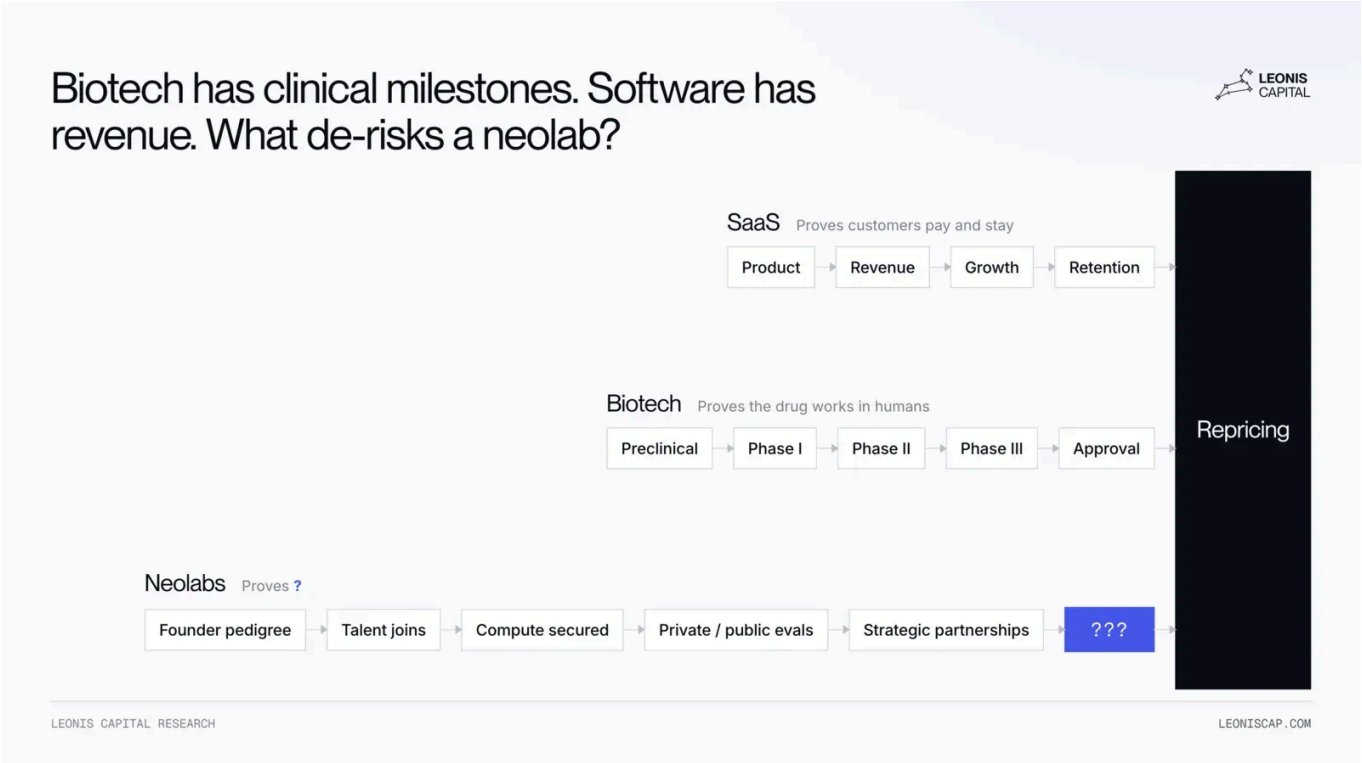

The consequence is unusually wide repricing even when public evidence changes little. Thinking Machines reportedly sought a $50 billion valuation in late 2025, roughly four times its previous valuation. The new round never closed, leaving $12 billion as the last completed valuation. A $38 billion range would normally imply dramatically different assumptions about the underlying asset. Here, the gap appears to reflect investor sentiment far more than any comparable public change in the research.

Revenue makes the underwriting vacuum especially visible. Among the 119 neolabs in our dataset, only 16 disclose revenue, while the vast majority do not. Revenue provides little common ground for comparing most of the category, and even where it exists, valuation does not map cleanly to commercial output. For example, Mistral reports roughly $400 million in ARR against a $14 billion valuation, while Skild discloses $30 million in revenue at the same $14 billion valuation. The same price is being applied to two companies whose valuations are anchored to very different evidence. Mistral is increasingly valued against commercial output, while Skild is still

valued primarily on technical potential. The market currently uses the same valuation vocabulary for both.

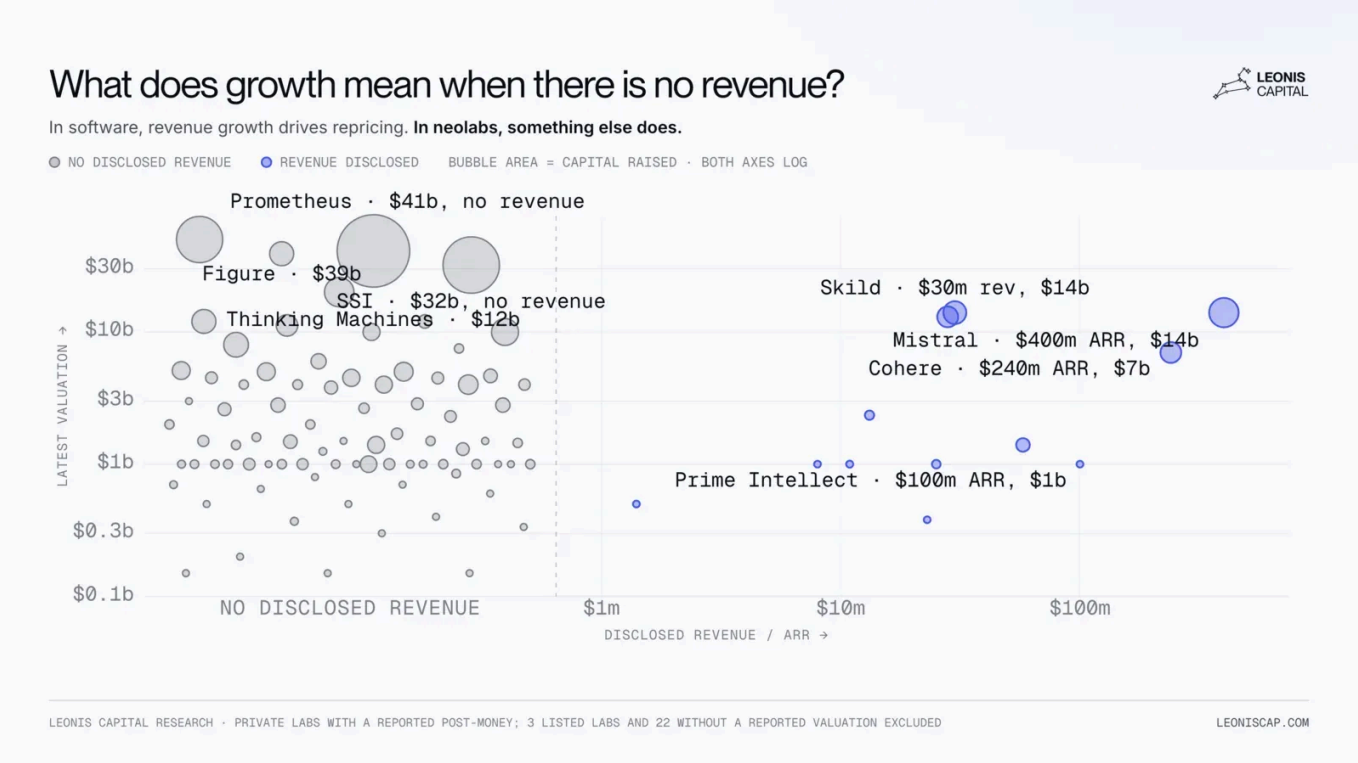

For the 103 labs without disclosed revenue, the market relies on proxy signals such as markup velocity, researcher headcount and pedigree, compute under contract, and usually private evals. These measures mostly describe inputs into the research program. Markup velocity is even more circular because investor enthusiasm becomes both the thing being measured and part of the evidence used to justify the next price.

The Global AI Talent Auction

Part of the pricing problem is a labor market problem. A neolab lets researchers turn conviction into an independent research budget. Capital can buy things like compute, but it cannot easily manufacture another researcher capable of leading a frontier training program. The scarce input is the small global population who can credibly lead frontier-adjacent training runs, plausibly numbering in the low thousands. **Neolabs are the price discovery mechanism for that labor.**

When a researcher's conviction exceeds the compute allocation inside an incumbent lab, outside capital can fund the bet directly. The supply of new labs is therefore partly downstream of frontier-lab org charts. Mira Murati leaves OpenAI and a $12 billion seed-stage company exists. Yann LeCun leaves Meta and Europe gains a billion-dollar lab. Researcher moves are the auction running in public.

The auction also highlights that the **valuation is an input to the business, not an output of it.** Pulling a senior researcher from OpenAI or DeepMind may require an equity package that competes with eight- and nine-figure retention offers. A high headline valuation and a large round help finance that compensation. In that sense, valuation itself becomes recruiting infrastructure.

Talent also creates a downside floor in the neolabs category. Inflection, Adept, and Character showed that a lab with a great team can retain strategic value even when the original company does not continue independently, while SAP's acquisition of Prior Labs shows that successful research can attract a strategic buyer. This floor compresses the loss distribution and can partially rationalize entry prices at the low end. If talent is both the primary input to a neolab and part of its residual value, then understanding where that talent comes from becomes central to understanding the category itself.

Our dataset covers 284 founders across three distinct talent ecosystems, the United States, China, and Europe. The U.S. and China account for 83% of disclosed neolab capital in our dataset.

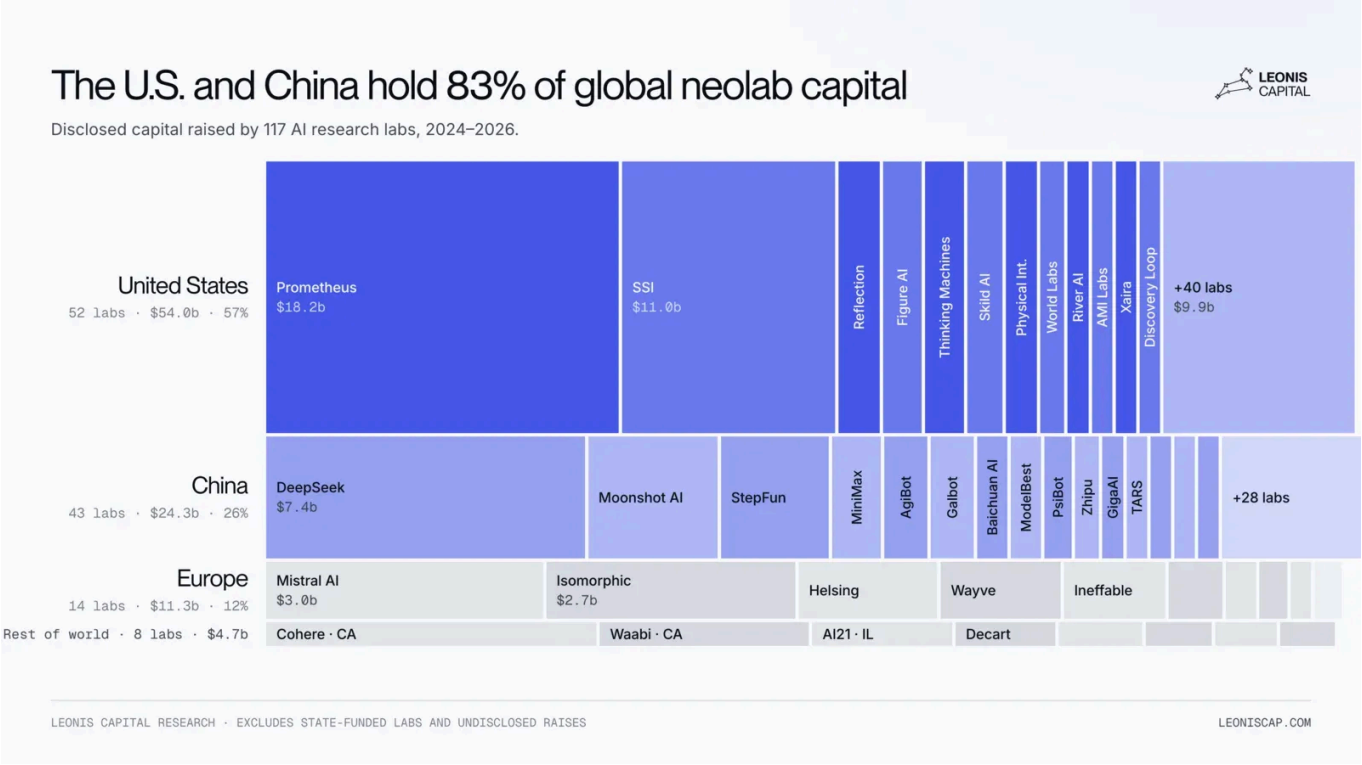

The supply of neolab founders reflects the institutional structure of research talent in each geography. In the U.S., frontier industry labs function as "finishing schools" for founders, with many coming through Google, OpenAI, and Meta. China's pipeline is more academic. Europe's ecosystem is disproportionately a DeepMind diaspora. The talent auction is global, but the institutions producing neolab founders differ by region.

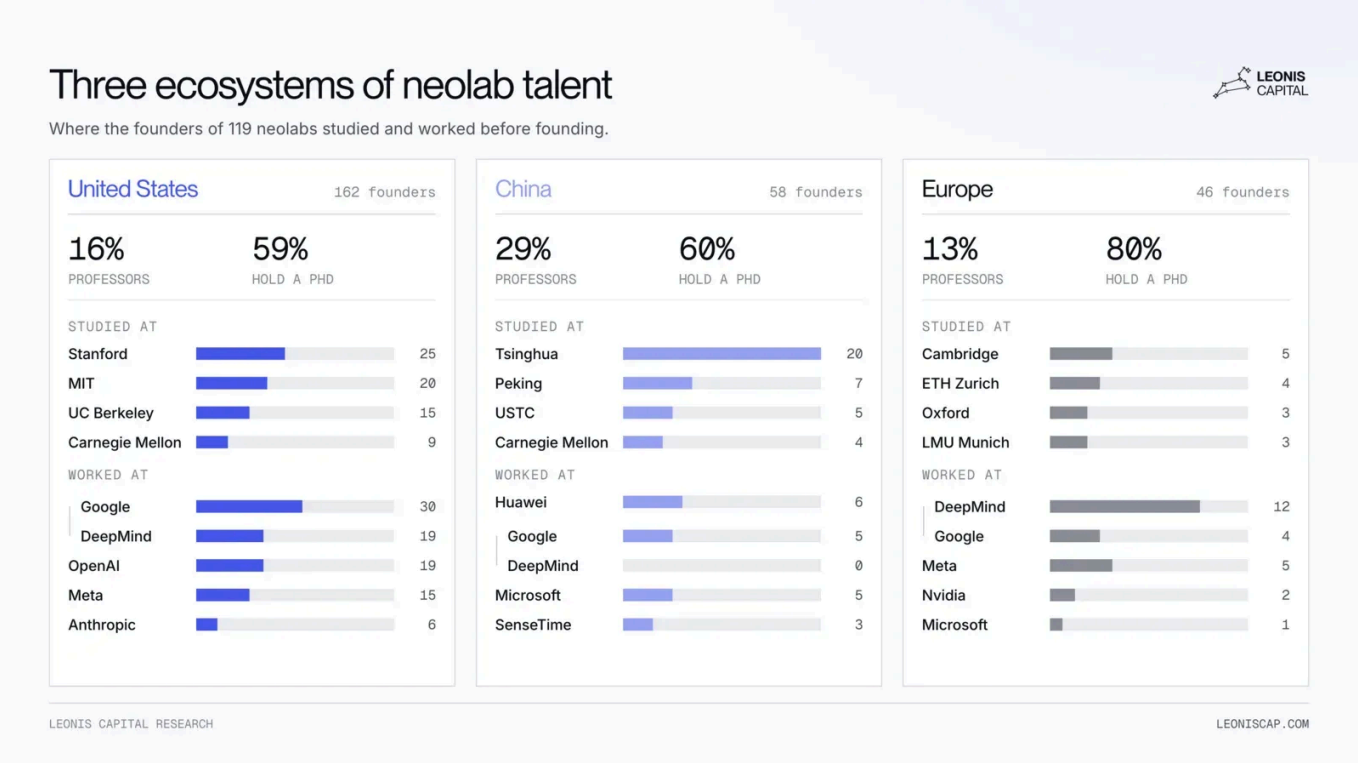

The common pattern is that lab formation follows researcher career structures at least as clearly as technology cycles. If you want to predict where the next wave of neolabs comes from, study who is ready to leave which frontier lab and what their equity is worth, not only the research frontier.

Who Should Fund Neolabs?

Technological importance and investor return are different variables. The neolab era may force the venture industry to relearn that distinction. Railroads transformed economies while repeatedly destroying investor capital. Telecom infrastructure carried the internet into existence even as many of the companies that built it went bankrupt. **Technological importance does not automatically accrue to the equity that financed it when commercialization requires repeated rounds of capital.**

Consider what the same $100 billion outcome means for two very different research companies. A superintelligence lab valued at $32 billion today and requiring another $15 billion before commercialization would return only about 2.5x to today's investor at a $100 billion exit. [2] A 10x return would require an outcome approaching $400 billion. A venture-shaped lab valued at $1 billion and needing only $300 million more to scale would return about 87x at the same $100 billion exit. **At the same $100 billion exit, entry price and the capital still required before commercialization produce radically different returns for today's investor.**

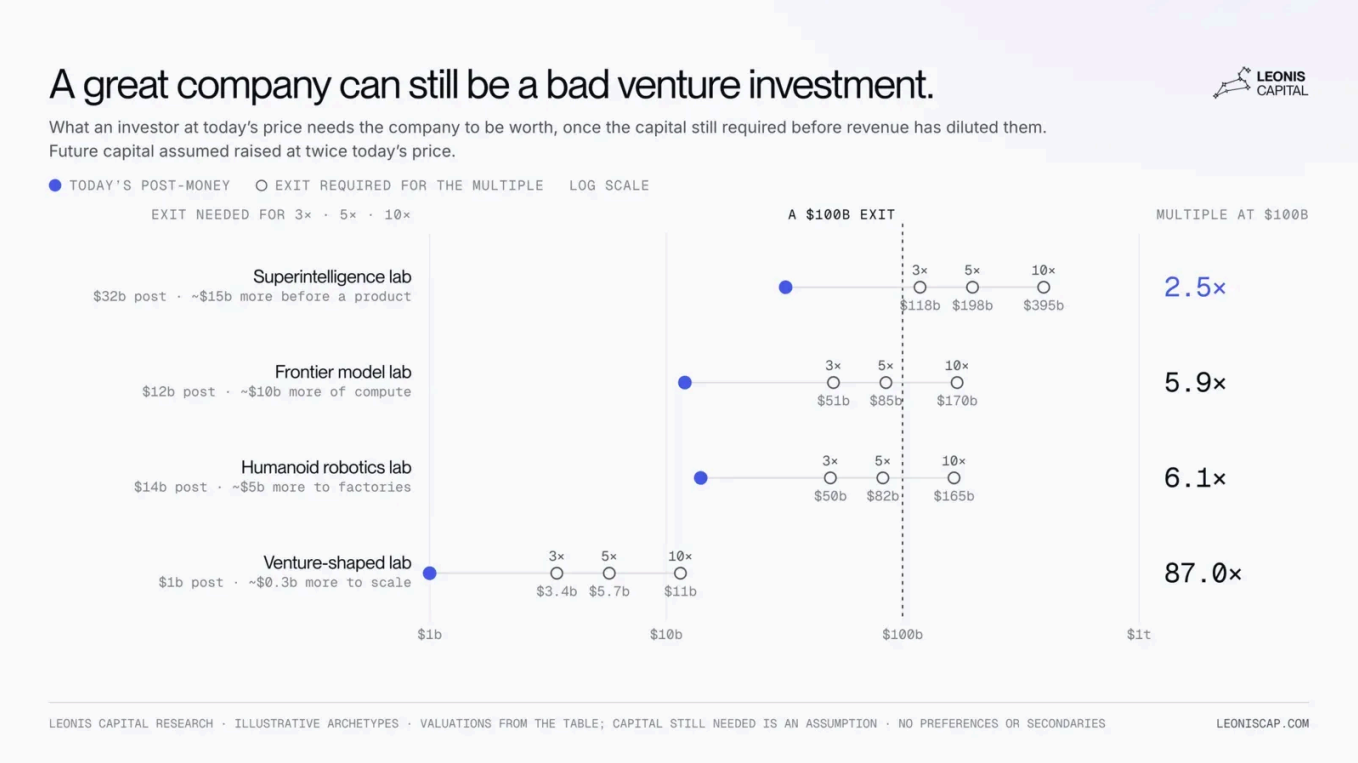

As a research program consumes more capital before commercialization, fewer outcomes can produce venture returns. Capital-intensive labs can still create extraordinary technological value while producing very different returns for today's equity holders. Pricing a neolab starts with understanding how much capital the research will consume before commercialization and how quickly a technical breakthrough can become a sustainable business.

Two dimensions shape the economics of a neolab: how much capital it takes to learn whether the research works, and how quickly a successful result can turn into revenue. Plotting those dimensions produces four economic shapes. The 2×2 below distinguishes research assets with very different financing burdens and commercialization paths, giving investors a starting point for pricing them.

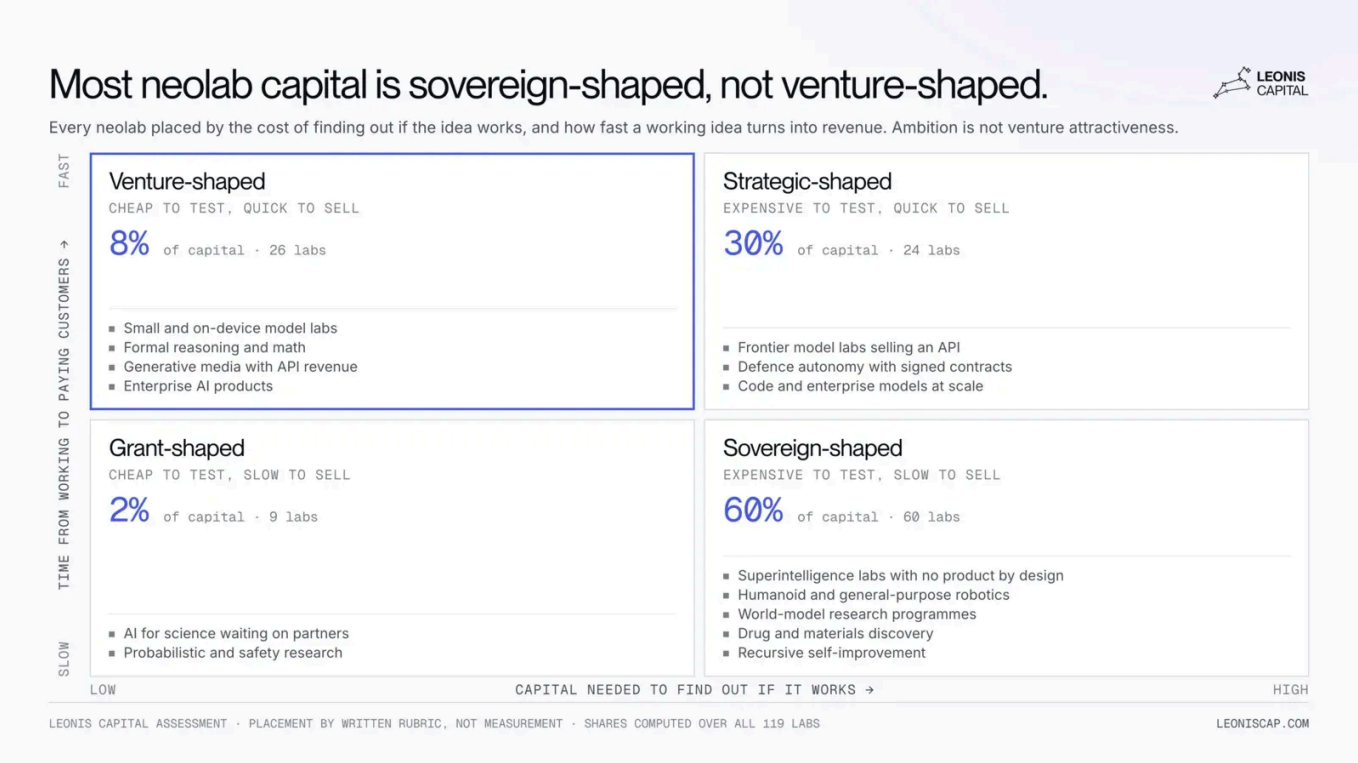

The distribution across the four shapes is highly uneven. 90% of disclosed neolab capital sits in research programs that are expensive to prove (right side of the 2×2). 60% sits in the sovereign-shaped quadrant alone. Only 8% of capital sits in the quadrant whose economics most naturally resemble traditional venture capital. **Sovereign-shaped labs are the true "flying cars," accounting for most of the capital in the category.**

Sovereign-shaped labs are expensive to test and slow to sell. Sixty of 119 labs fall into this quadrant, spanning superintelligence labs with no product, humanoid and general-purpose robotics, world models, drug and materials discovery, and recursive self-improvement. Their relevant milestone is a demonstrated capability that a generalist lab cannot replicate within a funding cycle. The historical analogue is the state. The Manhattan Project, Apollo, and DARPA-funded programs that were expensive, slow, and justified by capability rather than cash flow. The state could hold them because strategic capability itself was part of the return. That same logic helps explain why patrons and sovereigns already appear in the largest rounds, from Bezos anchoring Prometheus to state and industrial funds across Chinese humanoid labs.

Strategic-shaped labs are expensive to prove but quick to commercialize once the research works. They represent 24 labs and 30% of capital, including frontier model labs selling APIs, defense autonomy with signed contracts, and code models at scale. The relevant milestone is evidence that the technical capability can translate into durable unit economics or contracted demand. These companies are natural fits for strategic investors that capture value outside the equity return: Nvidia through compute demand, defense primes through capability, hyperscalers through platform pull-through, and corporate parents through proprietary IP. That gives them a

different payoff function from a venture fund. Nvidia has backed more neolabs than any venture firm in our dataset, appearing on 34 of 119 cap tables.



Venture-shaped labs are cheap to test and quick to sell. They account for 26 labs and 8% of capital in our dataset, spanning small and on-device models, formal reasoning, generative media with API revenue, and enterprise research products. Their economics preserve what made software attractive to venture: cheap experimentation, fast feedback, and fast monetization. The relevant milestone is revenue within roughly 24 months. Frontier research here can reach commercial feedback quickly enough to behave like a venture asset.

Grant-shaped labs are cheap to test and slow to sell. Nine labs in our dataset represent 2% of capital, mostly in AI for science, probabilistic methods, and safety research. The relevant milestone is a validated result that an institutional partner will pay to deploy. Historically, the owners in this quadrant would be universities and philanthropy, institutions built to fund useful research when the path from discovery to financial return is too indirect or too slow for venture investors.

These categories are not permanent and can change as the research de-risks. AI can compress research cycles, simulation can replace some physical testing, and a technical breakthrough can move a lab from sovereign-shaped to strategic-shaped faster than historical analogues suggest. The current distribution helps explain why the category has attracted venture funds, strategics, patrons, and sovereign investors. They can all rationally pay different prices for the same research asset because **they all have different payoff functions and incentives.**

How This All Ends

We do not think the Neolab boom necessarily ends in one broad collapse. More likely, it ends in a **re-sorting of capital**.

Today, venture funds, strategics, sovereigns, and patrons are often financing the same research programs despite having very different payoff functions. That can persist while capital is abundant and expectations for technical progress remain high. That becomes harder when the next round depends on proving that the previous billions meaningfully reduced technical or commercial risk.

Our expectation is that the four quadrants begin to separate. Venture fits best where technical risk resolves cheaply and commercial feedback comes quickly. Strategics are better suited to fund expensive research that strengthens an existing business, while sovereigns and patrons can fund programs whose value lies in capability rather than just financial return. This re-sorting process will lead to corrections in the neolabs category. The most vulnerable neolabs are those whose capital structure assumes venture-like returns while their research behaves like a decade-scale scientific program.

But the consequences of getting the capital wrong extend beyond investment returns. **Capital is not neutral to the research it finances.** A five-year research problem financed by investors who need visible step-ups every 12 to 18 months will eventually be asked to produce milestones on that cadence. A lab can start packaging immature research into products, chasing benchmarks that are legible to investors, or redirecting researchers toward commercial milestones before the underlying technical thesis has matured.

Capital that misunderstands what it owns does not just risk losing money. It can change what gets researched, on what timeline, and toward whose payoff function.

The twentieth century took decades to work out which kinds of science belonged inside universities, corporate labs, government programs, philanthropy, and venture-backed companies. Neolabs are forcing the market to relearn that institutional lesson in a fraction of the time.

Investors finally got their flying cars. Now they have to learn how to do price discovery without forcing every form of it into the same financial model.

FOOTNOTES

1. The definition turns on the direction of dependence between research and product. We consider Project Prometheus a neolab because it is a research program first, with applications to be found if the science works. Harvey is not because it trains models to serve a legal software business.
2. Illustrative return calculations assume an investor enters at the company's current post-money valuation and that all additional capital required before commercialization is raised in a single future financing at a pre-money valuation equal to 2× the current post-money valuation. Returns are calculated on a fully diluted basis after that financing and assume no further dilution. The analysis ignores liquidation preferences, option-pool expansion, secondaries, fees, taxes, and subsequent financings. These scenarios are intended to illustrate how entry valuation and the capital required before commercialization interact to shape investor returns, not to forecast specific company outcomes.