Fall 2025

Return On Investment for AI-Enabled R&D

A look at how lessons from HTS may guide our understanding of future value

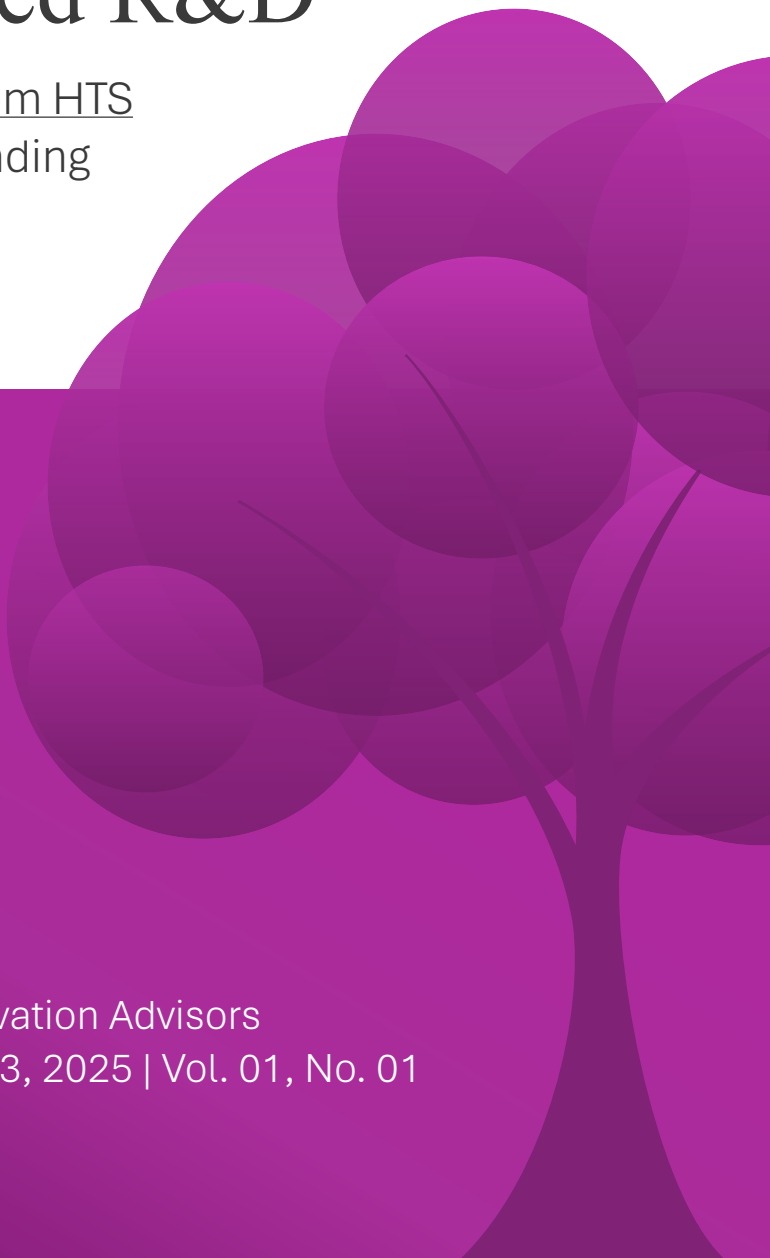

elm
Innovation Advisors

A white paper by Elm Innovation Advisors
Elm Idea Workshop | Dec. 3, 2025 | Vol. 01, No. 01

Executive Summary

Two decades ago, high-throughput screening (HTS) was hailed as the future of drug discovery. It promised to industrialize the identification of lead compounds through robotic automation and vast combinatorial libraries. Billions were invested; tens of thousands of assays were run; few marketed drugs emerged.

Today, artificial intelligence (AI) occupies a similar cultural and financial moment. Its promise is sweeping – to compress the discovery cycle, expand chemical space, and predict biology before experimentation. Yet, early evidence suggests that AI's measurable impact on clinical success is still unproven.

This paper argues that even if AI drug discovery initially mirrors the pattern seen with HTS, its broader capabilities provide a path beyond candidate discovery, enabling improved clinical outcomes and potentially shifting how therapeutic success is measured and achieved. Thus, early ROI similarities between HTS and AI-enabled discovery are likely to diverge as more data becomes available, and economic returns will be realized in materially different ways.

1. The HTS Lessons

In the 1990s, HTS was marketed as the engine that would "industrialize" biology. Automation, robotics, and microplates would enable millions of compounds to be screened rapidly. The logic was statistical: more shots on goal, more hits.

In practice, HTS generated immense volumes of data but modest progress toward clinical drugs. Common issues included the following:

- Library bias: Limited chemical diversity based on prior projects and published scaffolds, even with 106 compounds[1]
- Assay bias: Favored programs that fit into company assay automation over programs with clearer biological relevance and validation. Many assays had low signal-to-noise ratios producing many false positives
- Translational gaps: Poor physiological relevance of hit/lead series for ADMET, PK/PD, delivery and tissue targeting, and ultimately clinical efficacy

By the mid-2000s, HTS-focused start-ups, which touted game-changing drug generation abilities throughout the 1990s, began to restructure or be acquired. For emerging public HTS companies, the average time to consolidation from IPO was just 5.5 years (Table 1).

Company	IPO Year	Reorg Year	Years	Paid-In Capital	Exit Value	Value Delta	Consolidation Details
Millennium Pharma	1996	2004	7.5	$4,728	$8,800	86%	Acquired*
Pharmacopia	1995	2004	8.1	$53	$70	32%	Acquired*
Axys Pharma (Sequana/Arris)	1995	2001	6.2	$359	$156	-57%	Acquired*
Aurora Biosciences	1997	2001	4.6	$145	$592	308%	Acquired*
Evotec Biosystems	1999	2000	1.1	$139	$500	261%	Merged*
Discovery Partners	2000	2006	6.6	$109	$5	-95%	Acquired*
LION Bioscience	2000	2005	5.1	$302	$43	-86%	Reverse merger*
BioFocus	2000	2005	5.2	$56	$38	-32%	Acquired*

Table 1. Eight public HTS companies from the late 1990s and early 2000s identified by Elm. Additional Paid-in Capital (APIC) was gathered from publicly available fillings where available (For ex-US companies, Comon stock and Premium were used). For Biofocus Plc, an estimate from changes in the Galapagos NV balance sheet was used.
* Go to page 15 for more complete consolidation details.

Millennium Pharmaceuticals, then the leader in adopting HTS for novel targets, brought in billions from investors and licensing deal commitments. Their strategy involved running hundreds of screening campaigns to yield an enviable pipeline by pharma standards. Eventually, post-boom realities set in for Millennium, leading to restructuring and downsizing, without a single drug approved from these early HTS campaigns.[2]

Large pharma and established biotechs also began to build HTS platforms in-house (Pfizer's Groton, CT facility, GSK's Affymax integration). Eventually, HTS became an enabling infrastructure tool rather than a truly transformative discovery engine. Companies returned to a combination of rational design and structure-guided approaches alongside screening. Arguably, HTS filled the gap left from plucking all the low-hanging fruit of targets and natural structures discovered in the last century and thus kept the industry on its feet.

The net result was not an increase in clinical success rates, but rather a change in process that simply sustained the industry. In terms of industry output per dollar spent, the industry saw higher R&D costs per drug approved.[3] Although cheaper per assay vs. low throughput, wet-lab automation infrastructure and lab staff still added substantial costs for discovery as the volume of assays increased. In sum, greater spending on HTS did not produce more small molecule drugs at faster rates.[4]

Some of this can be seen in the data on R&D spending in 2004 vs. 2024. If there were an increase in productivity, albeit with a long lag, it would likely be seen in the drug approval output vs. spend across the industry. In 2004, 31 drugs were approved with $83 billion spent on R&D (inflation adjusted to 2024). Last year, $288 billion was spent on R&D and 50 drugs were approved. Comparing ratios as a yardstick on productivity: $2.7 billion per drug during the HTS rollout to $5.7 billion per drug after HTS integration. Though some portion of this increase in cost is due to more complex clinical trial design, it does not change the key observation that HTS did not lead to revolutionary improvements in costs or efficiencies.

HTS net result was not an increase in clinical success rates, but rather a change in process that simply sustained the industry

2. Signs of a Familiar Pattern: AI's HTS Parallel

AI now occupies HTS's old pedestal – this time powered by data rather than robotics. Its key claims are speed, novelty, and cost efficiency. Investment has poured into AI just like it did in the late 1990s for HTS, and large companies are busy integrating AI into their R&D process.

AI drug discovery partnerships now total more than $50 billion in commitments since 2017, and large pharmaceutical companies are integrating AI strategies across departments. Two of the early pioneers in AI drug discovery, Recursion and Exscientia, brought in nearly $2 billion in investment and more than $20 billion in licensing commitments.

Nearly five years out from a spike in biotech investment and an initial wave of AI funding, we are seeing signs of underperformance and consolidation. Recursion and Exscientia have now merged after initial clinical setbacks. BenevolentAI, another frontrunner in the AI space, recently restructured and delisted from the public exchange after falling 98% from peak stock levels. Smaller start-ups have also restructured: Atomwise downsized and renamed to Numerion Labs, and twoXAR rebranded as Aria Pharmaceutical, but has since folded. Although it is early by R&D timeline standards to determine the net contribution of AI in drug discovery, the initial volley of drugs moved into the clinic has not yet outperformed broad industry success rates.

3. Is this time different?

The recent boom cycle seen in AI, as much as it may echo with what was seen with HTS in the late 1990s and early 2000s, might be best viewed as the end of a chapter rather than the end of a book.

On the surface, AI's drug discovery narrative mirrors HTS's early hype about library scale and speed, as AI assumes that data scale yields success. However, unlike HTS, where the number of screens and hits was the main success metric, AI's role in drug discovery is much broader, with layers of applicability that make it less comparable to a single node along the R&D pathway.

AI, unlike HTS, can be harnessed before and after molecule discovery – from target identification to clinical trial enrollment. AI can self-correct, learn across biology-chemistry interfaces, and operate under emerging FDA frameworks for preclinical validation and trial design.

Even within the molecule discovery node itself, AI's feedback loops and hybrid human-machine collaboration can yield deeper mechanistic insight at scale. From a numbers perspective, the chemical space addressed with AI is orders of magnitude larger. HTS libraries were, and many still are, massive by human standards, but they also have limitations as mentioned above.

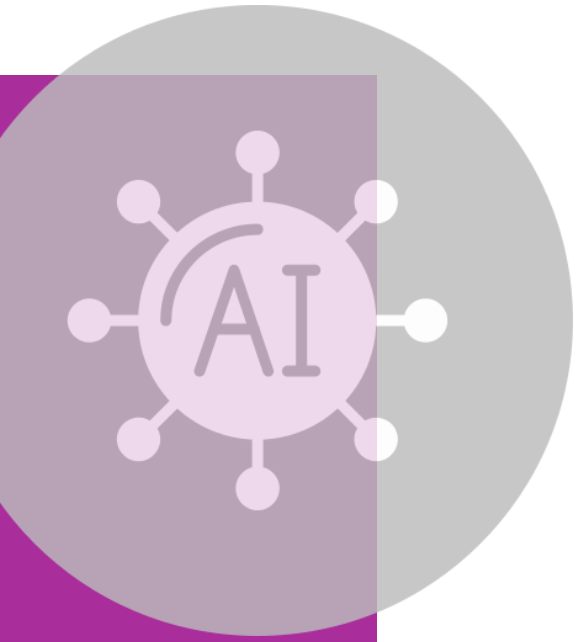

Secondly, new AI breakthroughs in structural biology greatly expand the number of druggable enzymes and receptor targets. Many of these targets intended for HTS could not be purified for assay in a wet lab (for HTS) or crystallography (hit to lead CAD) and thus remained outside the scope of traditional med chem efforts. Protein structure prediction models, such as Alphafold3, ESMFold, RoseTTAFold All-Atom, Boltz-1 have expanded the range of what is druggable and given AI modeling more to build on.

The "Hit to Lead" process has also fundamentally changed with AI, as dependency on a chemistry team's knowledge and algorithmic software packages are no longer limiting factors to creative chemical design. Non-obvious structure activity relationship (SAR) trends can be identified with AI by integrating millions of data points across unrelated scaffolds. This integrated search of chemical space extends to properties beyond target binding such as bioavailability and liver toxicity. Generative AI platforms like Insilico's Chemistry42, Exscientia's Centaur Chemist, and Nvidia's BioNeMo are aiming for faster lead optimization and better ADMET triage than what has existed prior. This creative side of AI is predicted to become more efficient each year.

Lastly, speed and cost are demonstrably different. There are now several examples of programs reaching IND in less than two years, with only a few hundred chemicals being tested in the wet lab.[6] The cost reduction from moving to *in silico* with a reduced workforce and lower CapEx is likely substantial. However, quantifying the cost savings from the shift away from large med chem teams and robotic equipment to smaller teams and low throughput validation assay capabilities is difficult.

	Pre-HTS	HTS	AI
Chemical Space	10^3 Natural product derivatives (ex: statins, SSRIs)	10^6 Higher volume library offering 10x speed of low throughput	10^{60} Order of magnitude > HTS
Target	Only known established targets (low hanging fruit)	Some novel targets (if wet lab isolated or cell assay)	All novel targets (no wet lab purification needed)
Hit to Lead	Limited to similar scaffolds found in nature, Non-parallel programs	False positives, biased library, dependent on chemists and software packages	Not dependent on chemists but rather on wet lab data feedback, quality, and AI learning
ADMET Insight	Low	Medium (secondary screens)	High
Cost	High	Average	Low
Time (Target to IND)	5-10 Years	4 Years	<2 Years
Clinical Success (NDA)	<10%	<10%	TBD

Table 2. Comparison across drug discovery paradigms: pre-HTS, HTS, and AI.

Are we seeing a difference in investor value creation? Of the eight HTS public companies listed in **Table 1**, four were acquired above the amount invested. The remaining four had exit values that were well below the total amount invested. Arguably, much of the discrepancy for this 50/50 split is based on the timing of the exits, rather than clinical failures. For example, Aurora was sold in 2001, well ahead of the 2002 biotech market nadir, whereas Lion Biosciences was acquired in 2005 after a 98% stock drop from its peak. Clinical failures were not the driving force for these cases as neither company reached the clinic.

Only two of the eight companies from the HTS cohort reached clinical-stage. OF the 13 clinical programs that began in the early 2000s, only two advanced to Phase 3, and none ultimately achieved FDA approval. Post acquisition and integration, a number of molecules reached clinical trials, but most failed to secure FDA approval. A notable exception was Aurora's screening for CFTR modulators, which eventually led to Vertex's blockbuster cystic fibrosis drug Kalydeco (ivacaftor).[7]

Comparing the fate of that HTS cohort to today's publicly traded cohort of AI drug discovery companies in **Table 3**, the current situation is similar from a timing basis. Many AI companies were able to raise substantial capital during the COVID years (2020-2022), but many have since seen precipitous stock performance, mirroring the 2000 genomics bubble and subsequent timing of the fallout years later.

Company	IPO Year	Reorg Year	Years Public	Total Raised	Exit Value / Mkt Cap	Value Delta	Consolidation Details
Exscientia	2021	2024	3.2	$879	$688	-22%	Acquired by Recursion
BenevolentAI	2022	2025	2.8	$528	$11	-98%	Delisted 2025, acquired by Otsuka
Relay Tx	2020	Active	5.3	$1,912	$1,173	-39%	Active
Recursion Pharma	2021	Active	4.6	$2,598	$2,174	-16%	Active (w/Exscientia APIC)
ABSCI	2021	Active	4.3	$623	$441	-29%	Active
XTALPI	2024	Active	1.4	$1,693	$5,344	216%	Active

Table 3. Six publicly traded AI drug discovery companies chosen by core focus and IPO date range. *Total raised amounts are Elm's estimate of net proceeds from capital raised based on gross amounts from press releases or Pitchbook. Exit value is based on reported price paid, market capitalizations are as of November 21st, 2025.

Market conditions aside, clinical failures have impacted half of the public AI companies listed in **Table 3**. Exscientia, Recursion and Benevolent AI have had numerous clinical failures and have yet to reach Phase 3, even as post-merger entities. Although Relay has seen clinical failures, one program has reached Phase 3. The outlier in the group is China-based XtalPi Holdings, which has no clinical programs and has survived the stock market downturn. XtalPi has gained in value through diversification and remaining a platform-based company (similar to HTS companies like Biofocus Plc and Evotec Biosystems AG). XtalPi's diverse platform approach spans quantum computing, robotics, and material science, while maintaining a respectable cadence of pharma partnerships (approximately 30 major partnerships, compared with Recursion's 7). This model has enabled XtalPi to generate substantial revenue while avoiding the capital burn and binary risk associated with running clinical trials.

Of the six companies listed, only one trades well above invested capital and only one traded well below. The rest trade below invested capital but given the recent upturn in the XBI index, the data set could ultimately approach the 50/50 success rate mark witnessed during the HTS cycle.

4. Defining Success with AI

How do we know if AI is capturing the value it promises? One way to assess the impact of AI on R&D is to use a two-pronged approach: Chemistry vs. Biology metrics.

Chemistry: Time is money in R&D, both at the discovery phase cash outlay and from a first-to-market perspective. Metrics that move away from hit and lead generation numbers to metrics surrounding INDs per program per year would be more meaningful. This metric covers molecule design time (hit to lead) but also captures preclinical model throughput which is anticipated to become more digital and less animal dependent. This metric will also incorporate a rising trend in AI: the streamlining of data management and regulatory-filing documentation. Paperwork flow should realize massive upside from LLMs, and this metric of INDs per year would capture that.

What would success look like for AI Chemistry? Past screening programs averaged 4 years from target isolation to IND, or 0.25 INDs per year, as disclosed by Pfizer.[8] New evidence suggests AI programs can compress target to IND to under 2 years, and thus would show a shift from 0.25 INDs per year to 0.5 INDs per year.[9] This is likely unprecedented in small molecule discovery and, cost savings aside, this shift in time to market could significantly impact value creation early in the R&D timeline.

Biology: Although time efficiency will be optimized by AI trial enrollment and data sharing, there will always be a fixed time required to show disease mitigation in the clinical trial setting. For example, reaching overall survival (OS) or progression-free survival (PFS) endpoints require protocol-defined assessment windows that cannot be shortened by process acceleration. Thus, a metric such as the proportion of candidates that reach a Phase 2 data readout may prove better than time-based metrics. The argument here is that value creation resides with endpoints over incremental time savings in the clinic. Cutting a month or two off a Phase 2 will not have the same impact as 24 months mentioned above, nor the impact of a positive readout.

What would success look like for AI Biology? Although no AI-originated drugs have made it into Phase III to date, we are still in the early stages of applying this technology to drug development. As of today, clinical translation from IND to drug efficacy remains limited: Of 20 AI-designed drugs that have entered clinical trials, 5 have been suspended, 13 are ongoing in Phase 1, and only one has shown positive Phase 2a results (Insilico's INS018_055 for IPF).

AI-Enabled R&D requires a new predictive process framework

Currently, the above numbers suggest 1 of 6 transitions have made it to positive Phase 2 readout, or 16.6%. Previous publications show IND to Phase II success at 15% (52.0% x 28.9% for Phase 1 x Phase 2 probabilities of success). This is certainly not the heralded game changer, but caveats exist. As mentioned above, given the short timeframe from when these were launched (in just the last few years), we have just a handful of trials. Secondly, a few of these early starts were quick AI entries into IND using monoamine drug scaffolds of well-known targets and one was terminated before Phase 2 started for business reasons.

The ultimate scorecard, beyond the near-term metrics above, is the number of FDA approvals stemming from AI-originated programs. A decade or two out, this will be the lens by which we will judge the AI era, just like we have done here with HTS.

AI can go beyond HTS's limitations by focusing on biology in addition to chemistry

Conclusion

AI stands where HTS once did: revolutionizing process, but not yet outcomes. However, AI can go beyond HTS's limitations by focusing on biology in addition to chemistry. Instead of focusing investment strategy on optimizing target-inhibitor interaction, an emphasis on decoding target-patient interaction could create meaningful value. In other words, molecule discovery using AI may mirror the HTS era unless the full stack of biology AI is adopted. Whereas HTS helped move faster from target to "lead," AI has the potential to take us from target all the way to "drug."

To move beyond hits and leads, additional investment should be focused on two fronts:

1. Biologic systems - AI models that outperform animal models and cell assays in predicting human cell, tissue, organ, and physiological responses

2. Patient stratification – AI models that improve clinical probability of success through biologically informed patient selection

Encouragingly, regulatory agencies such as the FDA are signaling openness to AI use across both domains, from IND to trial design. They understand that AI success will come with its ability to translate computational efficiency into biological insight.

Still, the field must insist on rigorous validation and shared benchmarking before valuations run ahead of evidence. To date, only a few clinical-phase programs have harnessed both molecular and biologic AI tools to ensure success. It may just be a few years out, but final proof will be the tally of FDA approved molecules that benefited from AI across the drug discovery and development spectrum. Until then, metrics such as time to IND and Phase 2 success should point to internal efficiencies on this journey.

HTS yielded many incredible inhibitors but few drugs. The lesson: throughput cannot substitute for biological understanding. The challenge then lies in unraveling the biology of disease, which we have learned is highly differentiated at the patient level. This means determining the right target for the right patient alongside working on the right drug for the right target.

So, is AI another HTS? No, not if AI is harnessed beyond chemistry into biology. The cyclical, learning nature of AI and its uses beyond singular early-R&D metrics are likely to make AI-enabled discovery ROI far more positive than HTS was. Though there is no current data to prove this definitively, the tea leaves are likely sanguine for investors and the biopharma industry.

About the Authors

Thomas Foster, Partner
tf@elminnovation.com
Thomas is a strategist who has focuses on commercialization and pricing & market access. He has received the PMSA Lifetime Achievement Award and has contributed The Economist, Wall Street Journal & others. Previously, he was VP of Consulting for IQVIA and Scholar at MIT Sloan.
Education: Brown University, BA

Kayvon Namvar, Partner
kn@elminnovation.com
Kayvon is an industry expert and financial analyst. He has forecast and valued products, patents, companies, complex securities and royalty interests for a full range of purposes. Previously, he was Principal at RNA Advisors and held senior positions at Hawthorne Health, BPM, and FTI.
Education: University of Southern California, BS

Jeff Williamson, Partner
jw@elminnovation.com
Jeff is a basic scientist and strategist whose work focuses on market assessment, competitive analysis, and forecasting. Additionally, he provides due diligence support for both the buy-side and sell-side. Previously, he held senior positions at RNA Advisors, LEK, and Trinity Pharma Solutions.
Education: UC Berkeley Haas, MBA | Cornell University, BA

Acknowledgments : We are thankful for the content contributions from David Thomas, Paul Saatsoglou, Laahn Foster MD, Meg Richards PhD, and Meera Rao PhD.

About Elm Innovation Advisors

Elm is a strategic and financial advisory firm. We serve the life sciences and healthcare sectors - alongside the investors, bankers, and legal advisors who support them. We guide early-stage and mature companies through moments of significant consequence and operate at the intersection of science, strategy, and capital.

- Strategy
- Transaction Support
- Valuation
- Advanced Modeling
- Market Characterization
- Pricing & Market Access
- Internal Operations
- Dispute Resolution

Endnotes

[1] Bleicher KH, et.al. Hit and lead generation: beyond high-throughput screening. Nature Reviews Drug Discovery. 2003; 2(5):369–378.

[2] Examples of HTS to approval programs include Incyte's JAK inhibitors and Vertex's CFTR drugs

[3] Scannell JW, et.al. Diagnosing the decline in pharmaceutical R&D efficiency. Nature Reviews Drug Discovery. 2012; 11(3):191–200.

[4] Much of the increase in FDA approval totals in the recent decade have come from a) more biologics (cell and gene therapy, antibodies, RNAi) approvals, and b) the net result of a larger industry R&D spend, number of biotechs, and expansion of the small molecule pipeline volume.

[5] Evaluate Pharma, World Preview 2025 report for 2024 data. CBO report, April 2021, Research & Development in the Pharmaceutical Industry.

[6] For biologics, the speed and cost benefits from AI have a steeper challenge. The antibodies generated during COVID-19 by Regeneron and Eli Lilly reached IND within 5 months from target ID. Thus, improvements promised by AI protein design will need to be measured beyond time to IND, and instead focus on comparative progress on hard to target antigens, selective epitopes, and desired on/off rates where traditional wet lab methods have failed.

[7] Welsh, M.J., The arc of discovery, from the description of cystic fibrosis to effective treatments. J Clin Invest. 2024 Oct 1;134(19)

[8] Pereira DA, et.al. Origin and evolution of high throughput screening. British Journal of Pharmacology. 2007; 152(1):53–61.

[9] Recursion's REC-1245 (18 months), Exscientia's EXAI's DSP-1181 (12 months), EXAI's EXS-21546 (9 months), Insilico's INS018_055 (18 months).

* **Millennium** - Restructured to oncology focus (2004), acquired by Takeda (2008)
* **Pharmacopeia** - Restructured into two entities (2004, PDD and Accelrys), PDD acquired by Ligand (2008)
* **Axys Pharmaceuticals (Sequana/Arris)** - Axys (from Sequana/Arris merger) acquired by Celera Genomics (2001)
* **Aurora Biosciences** - Acquired by Vertex (2001)
* **Evotec Bio** - Merged with Asymmetry International (OAI) (2000) to become Evotec AG
* **Discovery Partners International (DPI)** - Restructured after Pfizer deal ended, sold HTS to Galapagos, acquired by Infinity (2006)
* **LION Bioscience** - Reverse merger with SYGNIS Pharma AG (2005)
* **BioFocus Plc** - Acquired by Galapagos NV (2005)

elm

Innovation Advisors

elminnovation.com
8400 West Sunset Road
Suite 300
Las Vegas, NV 89113
USA