



AI Drug Discovery

Who Sells the Picks and Shovels?

If AI accelerates drug discovery, where does scarcity — and economic value — actually move?

Public Research #001 · August 2026

Pál Darabos · MIOS Research

miosresearch.com

1. Executive Summary

Artificial intelligence has demonstrated meaningful productivity gains in the earliest stages of drug discovery. Prospective, experimentally validated studies show that AI systems can generate novel, synthesizable and biologically active molecules and improve chemical-space exploration for selected targets.¹²³ Yet the evidence weakens sharply further down the development chain: Computational Productivity is PROVEN, Discovery Productivity is EMERGING, and Drug-Development Productivity remains UNPROVEN.

That distinction changes the investment question. The original hypothesis was straightforward: if AI generates dramatically more viable drug candidates, downstream physical and clinical systems may become new bottlenecks. But the evidence does not support Candidate Flood as the baseline case. An equally plausible outcome is Better, Not More: improved selection may allow more digital hypotheses to produce fewer, but better, physical development programs.

The downstream system is not passive. AI and automation may increase demand for experiments and clinical execution while simultaneously expanding the capacity to perform them. The relevant variable is therefore not gross downstream demand, but persistent net scarcity — demand that remains after AI-enabled efficiency and capacity expansion are taken into account.

MIOS View — A bottleneck is not an investment thesis. The classic picks-and-shovels framing works only if scarcity persists, transmits economically to a company, can be captured rather than passed through, remains defensible, and is not already fully reflected in valuation. The most interesting exposures may therefore not be the companies that simply process more candidates, but the systems that remain difficult to scale after AI is deployed — and that can retain the economic value of that scarcity.

2. The Question

The starting point is simple. If AI can rank biological targets, design molecules and select potential drug candidates faster and more cheaply, what happens to the systems that must turn those digital outputs into biological, preclinical and ultimately clinical evidence?

AI/model capability ↑ → candidate-generation productivity ↑ → experimental and clinical demand ↑ → downstream bottlenecks → scarcity → potential value capture

This was the starting hypothesis — not the conclusion.

At the outset, we also stress-tested an extreme version of the thesis: what if AI eventually pushed ten times as many viable drug candidates into development? The 10× case was never a forecast or a baseline. It was a thought experiment designed to expose where capacity constraints might appear.

The research therefore has to answer two separate questions: where can persistent scarcity emerge, and who can turn that scarcity into durable economic value?

Refined research question — If AI materially improves productivity in the early stages of drug discovery, which downstream resources become relatively scarcer after AI-enabled efficiency and capacity gains are considered — and where can that scarcity translate into durable economic value?

3. Is AI Actually Accelerating Drug Discovery?

“AI is accelerating drug discovery” is too broad a statement to evaluate as a single claim. We separate three levels of productivity: Computational Productivity, Discovery Productivity and Drug-Development Productivity. The evidence gives a different answer at each level.

3.1 Computational Productivity — PROVEN

The strongest evidence is at the computational and early experimental layers. This is no longer limited to retrospective benchmarks: in multiple prospective studies, AI-proposed molecules were actually synthesized and tested in the laboratory.^{1,2,3}

The GENTRL/DDR1 work demonstrated AI-generated molecules that were synthesized and experimentally active. SyntheMol produced structurally novel compounds with experimentally observed antibacterial activity. Lyu and colleagues showed that AI-predicted protein structures could be used in prospective virtual screening to identify experimentally validated ligands for selected targets.^{1,2,3}

These are strong demonstrations that AI can improve selected early-discovery tasks. They do not show that AI generally produces better clinical drugs, nor that industry-wide R&D productivity has already improved.

3.2 Discovery Productivity — EMERGING

The next evidentiary step is harder: does better computational output translate into developable drug candidates faster, or with a higher probability of success? Faster molecule design is not the same as better drug-development productivity.

Rentosertib is an important step forward because an AI-assisted target-and-molecule program reached a randomized Phase 2a study. The 71-patient, 12-week study provided early clinical signals, but safety was the primary endpoint. It is therefore evidence of clinical progression — not proof that AI-designed drugs deliver superior clinical efficacy.⁴

Rentosertib	Research boundary
What it proves	An AI-assisted program can progress into meaningful randomized clinical testing.
What it does not prove	AI-designed drugs are clinically superior.


3.3 Drug-Development Productivity — UNPROVEN

A 2024 analysis of a small cohort of AI-originated clinical programs reported a notable Phase I success signal, while the Phase II result was much less conclusive relative to historical benchmarks.⁵

Early AI-native cohort	Reported result
Completed Phase I	21 / 24 successful (N=24)
Completed Phase II	4 / 10 successful (N=10)

The Phase I result is notable. The Phase II evidence does not show a clear advantage over historical baselines, and the cohort is small and selected; selection and publication bias cannot be ruled out.

A 2026 critical review points to the same evidence gap: beyond model benchmarks and selected programs, the field still lacks sufficiently mature clinical evidence to establish a broad AI-driven uplift in drug-development productivity.⁶



The evidence is therefore asymmetric. The closer we move from digital molecule design toward clinical outcomes, the weaker the evidence becomes. AI has already demonstrated meaningful improvements in selected early-discovery processes; it has not yet demonstrated that those improvements translate into higher productivity across the full drug-development system or into more successful medicines.


Taken together, the evidence supports an emerging but not yet proven improvement in discovery productivity.

4. The Drug-Development Funnel

A promising molecule is still a long way from an approved medicine. After early discovery, it must move through preclinical development, regulatory steps and multiple clinical phases — with substantial attrition along the way.

Target → Hit → Lead → Candidate → IND → Phase I → Phase II → Phase III → NDA/BLA → Approval

In the BIO/Informa/QLS 2011–2020 historical dataset, transition success was 52.0% from Phase I to Phase II, 28.9% from Phase II to Phase III, 57.8% from Phase III to NDA/BLA, and 90.6% from NDA/BLA to approval. Overall Phase I-to-approval success was 7.9%.⁷



These percentages are conditional transition probabilities for programs that have already reached each stage. The largest attrition in this dataset occurs at Phase II → Phase III, where only 28.9% advanced. This is a historical, dataset-specific benchmark — not a universal biological constant.

The downstream process is not only selective; it is also long and capital intensive. In the same 2011–2020 historical cohort, average Phase I-to-approval time was approximately 10.5 years. In a separate cohort of pivotal efficacy trials supporting FDA approvals in 2015–2016, estimated median direct trial cost was approximately \$19 million, with a very wide range of roughly \$2 million to \$347 million.^{7,8}

Historical benchmark — not a 2026 industry estimate.

Historical warning: Eroom's Law

Drug R&D history is a warning against assuming that better technology automatically produces better end-to-end productivity. Scientific and technological progress has historically coexisted with high attrition, rising therapeutic standards and regulatory complexity.⁹

Historical framing — not a claim that the trend must continue.

The downstream system is therefore long, expensive and highly selective. But that does not mean faster AI discovery automatically creates more downstream volume. If AI mainly improves selection, it could reduce the number of weak programs entering the most expensive stages of development.

That leads to the next question: Candidate Flood — or Better, Not More?

5. Candidate Flood or Better, Not More?

The simplest scenario is that AI creates more drug candidates, more of them advance into preclinical and clinical development, and new capacity constraints emerge. This is the Candidate Flood world, and at first glance it provides the strongest foundation for a conventional picks-and-shovels thesis.


The problem is that there is not yet sufficient evidence to treat this as the baseline.

10× candidate generation ≠ 10× viable clinical candidates.

Digital hypothesis generation can scale at relatively low marginal cost, but a selection layer sits between digital output and physical workload. Ranking, filtering, developability prediction, toxicity and risk prediction, and patient or biomarker selection can all reduce the probability that more digital output translates into a proportional increase in physical programs.

FIGURE 3 — THREE POSSIBLE WORLDS

AI Discovery Productivity Can Lead to Three Very Different Worlds



Note: Scenarios are mutually exclusive and represent directional hypotheses.

Source: See accompanying publication and References.

MIOS RESEARCH | PUBLIC RESEARCH #001 | AUGUST 2026

5.1 Historical analogue: biomarker preselection

Historical drug-development data provide a useful analogy for Better, Not More. In the 2011–2020 cohort, programs using biomarker-based patient preselection had materially higher transition success rates than programs without biomarkers: 15.9% versus 7.6% from Phase I to approval, and 46.3% versus 28.3% from Phase II to Phase III.⁷

2011–2020 historical cohort	Biomarker	No biomarker
Phase I → Approval	15.9%	7.6%
Phase II → Phase III	46.3%	28.3%

Historical analogy — not evidence of AI performance.

The point is not that AI will reproduce these numbers. The analogy shows that better selection can materially change funnel outcomes without requiring more programs to pass through the system.

On current evidence, Candidate Flood should not be treated as the baseline. Better, Not More is at least as plausible, while a Digital Productivity Trap (strong computational progress with limited clinical translation) cannot yet be excluded.

For the downstream investment thesis, the key question is not how many candidates AI can generate digitally. It is how many additional physical development programs AI ultimately creates.

6. The Bottleneck Migration


Even if AI increases downstream workload, a new bottleneck does not necessarily emerge.

For scarcity to increase, demand must grow faster than available capacity.

This matters especially in AI because the same technology wave that may create new demand can also expand supply: automating experiments, improving clinical-site selection and patient recruitment, and reducing the resource intensity of execution.

Net Bottleneck Migration = AI-Driven Demand Expansion – AI-Enabled Efficiency & Capacity Expansion

MIOS analytical framework — not an empirical equation.



Qualitative directional markers only; not quantitative estimates. The same AI wave may increase both downstream demand and downstream capacity.

6.1 Clinical Execution — CONDITIONAL

Patient recruitment and site access are among the best-documented real-world constraints in clinical execution. More clinical programs could increase demand on this infrastructure. At the same time, AI-assisted patient-trial matching and recruitment may improve throughput, although current evidence remains heterogeneous and requires further real-world validation. Clinical execution is constrained today; an AI-driven increase in net scarcity is not yet proven.¹²

6.2 Biological Validation — EMERGING

The value of a digital prediction ultimately depends on whether it survives contact with biological reality. Disease-relevant, reproducible biological validation therefore remains necessary. Necessity, however, is not the same as scarcity.

Self-driving laboratories, robotics and closed-loop experimentation may expand experimental capacity and reduce cost per experiment. The 2026 peer-reviewed evidence also highlights important limitations in scale, generalizability and provenance; these systems should not be treated as system-wide substitutes for biological validation.^{10,11}

AI-native discovery platforms may also internalize or automate part of the validation layer themselves. In other words, the same platforms generating more hypotheses may expand the capacity to test them.

6.3 Regulated Preclinical Validation — AMBIGUOUS

Regulatory-grade toxicology, safety and other IND-enabling evidence remain necessary. More viable programs could increase demand, while better early selection and automation could reduce the physical work required per program. The net volume effect cannot currently be established with confidence.

6.4 Longitudinal Multimodal Data — EMERGING

Not all data are equally valuable to AI. Longitudinal datasets that connect multiple biological and clinical modalities to real-world outcomes may be substantially harder to create than additional raw data. That makes them a plausible source of scarcity.

But scarcity and moat are not yet proven. Interoperability, data federation, public datasets and new data-generation technologies may reduce exclusivity over time.

6.5 Raw Biological Measurement — NOT PROVEN

Rising demand for biological measurement does not automatically make measurement capacity scarcer. If throughput rises, unit costs fall, competition intensifies and new capacity is added, even strong volume growth may fail to create scarcity rents.

Demand growth $\neq$ scarcity.

6.6 A hidden bottleneck: Capital Allocation

Even if AI creates more developable programs, someone still has to fund and prioritize them. If available capital and development budgets do not expand at the same rate, some candidates will be filtered out before they create downstream physical demand.

The research therefore does not identify a single proven “AI drug discovery bottleneck.” It identifies several conditional scarcity candidates in which demand and capacity are moving at the same time.

That changes the picks-and-shovels question. The task is not to find companies adjacent to a growing downstream market. It is to identify resources whose net scarcity can persist and then determine whether a company can capture the economics.

7. Who Sells the Picks and Shovels?

Identifying a bottleneck still does not tell us who profits from it. The company examples below are therefore not stock recommendations. They are economic-transmission case studies: examples of how a structural industry change may or may not reach a public company's economics.

Structural exposure ≠ AI beneficiary ≠ good investment.

FIGURE 5 — FROM BOTTLENECK EXPOSURE TO ECONOMIC CAPTURE

Structural exposure ≠ AI beneficiary ≠ good investment.

LAYER (BOTTLENECK)	EXAMPLE EXPOSURE (PUBLIC COMPANIES)	WHAT IS SUPPORTED (EVIDENCE TODAY)	MISSING LINK TO ECONOMIC CAPTURE (THE REAL QUESTION)
1  CLINICAL EXECUTION	IQVIA MEDPACE ICON	 Material execution exposure	 AI demand → rent capture
2  REGULATED PRECLINICAL VALIDATION	CHARLES RIVER	 Necessary regulated capability	 AI → persistent volume / scarcity
3  LONGITUDINAL MULTIMODAL DATA	TEMPUS	 Material data monetization	 Data → durable moat
4  RAW BIOLOGICAL MEASUREMENT	ILLUMINA 10x GENOMICS	 Enabling measurement exposure	 Volume → scarcity / pricing power



Company names are economic-transmission case studies, not recommendations.

Source: Company filings, investor presentations, industry reports, expert interviews, MIOS Research.

MIOS RESEARCH

PUBLIC RESEARCH #001

AUGUST 2026

7.1 Clinical Execution — exposure is not rent ownership

IQVIA, Medpace and ICON have direct structural exposure to clinical-trial execution. If more programs enter clinical development, demand for their services could rise. Current evidence supports the structural exposure — not incremental AI-driven demand, nor the ability to retain the economics of that demand.^{13,14,15}

A portion of clinical-trial spending, including certain site, investigator and other third-party costs, can be reimbursable or passed through under CRO contract structures. Higher trial spending therefore does not necessarily translate one-for-one into higher-quality revenue, pricing power or margin expansion for the CRO.^{13,14,15}

The relevant question is not whether more money is spent on clinical trials. It is how much economic rent the service provider can retain.

7.2 Regulated Preclinical Validation — necessity without guaranteed volume

Regulated preclinical infrastructure remains necessary, which gives providers such as Charles River natural structural exposure. But necessity does not prove volume growth: better AI-based selection may prevent weak programs from advancing, while automation may expand available capacity.¹⁷

Necessary infrastructure ≠ automatic beneficiary.

7.3 Longitudinal Multimodal Data — potential moat, not proven moat

Tempus illustrates a different type of exposure. Linking diagnostics, molecular data and clinical connectivity can create longitudinal, multimodal datasets that may be more valuable to AI systems than easily generated raw data.

Proprietary data is not automatically a moat.

Uniqueness → Reuse → Customer dependence → Renewal / expansion → Monetization → Margin → Defensibility

Tempus Q2 2026 company data support the narrower claim that the Data & Applications business already represents material and growing monetization. That is business evidence — not proof of a durable data moat, and not evidence that AI drug discovery caused the growth.¹⁶

7.4 Raw Biological Measurement — the volume trap

Higher biological-measurement volumes initially look like a natural picks-and-shovels opportunity. But if throughput rises quickly, unit costs decline and competition remains strong, higher demand may not create persistent scarcity or pricing power.


Volume beneficiary ≠ scarcity owner.

Exposure to a bottleneck is not the same as ownership of its economics.

The next step is therefore no longer about the industry layer. It is about the company: can structural advantage be translated into pricing, margins, cash flow and a defensible competitive position?

8. From Bottleneck to Investment

A real bottleneck is still not an investment thesis. For a structural industry constraint to become durable shareholder value, it has to pass through a series of economic gates.



Clinical execution illustrates why Economic Transmission matters. Patient or site scarcity may be real while part of the incremental cost remains pass-through. Physical scarcity can therefore exist without equivalent pricing power or margin expansion at the CRO.

The same logic applies to data. A hard-to-replicate data asset is not automatically a durable competitive advantage. The next questions are whether the data can be reused, whether it deepens customer dependence, whether it improves monetization, and whether competitors can reproduce the asset or its utility.

Valuation is the final but indispensable filter. PR #001 is not a company-specific valuation report and does not provide price targets or buy/sell ratings. The methodological conclusion is simpler: even a real bottleneck, a durable moat and strong business economics can still be a poor investment if the market already prices in an even larger fundamental uplift.

The bottleneck is where the research begins — not where the investment decision ends.

9. What Would Prove Us Wrong?

The thesis can fail at more than one point. Each major link in the causal chain is independently testable:

AI capability → discovery productivity → candidate quality / volume → downstream evidence demand → net scarcity → value capture → business economics → investment return

It is not enough for the thesis to be right once. The important links in the chain have to hold.

Kill criterion	What it would mean
NO DISCOVERY UPLIFT	Prospective evidence fails to show meaningful improvements in cycle time, decision quality or experimental efficiency.
NO CLINICAL TRANSLATION	Maturing AI-originated Phase II/III cohorts fail to improve; AI value remains concentrated in early discovery.
BETTER, NOT MORE DOMINATES	Drug-development productivity improves without a broad downstream volume boom; the volume-driven picks-and-shovels thesis weakens.
SUPPLY SCALES FASTER	Automation, self-driving labs, recruitment/site technology and data infrastructure expand capacity faster than AI expands demand.
NO RENT CAPTURE	Scarcity emerges, but pricing, margins or ROIC do not improve, or the economics accrue elsewhere.
VALUATION ABSORBS THE THESIS	Market expectations already imply a larger fundamental uplift than the evidence supports.

Conclusion — MIOS View

AI has already demonstrated that it can improve selected computational and early experimental processes in drug discovery. There is not yet sufficient evidence that these gains translate into higher productivity across the full drug-development system or into more successful medicines.

Candidate Flood should therefore not be treated as a proven baseline. Better, Not More is at least as plausible: AI may create value primarily by improving selection, allowing more digital hypotheses to produce fewer, but better, development programs.

The downstream investment opportunity is therefore not determined by gross demand growth. The key variable is persistent net scarcity: which resources experience demand growth faster than AI and automation can expand their capacity?

Even that is not enough. Scarcity must transmit economically to a company; the company must be able to retain the value; the advantage must be defensible; and valuation still matters.

That is where the picks and shovels may be — if they exist at all.

References & Source Notes

1. Zhavoronkov et al. (2019), Nature Biotechnology — GENTRL / DDR1 prospective molecule generation. doi:10.1038/s41587-019-0224-x
2. Swanson et al. (2024), Nature Machine Intelligence — SyntheMol prospective antibiotic discovery. doi:10.1038/s42256-024-00809-7
3. Lyu et al. (2024), Science — AI-predicted structures for prospective ligand discovery. doi:10.1126/science.adn6354
4. Xu et al. (2025), Nature Medicine — rentosertib randomized Phase 2a. doi:10.1038/s41591-025-03743-2
5. Jayatunga et al. (2024), Drug Discovery Today — AI-originated clinical success analysis. doi:10.1016/j.drudis.2024.104009
6. Bender et al. (2026), Nature Reviews Drug Discovery — critical review of clinically relevant AI impact. doi:10.1038/s41573-026-01496-2
7. BIO / Informa / QLS (2021) — Clinical Development Success Rates and Contributing Factors 2011–2020.
8. Moore et al. (2018), JAMA Internal Medicine — pivotal efficacy trial direct cost distribution. doi:10.1001/jamainternmed.2018.3931
9. Scannell et al. (2012), Nature Reviews Drug Discovery — historical R&D productivity / Eroom framing. doi:10.1038/nrd3681
10. Canty & Abolhasani (2026), Nature Reviews Chemistry — self-driving laboratory review. doi:10.1038/s41570-026-00847-2
11. Pilon et al. (2026), Nature Synthesis — flexible self-driving lab demonstration. doi:10.1038/s44160-026-01053-0
12. Chen et al. (2025), JCO Clinical Cancer Informatics — scoping review of LLM-based patient-trial matching; recruitment bottleneck and emerging AI-assisted matching evidence. doi:10.1200/CCI-25-00071
13. Medpace Holdings, Inc. (2026), Form 10-K for fiscal year ended December 31, 2025 — contract structure, reimbursable third-party costs and investigator-site payments.
14. ICON plc (2026), 2025 IFRS Annual Report — investigator payments and other reimbursable costs under client contracts.
15. IQVIA Holdings Inc. (2026), Form 10-K for fiscal year ended December 31, 2025 — reimbursed expenses and investigator-payment economics in clinical research services.
16. Tempus AI, Inc. (2026), Q2 2026 SEC-filed earnings release — Data & Applications business evidence, not moat proof.
17. Charles River Laboratories International, Inc. (2026), Form 10-K for fiscal year ended December 27, 2025 — regulated and non-regulated Discovery and Safety Assessment services supporting preclinical drug development and regulatory-required safety testing.
18. Illumina, Inc. (2026), Form 10-K for fiscal year ended December 28, 2025 — sequencing instruments, consumables and biological measurement platform exposure.
19. 10x Genomics, Inc. (2026), Form 10-K for fiscal year ended December 31, 2025 — single-cell and spatial instruments and consumables enabling high-resolution biological measurement.