The Change Worth Getting Out of Bed For

Decision Intelligence Is the Next Transformation in Drug Development

SM
Sukant Mittal, PhD
SVP, Corporate Development · QuantHealth

EXECUTIVE SUMMARY

Generative AI adoption in drug R&D is accelerating sharply — early-adopter budgets have expanded significantly over the past 18 months, and nearly half of life sciences leaders plan to raise AI spend by more than 10 percent this year. As infrastructure providers, frontier model companies, and specialized application vendors deepen their presence across the R&D value chain, the industry's focus has shifted: from simple chat interfaces to multi-agent orchestration, and from AI-as-efficiency-tool to AI-as-ROI-driver in complex, human-in-the-loop decision-making.

Sustaining this momentum in a fast-moving environment depends on three things: proving real value, mobilizing the surrounding ecosystem, and shifting the relationship between human experts and AI systems.

1. **Proof of Technology and Value.** The bar for pharma attention is high — one executive noted it's "not worth getting out of bed" for anything under $100 million in impact. That scale of value lives in drug development's hardest problems: near-billion-dollar investments, decade-long timelines, and roughly a 10% probability of success. Model-driven outcome predictions available months or years before trials begin — combined with optimized endpoints and subgroup selection — are improving the odds of success. Given $250–300 million budgets per trial phase, even marginal design improvements compound significantly; the "fast-fail" case frees resources for higher-probability assets.
2. **Mobilizing the Ecosystem.** No single company can shift an entire industry's way of working — the ecosystem has to move together across three layers: **government as legitimizer** (regulators now inviting computational evidence via the MIDD Paired Meeting Program, ICH M15, FDA AI credibility guidance, and Operation TrialBlazer), **infrastructure as foundation** (GPU-scale modeling made accessible through partnerships like NVIDIA's BioNeMo on AWS), and **intelligence as distribution** (frontier model companies specializing, as with Anthropic's Claude Science).
3. **Human-in-the-AI-Loop vs. AI-in-the-Human-Loop.** Trust-building is the slowest but most essential shift. The goal is moving from AI assisting at the margins toward AI taking the first pass while experts review and approve. Professionals still favor human judgment over AI for high-stakes work by roughly nine to one — reasonable caution given the financial stakes. Comfort accrues gradually, one validated decision at a time.

Bottom line. As AI systems demonstrate they can carry workflows to validated results, commercial value is shifting from software access to outcomes — driving a broader move toward milestone-based, risk-sharing deal structures. Companies that build credibility across all three dimensions will define the next era of drug development.

In drug R&D, excitement for generative AI has increased the adoption of technology at an unprecedented rate. In fact, over the past 18 months, budgets for early adopters have expanded, and most life sciences leaders plan additional increases this year — nearly half intend to raise AI spend by more than 10 percent.[1] At the same time, infrastructure players, frontier model companies, and smaller application-focused companies are extending their footprints across the R&D value chain. The focus has also moved from chats to multi-orchestrated agentic solutions to harnesses around agentic frameworks, and the industry has shifted its expectations for AI from efficiency alone to how solutions deliver ROI in complex use cases with humans-in-the-loop.

In a high-velocity environment, accelerating adoption requires three things: proving the technology and delivering value; mobilizing the ecosystem; and getting comfortable with human-in-the-AI-loop in place of AI-in-the-human-loop.

Proof of Technology and Delivering Value

A senior executive from a top 10 pharma put it plainly at a recent QuantHealth industry summit: *"It's not worth getting out of bed for us unless we're talking about an impact worth at least $100 million."* That level of value lives in the hardest problems, specifically drug development, where investments run high (~ billion dollars), approval timelines are long (~ decade), and the probability of success is low (~10 percent).

The industry has long known that changing the trajectory of drug development — from its high costs, timelines, and patient impact — starts with changing which clinical trials are ultimately pursued. Being able to quantify the potential of a drug for smarter decision-making and confidence, though, was largely a theoretical exercise.

But this change is finally upon us. Deep, model-driven outcome predictions and decision intelligence available months or years before trials begin, along with the ability to pinpoint areas of improvement — such as selecting endpoints and subgroups — while iteratively collaborating with experts from various disciplines, can meaningfully help improve the odds of a drug succeeding and have a real impact for patients who invest time and effort into trials.

An important note: the economic value is clear if the drug succeeds, but the "fast-fail" approach reduces the opportunity cost, cutting spend on an asset more likely to fail and making space in the same infrastructure for the ones that might have a better probability to succeed (economic value compounds fast within the lifespan of an asset, and even faster on a portfolio level). This computation, resulting from fit-for-purpose, deep computational models augmented by LLMs, can help guide a decision to retire a drug asset earlier and clear a path to multi-million-dollar savings without strain.

While the examples in drug development of deep frontier technology and models integration are rare, in research recent examples provide guidance: CHAI Discovery's CHAI-2 addressed the genuinely hard problem of designing antibodies for 52 targets with no known binders, with confirmed hits in under two weeks,[3] and Isomorphic's collaboration with Lilly and Novartis[4] — structured as relatively modest upfront milestones (compared to the overall potential payout) that pay only as frontier tech proves itself — are good benchmarks. Earn conviction on a hard problem first by proving the value of your technology, then let the commercial terms follow the evidence.

Mobilizing the Ecosystem

No company alone can prove a new way of working; the surrounding system has to move with it. This is exactly what's happening in drug development now.

> **Government as the legitimizer.** Regulators have moved from tolerating computational evidence to inviting it, and the binding constraint has shifted downstream — from building a model to establishing that its output can be trusted, through the validation, provenance, and context that turn a prediction into evidence. Several initiatives now formalize that credibility layer: the MIDD Paired Meeting Program;[5] the ICH M15 harmonized framework and terminology for how model-derived evidence is submitted and judged;[6] the FDA's risk-based AI credibility guidance for AI used to support decisions on drug safety, efficacy, and quality;[7] and Operation TrialBlazer, an HHS-wide, expedited-IND effort with QSP-based first-in-human dose guidance and updated master-protocol guidance.[8]

> **Infrastructure as the foundation.** The compute to train and run models over genomic, clinical, and real-world data sits with a few providers, now partnering downward into a stack of raw compute, foundation models, and applications. NVIDIA's BioNeMo, running on AWS, puts GPU-scale training within reach of companies that could never build it alone.[9]

> **Intelligence as the distribution.** The frontier model companies are moving from horizontal tools toward specialization. Anthropic's Claude Science — a workbench that connects models to the tools researchers already run, such as PubMed, Jupyter, and open biomolecular models like Boltz-2 — launched alongside an in-house program to pursue drugs for neglected diseases.[10] Specialist connectors extend that reach. This intelligence-layer-as-distribution is the ultimate mobilizer of the ecosystem.

Human-in-the-AI Loop vs. AI-in-the-Human-Loop

The third source of trust takes the longest to build because it is about us. Drug development is multidisciplinary, and therefore a useful technology system has to speak medicine, science, regulation, and engineering, and interface with multiple stakeholders who each hold a piece of the decision.

Only a handful of companies working alongside pharma have assembled solutions that reflect the needs of these diverse groups, and even fewer have the same range of talent as their core DNA. That presence is what moves an organization from AI-in-the-human-loop, where the expert does the work and the model assists at the margins, toward human-in-the-AI-loop, where a frontier model carries the first pass and the expert becomes its reviewer and approver. Professionals still prefer a human to AI for complex, high-stakes work by roughly nine to one,[11] and where a wrong call costs hundreds of millions of dollars, that caution is well placed. Closing the gap is slow work: staying beside the teams as the technology advances, keeping AI in the human's line of sight, and letting comfort accrue one decision at a time.

These necessities — proof, ecosystem, and people — are also changing what is bought and sold. When a system can carry a workflow to a validated result, value moves from accessing software toward the outcome — and commercial terms move toward the milestone-based, risk-sharing structures now forming across AI-driven R&D.[12] This is why QuantHealth has built its product portfolio with a combination of deep foundational models trained by scientific, medical, and biology experts, with recommendations provided with further expertise — or, human-in-the-AI loop.

The companies that earn trust on all three fronts — proving value where the stakes are highest — will be the ones that win, and inspire us to get out of bed in service of bringing life-changing medications to patients: those participating in clinical trials, and those awaiting a successful outcome.

References

1. "State of AI in Healthcare and Life Sciences: 2026 Trends" (NVIDIA, 2026) — AI budget increases.
2. AI-driven clinical-trial simulation and in-silico validation — conference coverage, Clinical Trial Innovation Summit 2025.
3. Chai Discovery, "Zero-shot antibody design in a 24-well plate" — Chai-2; 16% de novo antibody hit rate across 52 targets with no known binders, validated in under two weeks (bioRxiv, 2025).
4. Isomorphic Labs, strategic collaborations with Eli Lilly and Novartis — combined potential value near $3B, modest upfronts against milestone payments (2024).
5. FDA Model-Informed Drug Development (MIDD) Paired Meeting Program made permanent — dose selection, drug–disease–trial simulation, and predictive/mechanistic safety; requests accepted quarterly through June 2027.
6. ICH M15 — harmonized assessment framework and terminology for how model-derived evidence is submitted and judged.
7. FDA, "Considerations for the Use of Artificial Intelligence To Support Regulatory Decision-Making for Drug and Biological Products" — risk-based credibility-assessment framework (draft guidance, Jan 2025).
8. U.S. FDA / HHS Operation TrialBlazer — pilot to speed early-stage clinical trials (STAT, June 2026).
9. "NVIDIA BioNeMo Enables Generative AI for Drug Discovery on AWS" (NVIDIA, 2024).
10. Anthropic, "Claude Science, an AI workbench for scientists," and internal neglected-disease drug-discovery program (2026).
11. ZS, "Scaling AI in pharma and biotech: 2026 CDIO research."
12. BCG, "Reimagining Business Models: Biopharma Trends 2026."