ELIQUENT LIFE SCIENCES

AN EXECUTIVE BRIEFING · LIFE SCIENCES · ARTIFICIAL INTELLIGENCE

Beyond the Model

Why Expert Judgment, Regulatory Intelligence and Trusted Relationships are the Keys to Unlocking Value in AI-Enabled Life Sciences

Authored by Andy Greenberg, Chief AI & Digital Officer ELIQUENT Life Sciences

July 2026

THE GOVERNING IDEA

As AI capabilities expand and commoditize, the greatest value in life sciences is created not by model access but by expert judgment, proprietary intelligence, trusted relationships, and validated AI infrastructure. AI is not eliminating the need for expertise, it is refocusing it, lowering the value of routine, codifiable work and raising the value of judgment, accountability and domain intelligence. This is where competitive advantage follows and, more fundamentally, where impact is realized.

The shift in value realization

In recent years, the strategic conversation in AI has centered on which model is best and what features are next to be deployed. The answers to those questions are losing their force. As frontier performance converges, access barriers fall, and costs decline, the source of value is moving away from the intelligence itself and toward the **people and systems that govern, validate and apply** it.

In regulated life sciences the distinction is decisive. Regulatory submissions, quality decisions, pharmacovigilance assessments and medical content all carry accountability that cannot be delegated to an algorithm. AI may accelerate analysis and generate recommendations, but responsibility for compliance, patient safety and product quality remains human. **The durable asset is rarely the model. It is the judgment, intelligence, relationships and validated infrastructure built around it.** That is where value is created, and where competitive advantage ultimately follows.

This briefing makes a single argument in four parts. **The Elsa story** shows what survives when a model is replaced; **the judgment thesis** explains why value is migrating from reasoning to judgment; **the forces of value** describe the four capabilities that compound into a moat; and a **seven-question diagnostic** gives leaders a practical test. The conclusion for executives is that AI adoption is no longer principally a technology decision. It is an organizational capability decision.

1 The Elsa story: what survives when a model is replaced

The most instructive AI case study in life sciences is not a pharmaceutical company or a technology vendor. It is the regulator. When the U.S. Food and Drug Administration launched Elsa in 2025 — an AI-enabled platform to support regulatory review — it signaled that AI had moved from experimentation into operational practice inside a highly controlled, accountable environment.[1]

Less than a year later, a federal directive required the agency to migrate away from the foundation model on which Elsa had been built and validated.[2] Public attention fixed on the change of provider. The more important lesson lay in everything surrounding the model. The agency had to reassess prompts, retrieval systems, embeddings, evaluation criteria and validation evidence before the system could operate with confidence again. The episode exposed a misconception that persists across most AI programs: **the model is treated as the asset**.

> ***The model is often the most replaceable component of the system. The enduring value lies elsewhere.***

What endures is the expert-defined acceptance criteria that determine whether an output is usable; the governance that controls change; the evidence-tracing architecture that enables explainability; the validation documentation that demonstrates reliability; and the people who understand the consequences of being wrong. These elements survive a model change. The model itself may not.

Two shifts the migration revealed

Intelligence is becoming a commodity. Open-weight alternatives are improving and pricing keeps falling, so differentiation depends on what organizations do with intelligence rather than the intelligence itself[3] — that is, the value they create from it. The question moves from "which model should we use?" to "what capabilities remain valuable regardless of which model we use?"

Regulators are becoming AI-enabled. Submissions must increasingly perform well inside AI-assisted review environments, not only for human reviewers — which means content that is structured, traceable, evidence-based and machine-readable. AI-ready submissions are becoming a strategic capability, not an operational preference.

2 The judgment thesis: value is migrating from reasoning to judgment

Through the current cycle, value has been equated with access to intelligence. As frontier models converge and once-exclusive capabilities become widely available, that assumption is failing. AI is not becoming less important; the source of value is moving **from reasoning to judgment**.

Reasoning and judgment are not the same

AI excels at reasoning — summarizing information, identifying patterns, drafting and synthesizing evidence at superhuman speed. But reasoning alone does not determine outcomes. The most consequential decisions in regulated environments are rarely about finding a technically plausible answer. They are about whether an action should be taken, under what conditions, what risk remains unresolved, and who accepts accountability. Judgment operates where information is incomplete, precedent is ambiguous, trade-offs are unavoidable and the cost of error is asymmetric.

> ***A model can recommend. A professional must decide.***

AI is refocusing expertise, not removing it

Routine activity — initial drafting, literature review, summarization, data extraction — is being automated, and its value is falling. Demand for senior expertise remains strong[4], because organizations still need individuals who interpret regulatory expectations, challenge assumptions, evaluate risk and provide accountability. *A regulatory strategy is a judgment*

about how precedent applies in context; a safety signal is a judgment about causality and action; a quality risk assessment is a judgment about patient impact. The faster AI processes information, the more these judgment layers matter.

EXHIBIT 1 | **AI lowers the value of codifiable work and raises the value of judgment**

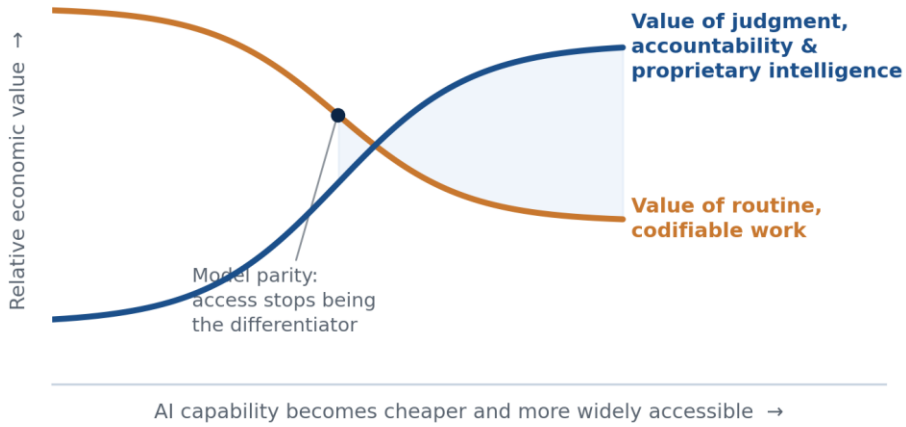

Source: Eliquent Life Sciences analysis, drawing on Stanford HAI, AI Index Report 2026, and Brynjolfsson, Chandar & Chen (Stanford, 2025).

A model trained on public data may know what guidance says, but not how a particular review division will interpret it. These are not failures of AI; they reflect that judgment is fundamentally different from retrieval. The implication is direct: if everyone can access similar intelligence, value is created by combining intelligence with judgment. **AI generates possibilities; experts determine which possibilities become actions.**

3 The forces of value: four capabilities that compound into a moat

Four capabilities convert abundant intelligence into durable value. Individually each creates value; combined, they form a moat that becomes **more valuable as AI itself becomes more available**.

EXHIBIT 2 | **The four capabilities**

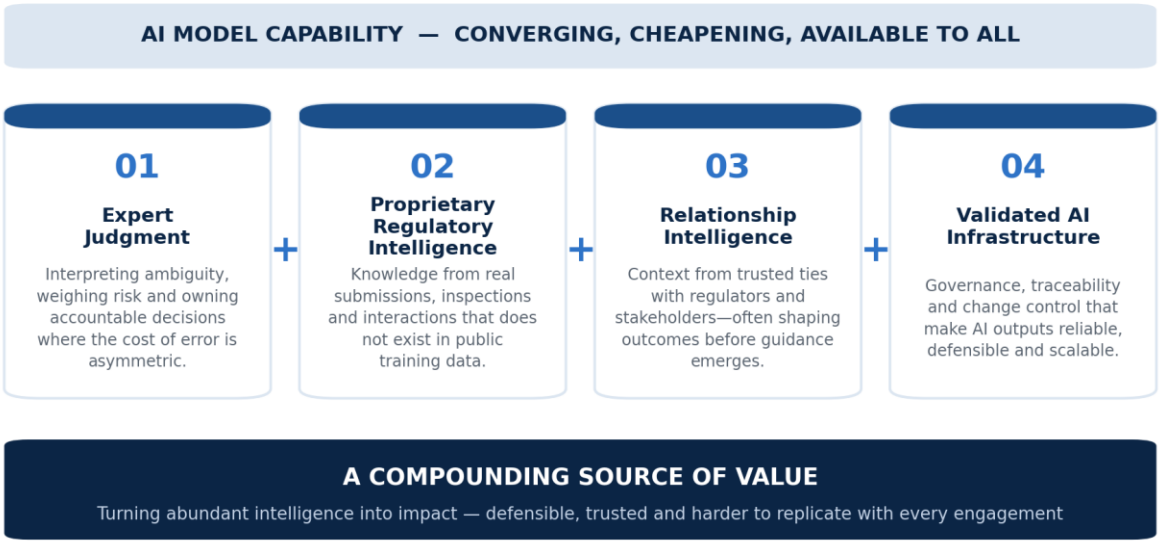

Source: Eliquent Life Sciences analysis.

The stakes are higher in regulated life sciences

In most sectors, an AI error is a commercial problem. In regulated life sciences it can affect patient safety, product quality, approvals and credibility with health authorities. The question is therefore not only whether AI works, but whether it can be **trusted** — and trust is established through evidence, validation, traceability and governance. The direction of travel is now explicit: in January 2026 the FDA and EMA jointly published their Guiding Principles of Good AI Practice in Drug Development, with ISO/IEC 42001 and the NIST AI Risk Management Framework reinforcing the same disciplines.[5] Crucially, accountability does not transfer: an AI-assisted submission remains the sponsor's submission.

Where value splits across the functions

The pattern is consistent across the functions that shape regulatory success. AI compresses the cost of analysis; expert judgment determines the highest-value decisions; and proprietary intelligence accumulates with every engagement.

EXHIBIT 3 | **Where AI assists vs. where judgment decides**

FUNCTION	WHERE AI ASSISTS	WHERE JUDGMENT DECIDES
Regulatory Affairs	Intelligence monitoring, precedent analysis, submission drafting and horizon scanning.	Filing strategy, interpreting agency feedback, when to challenge a regulator, balancing speed against approval risk.
Quality & Compliance	Inspection readiness, CAPA and deviation trending, complaint and audit monitoring.	Which observations matter most, which remediation will satisfy regulators, where inspectors will focus.
Pharmacovigilance	Case triage, duplicate detection, literature surveillance, multi-source signal integration.	Whether a signal is causal, changes the benefit–risk profile, or warrants regulatory or labeling action.
Medical Writing	Literature review, drafting, consistency and compliance checks, template population.	Which evidence matters, how it is framed, whether conclusions are supported and defensible.

Source: Eliquent Life Sciences analysis.

Validated infrastructure: the barrier to replication

New entrants can buy technology, but they cannot buy years of validation history, inspection-tested governance or accumulated trust. Every validated workflow produces documented evidence, reviewer confidence and reusable controls that compound over time. This is why validated AI infrastructure is more than a compliance requirement — it is the mechanism through which **trust, the most valuable asset in regulated life sciences, is created and made defensible**. It is where validated infrastructure turns effort into realized value.

From functional excellence to intelligence networks

The largest opportunity lies in connecting intelligence across the lifecycle. A quality observation may inform regulatory strategy; a safety signal may reveal manufacturing risk; a regulatory interaction may surface issues that later appear in inspection. AI changes the economics of integration, surfacing these relationships across domains at scale. The result is ***lifecycle intelligence*** — a trusted network that continuously learns and improves decision quality. The compounding value comes not from smarter models, but from a smarter organization.

4 Questions leaders should ask

When considering the use of AI in Regulatory, Quality and Safety, asking the right questions can unlock scalable value while limiting risk and technical debt. The wrong questions push organizations toward technology selection while overlooking the capabilities that create lasting value. The following diagnostic applies equally to internal initiatives and external partners. It separates an AI strategy that creates durable value from one that merely automates existing activity.

EXHIBIT 4 | **A diagnostic for the C-suite**

1	**Can every output be defended?** Traceable to source evidence, transparent in its reasoning, explainable to a regulator or auditor.
2	**Does the workflow survive a model change?** Model-agnostic architecture; you own the validation evidence and acceptance criteria, not the vendor.
3	**Where does expert judgment enter the process?** The points of interpretation, prioritization and accountability are designed in, not bolted on.
4	**What intelligence can competitors not replicate?** Knowledge from real engagements, inspections and interactions that compounds with every project.
5	**How is trust maintained over time?** Active governance of cybersecurity, drift, bias and hallucination — not a one-time validation event.
6	**Are you building intelligence or only automating?** Each engagement should make the next one better; expertise is captured, not lost when people leave.
7	**Where is value migrating — and are you ahead of it?** Talent and investment shift toward judgment and proprietary intelligence, away from commoditizing tasks.

Source: Eliquent Life Sciences.

The leadership challenge

The most important AI decisions facing life sciences organizations are no longer technology decisions — they are capability decisions. Leaders who focus exclusively on technology risk building temporary advantages that disappear with the next model release. Leaders who focus on judgment, intelligence and validated operating models build capabilities that compound in value and are far harder to replicate.

> ***AI is rapidly making intelligence more abundant. Judgment remains scarce.***

The most consequential decisions are rarely determined by information alone; they require interpretation, context, accountability and trust — attributes that cannot be downloaded or licensed and must be developed over time. The future of AI-enabled life sciences is not, finally, a technology story. It is a **judgment story** — and the firms that scale judgment, intelligence and trust responsibly will redefine leadership in the next generation of the industry. It is, above all, a story about where value is created — and value, not the model, is what endures.

WHAT THIS MEANS FOR LEADERS

- → Prioritize capabilities that survive model change.
- → Treat validation, cybersecurity and Responsible AI as strategic assets, not compliance exercises.
- → Capture expert judgment before it remains trapped as tacit knowledge.
- → Build intelligence networks across functions, not isolated automation use cases.
- → Create lasting value through trust, defensibility and decision quality — not by competing on model access.

THE ASSETS THAT LIE BEYOND THE MODEL

Expert judgment — guides decisions under uncertainty.

Proprietary intelligence — accumulated through real-world experience.

Trusted relationships — context no model can infer.

Validated infrastructure — turns expertise into scalable, defensible capability.

These are the assets that compound in value over time — and they lie beyond the model.

Sources & notes

Superscript numbers in the text correspond to the numbered sources below.

1. U.S. Food & Drug Administration, *"FDA Launches Agency-Wide AI Tool to Optimize Performance for the American People,"* press release, 2 June 2025. Elsa was built by Deloitte on Anthropic's Claude and deployed in FedRAMP-High AWS GovCloud. fda.gov/news-events/press-announcements/fda-launches-agency-wide-ai-tool-optimize-performance-american-people
2. K. Chew & M. Yang, *"FDA's Elsa AI Switches From Claude to Gemini: What Sponsors Need to Know,"* Clinical Leader, 12 March 2026 — the migration followed a 27 February 2026 federal directive to cease federal use of Anthropic's Claude. clinicalleader.com/doc/fda-s-elsa-ai-switches-from-claude-to-gemini-what-sponsors-need-to-know-0001
3. Stanford Institute for Human-Centered AI (HAI), *AI Index Report 2026* — the performance gap between leading open- and closed-source models has narrowed to roughly two points, while frontier per-token pricing has fallen by close to an order of magnitude year-over-year.
4. E. Brynjolfsson, B. Chandar & R. Chen, *"Canaries in the Coal Mine: Six Facts about the Recent Employment Effects of Artificial Intelligence,"* Stanford Digital Economy Lab, updated November 2025 — relative employment declines for early-career workers in highly AI-exposed roles, with senior cohorts insulated; evidence of a structural repricing of judgment.
5. FDA & EMA, *"Guiding Principles of Good AI Practice in Drug Development,"* 14 January 2026 — ten joint principles emphasising human oversight, risk management, data governance and lifecycle control. Reinforced by *ISO/IEC 42001:2023* (AI management systems) and the *NIST AI Risk Management Framework (AI RMF 1.0)*, January 2023. fda.gov/about-fda/artificial-intelligence-drug-development/guiding-principles-good-ai-practice-drug-development

ADDITIONAL REFERENCES

On the deployment-to-value gap and the role of domain integration: METR (Becker et al.), *"Measuring the Impact of Early-2025 AI on Experienced Open-Source Developer Productivity,"* July 2025; F. Dell'Acqua et al., *"Navigating the Jagged Technological Frontier,"* Harvard Business School / BCG, 2023; and enterprise-adoption analyses from BCG (2025), Bain & Company (2025) and Gartner. On the market value of the expert-annotation layer: Mercor Series C, October 2025 ($350M at a $10B valuation).