modak

The Hidden Tax on R&D: The Compounding Cost of Ungoverned Biomarker Data

Table of Contents

1 Foreword

Of more than 40,000 biomarkers currently under pre-clinical and clinical investigation globally, fewer than 400 have achieved regulatory qualification through FDA or EMA approval, drug label inclusion, or companion diagnostic designation. Published evidence from the NCBI places clinical adoption rates at one to two percent of all biomarkers appearing in the literature, a qualification ceiling that has remained structurally static despite significant advances in AI-powered discovery capability.

The constraint is not a scientific ambition. Biomarker discovery has accelerated materially: fifty percent of all research articles referencing new cancer biomarkers have been published within the last decade alone, a rate of output that reflects both the power of modern omics technologies and the proliferation of large-scale biobanks holding multi-modal data from millions of participants. The constraint is evidentiary infrastructure, the absence of scalable, governed mechanisms to translate the volume of discovered signals into the decision-grade evidence that qualification demands.

The organizational consequence is well-documented and consistently underpriced. When discovery processes are distributed and generative, but qualification processes remain selective, formal, and resource-intensive, the gap between the two does not stay static; it compounds. Program advance weak candidates based on informally assessed signals. Teams independently re-curate the same associations, each applying locally defined acceptance criteria. Interpretations diverge across functions. Review cycles extend. None of this registers as a discrete budget line, but each iteration consumes the expert bandwidth and program capital that should be directed toward advancing the strongest candidates.

Addressing this gap does not require additional discovery investment. The signals are already there. What is required is a disciplined framework for assessing evidentiary strength, preserving adjudication decisions, and making those decisions reusable across every program that encounters the same evidence.

> Discovery without governance is not a pipeline advantage. It is a compounding liability.

The Cost of Evidence Risk

<2% : Biomarkers achieve clinical adoption

Despite decades of discovery investment, clinical adoption by healthcare providers remains below 2% of published biomarkers.

NCBI/ Javaid et al., 2025

$2.6Bn : Average cost per approved drug

The fully capitalized cost of bringing a drug to market, is a figure directly compressed by evidence quality at every upstream decision point.

DiMasi et al., J.Health Economics, 2016

~90%: Drug candidate failure rate

Evidence quality at the target selection stage is a leading variable in clinical failure. Governance compresses that risk.

Mullard, A. Parsing clinical success rates. Nat Rev Drug Discov 15, 447 (2016).

02 The AI Paradox

AI Found the Signal. Governance Decides What to Do With It

Artificial intelligence has materially restructured the front end of biomarker research. Deep learning models demonstrate the ability to identify non-intuitive associations within high-dimensional datasets such as genomic, transcriptomic, proteomic, imaging, and electronic health record data, at a scale and speed that was not achievable through hypothesis-driven approaches. AI-powered frameworks have improved the precision of cancer screening, accelerated novel target identification, and enabled biomarker-informed patient stratification strategies that enhance the probability of trial success.

The published record substantiates these gains. Multi-omics frameworks such as MILTON leverage ensemble machine learning across genomics, proteomics, and EHR data to predict more than 3,000 diseases, outperforming conventional polygenic risk scores. AI-based models have been shown to identify misdiagnosed and undiagnosed cases, generating augmented case-control cohorts with greater statistical power to detect previously undetected genetic signals. Approximately 950 AI and machine learning-enabled medical devices have already passed FDA review, demonstrating that the regulatory pathway is viable.

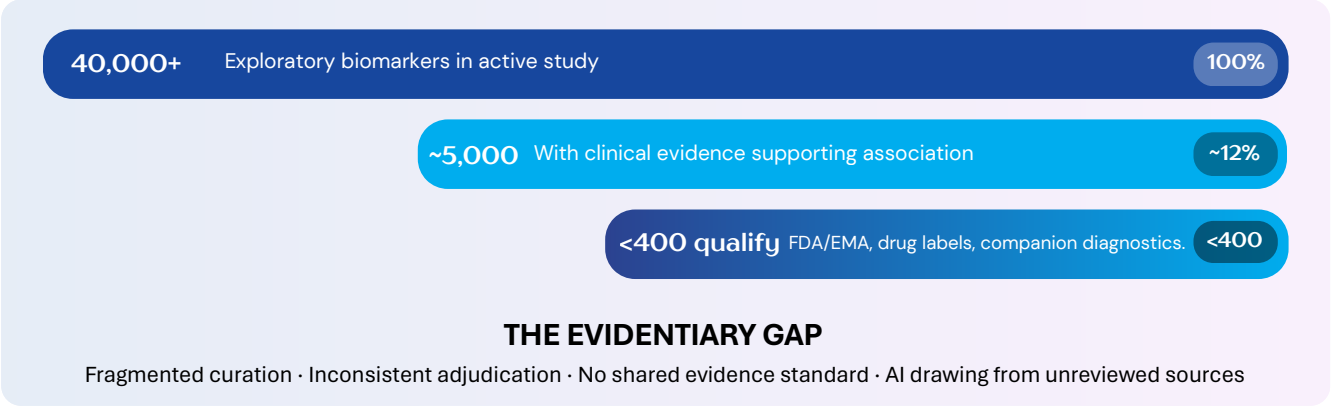

What AI Has Delivered

- **Fast, scalable interrogation of large scientific corpora.** AI-driven retrieval across publications, patents, clinical trial data, and internal reports now operates at a speed and coverage that manual literature review cannot approximate.
- **Association mapping across modalities.** Multi-modal frameworks integrating omics, imaging, and EHR data surface biomarker-disease associations that conventional single-modality approaches would not detect.
- **Conflict detection and candidate flagging.** AI can identify discordant signals, ambiguous ontology mappings, and divergences between internal findings and published literature, routing these to human reviewers.

What AI Cannot Resolve

- Deciding what the organization believes about a contested association.
- Assigning accepted, provisional, or disputed status with documented scientific rationale.
- Determining which downstream systems may consume sensitive or contested findings.
- Setting adjudication criteria for a specific program context of use.

These are organizational judgment decisions and they require governance infrastructure, not better models.

AI acceleration without governance compounds the problem rather than resolving it. Published research identifies four structural barriers to clinical translation of AI-powered biomarkers: siloed real-world evidence without harmonization, insufficient population diversity in training data, black-box interpretability gaps, and the absence of scalable infrastructure for evidence reuse. Each is a governance failure, not a model limitation.

When AI retrieval tools surface evidence from an ungoverned corpus, they cannot distinguish a reviewed association from a raw extraction. Downstream systems such as copilots, trial design tools, regulatory workflows, draw from the same pool indiscriminately. The result is not operational inefficiency alone but a structural risk to the scientific credibility of every decision that evidence informs.

03 The Governing Framework

The Three Criteria That Separate Signal from Qualification.

Rational prioritization of biomarker candidates across the discovery-to-qualification continuum rests on three evidentiary pillars, each of which must be explicitly structured and preserved to support downstream decision-making at the standard qualification requires.

Across these domains, three evidentiary dimensions consistently determine whether a biomarker progresses from exploratory signal to decision-grade asset. For prioritization to be scalable, these dimensions must be evaluated explicitly, structured systematically, and preserved for reuse rather than reconstructed at each program boundary.

1

Pillar 1: Demonstrated linkage to disease biology within a defined context of use

Mechanistic relevance within a specified context of use is the foundational criterion for any biomarker entering a qualification pathway. The FDA-NIH BEST framework classifies biomarkers across seven functional categories: susceptibility and risk, diagnostic, monitoring, prognostic, predictive, pharmacodynamic and response, and safety, each with distinct evidentiary requirements tied to the stage of the disease journey it is intended to address.
A biomarker assessed without explicit reference to its context of use cannot be meaningfully compared against others, cannot be weighed against regulatory qualification criteria, and cannot be prioritized for clinical investment with any rigor.

2

Pillar 2: Reproducible association with clinically meaningful outcomes

Demonstrated association with clinically meaningful endpoints: overall survival, progression-free survival, treatment response, immune infiltration, disease recurrence, constitutes the quantitative layer on which qualification decisions ultimately rest. The strength of this association must be assessable across multiple independent references, with explicit capture of positive, negative, and inconclusive signals across studies. A single-source finding, however promising, carries insufficient evidentiary weight for qualification investment. The evidence object must preserve the full distribution of associations and not collapse them into a single confidence score that obscures the underlying scientific debate.

3

Pillar 3: Analytical verification confirming fitness for decision-making

Measurement reliability, reproducibility across analytical platforms, and fitness for the decision context in which the biomarker will be applied complete the evidentiary triad. This pillar is particularly consequential for biomarkers entering AI-assisted workflows, where the quality of downstream outputs is directly bounded by the quality of the evidence on which those outputs are grounded. A biomarker association that is scientifically plausible but analytically unverified introduces downstream risk that compounds at every stage of program advancement.

These three pillars do not represent a new qualification standard. They reflect the criteria that regulatory bodies and clinical guideline committees have applied: inconsistently, manually, and at considerable cost for decades.

The strategic opportunity is to operationalize these criteria systematically, preserving the assessment as a reusable asset rather than repeating it independently at every program encounter with the same evidence.

04 Introducing the Governed Evidence Object

A Unit That Closes the Evidentiary Gap: Preserved and Reusable

A governed evidence object is not an analytical output. It is a decision-grade asset, a structured, versioned, provenance-tracked artefact that carries the organization's adjudication decision and makes it reusable across every program that subsequently encounters the same evidence. The distinction is consequential.

A document answer is session-bound. An AI extraction result reflects model confidence. A governed evidence object reflects organizational confidence, and it travels with the reasoning that produced it.

At enterprise scale, a properly constituted evidence object must preserve the full complexity of scientific reality. Ambiguity is not a defect to be resolved; it is a condition to be documented.

Contested associations, provisional interpretations, and evolving regulatory positions must be carried explicitly and not collapsed into a false binary of accepted or rejected, so that downstream systems and human reviewers can apply appropriate epistemic weight to every finding they encounter.

The Six Attributes of an Enterprise-Ready Evidence Object

Identity and mapping context.

Preferred terms, accepted and alternate ontology mappings, and therapeutic context, ensuring consistent interpretation across programs while preserving legacy or disputed representations. Alternate mappings are retained, not discarded, because ontology frameworks evolve and traceability of prior mappings carries regulatory value.

Conflicting and convergent evidence signals.

Explicit capture of positive, negative, and inconclusive findings across internal studies, published literature, and IP sources, rather than collapsing disagreement into a single confidence score. This is the attribute that makes the evidence object auditable rather than merely usable.

Validity status and qualification intent.

Clear designation of whether the evidence is provisional, accepted with caveats, disputed, or excluded, explicitly tied to a defined context of use. The BEST framework's seven biomarker categories provide the regulatory vocabulary within which this designation is made.

Provenance and decision lineage.

Traceable links back to source passages, internal analyses, reviewer notes, and adjudication rationale, enabling regulatory defensibility and allowing any downstream user to reconstruct the reasoning behind any prior decision without consulting the original reviewer.

Publication and consumption controls.

Defined rules governing where and how the evidence may be used such as, analyst workbenches, trial design workflows, AI copilots, regulatory submissions, preventing inappropriate downstream automation of findings that require human validation before use.

Re-evaluation triggers.

Explicit conditions under which evidence must be revisited: new trial readouts, updated regulatory guidance, ontology changes, or the passage of defined time thresholds. These triggers ensure that confidence does not silently degrade over time as the evidentiary landscape evolves.

Illustrative Biomarker Evidence Object

A reusable evidence asset that preserves ambiguity, provenance, review state, and publication controls.

KRAS G12C | **Discovery workflow** | **Disputed / accepted with caveat** | **Not approved for copilot answers**

1. Identity and Mapping

Field	Value
Preferred term	KRAS G12C
Entity class	Biomarker / mutation
Accepted mapping	HGVS p.G12C (workflow default)
Alternate mapping	Legacy local alias retained
Therapeutic context	NSCLC discovery evidence assembly

2. Conflicting Evidence Signals

Publication Positive association in assay A	Signal observed in one NSCLC cohort, but effect strength is assay-specific and not replicated across all subgroups.
Internal Study No reproducible signal in assay B	Translational team reports inconsistent internal readout; reviewer flags method mismatch before reuse.
Patent Family Broad claim, low practical specificity	Signal observed in one NSCLC cohort, but effect strength is assay-specific and not replicated across all subgroups.

3. Publication Controls

Field	Value
Reviewer state	Accepted with caveat; alternate mapping retained
Approved downstream use	Research + analyst workbench only
Excluded from	Copilot answer layer until re-review
Re-review trigger	New phase II readout or ontology update

4. Provenance and Lineage

Field	Value
Source passage	Results section, paragraph 3, sentence 2
Conflicting internal note	Assay B appendix, reviewer comment retained
Access policy	Discovery-TA-01
Lineage run ID	ELA-2026-04-04-01
Current version	1.3
Last review	2026-04-04

05 Proof from Practice

An Operating Model Validated in Pharma Evidence

The governance discipline underlying a structured evidence layer is not hypothetical. It has been tested and validated in the most demanding evidence environment in life sciences such as regulated clinical trial data onboarding, where the evidentiary burden is formally defined, the adjudication is legally accountable, and the downstream consequences of an ungoverned data release include regulatory rejection of a submission.

In a documented engagement with a Fortune 500 pharmaceutical organization managing thousands of legacy clinical trial datasets alongside newer studies standardized to SDTM, Modak addressed a structurally identical problem: multiple source systems each describing the same scientific object through different identifiers, terminologies, and completeness levels, with transformation logic resident in institutional memory rather than in a reusable, inspectable contract.

5 Operating Components — and Their Biomarker Evidence Equivalents

	Operating Component	What Modak Did in Clinical Data	What This Means for Biomarker Evidence
1	Cross-system source profiling and intake.	SAS binary datasets and metadata ingested across legacy and modern study formats, with source inventory established before any mapping work began.	Inventorying publications, internal translational studies, and AI-surfaced findings for a given indication before adjudication begins.
2	ML-guided mapping with human approval gates	Guided mapping proposed alignments and flagged ambiguous cases for human review. The system routed decision-ready candidates to named owners; it did not make acceptance decisions unilaterally.	AI surfacing ontology candidates and routing contested mappings to the domain steward for decision.
3	Exception adjudication with retained evidence	Each contested mapping was documented, adjudicated by a named reviewer, and recorded with rationale preserved in the release evidence package.	The adjudication decision travels with the evidence object rather than residing in an email thread or institutional memory.
4	Automated provenance and documentation generation	Documentation was generated as a by-product of the governed workflow rather than assembled retrospectively.	The audit trail enabling a regulatory reviewer to trace any AI output to a specific reviewed evidence object with reviewer attribution and version history.
5	Governed downstream integration without remapping	Standardized outputs delivered into existing analytics infrastructure. Downstream consumers drew from the governed layer rather than re-extracting from source.	AI copilots, trial design dashboards, and submission workflows consuming adjudicated evidence rather than raw extractions.

The measured outcomes of this engagement establish a practical performance baseline for the governance model. These are illustrative target-setting benchmarks derived from Modak methodology

Which is applied in a regulated pharma context — they represent what the operating model is designed to achieve, not guaranteed outcomes.

25-40% reduction

Lead time · source receipt to approved release

The primary driver is elimination of rework: reviews begin from a validated baseline rather than from raw source material.

20-35% reduction

Manual mapping and reconciliation effort per release cycle.

Automation handles high-confidence alignments; human review is concentrated on genuinely contested cases.

40-60% reduction

Evidence-pack assembly time per release.

Documentation generated as a workflow by-product rather than produced manually after the fact.

0

Critical unresolved conformance exceptions at release.

Every contested case is resolved and documented before the evidence package exits the governed layer.

95% or above

Required lineage and metadata fields populated at release

the completeness threshold that supports downstream audit and regulatory defensibility.

06 The ROI Case and Operational Impact

What Governance Pays Back and What It Changes.

Evidence governance is a productivity investment. The return compounds across every Program in the portfolio. It materializes in three measurable dimensions: the elimination of redundant re-curation, the improvement of downstream AI output reliability, and the reduction of regulatory audit exposure.

Each of these returns is quantifiable from the first quarter of deployment using the KPI framework in the following section.

25 - 30% Review cycle time recovered within 12 months	Achieved through the elimination of evidence packet reconstruction at the start of each review cycle. Analysts begin from a governed baseline rather than raw literature.
3x - 5x Return on evidence curation investment	A single governed evidence object consumed by three to five downstream workflows recovers its curation cost three to five times over. Without governance, that cost is paid again at every workflow encounter with the same evidence.
40 - 60% Reduction in evidence-related AI rework	AI Programs grounded in governed evidence objects produce consistent, auditable outputs. Those drawing from ungoverned corpora require rework at every downstream stage. Governance removes the root cause.

What Changes When Governance Works

The operational impact of a governed evidence layer is most precisely understood through the workflow transformations. It produces across the five most consequential failure modes in ungoverned evidence environments.

Workflow	Without Governance	With Governed Evidence Layer
Review cycle time	Analysts reconstruct evidence packets from raw literature at the start of each review cycle. The same associations are curated independently across multiple teams simultaneously. Review cycles run four to six weeks.	Reviews begin from a governed, validated baseline. Cycle time reduces to three to four weeks within six months. A 25 to 30% time recovery is sustained by end of Year 1, compounding across every Program in the portfolio.
Cross-functional consistency	Acceptance criteria are implicit and locally applied. The same biomarker association produces different conclusions across functions. Interpretive misalignment surfaces late in Program governance, at the stage where correction is most expensive.	Organizational consensus is encoded in the evidence object and propagated across all functions. Interpretive drift is made visible and traceable before it escalates into a Program decision conflict.
AI tool reliability	AI tools surface reviewed and unreviewed material without distinction. The same scientific query returns different answers across sessions and tools. Outputs cannot be validated against a shared evidence standard, undermining confidence in AI-assisted decision-making.	AI tools are required to draw exclusively from the governed evidence layer. Output consistency reaches 85% or above on controlled query sets within six months. AI investment performance becomes measurable and reportable.
Regulatory submission	Evidence used in submissions lacks traceable provenance or documented review history. Regulatory queries require expensive retrospective reconstruction. Audit gaps create submission risk that is difficult and costly to remediate after the fact.	Every evidence object carries a complete audit trail: source documents, review status, reviewer attribution, version history. Regulatory queries are answered directly from the governed layer. Retrospective reconstruction is eliminated.
Expert bandwidth utilization	Domain experts spend significant time answering questions that a governed evidence system would answer directly. Strategic scientific work is deferred to manage operational queries.	Routine evidence questions are resolved from the governed layer. Expert bandwidth is redirected to contested adjudications, program decisions, and scientific leadership activity.

07 KPI Tracking Framework

What to Measure, What Good Looks Like, and What It Unlocks

Governance without measurement is policy without accountability. The KPI framework below is structured for quarterly board and program governance reporting. Each metric includes a Year 0 baseline that can be established before any tooling is deployed.

A Year 1 target, and the specific business value that achieving it produces. These metrics are the instrument by which a governance Program becomes a measurable strategic investment rather than an operational cost center.

Program Performance KPIs

KPI	Baseline (Year 0)	Target (Year 1)	Business Value Unlocked
Review cycle time	4 to 6 weeks per target review cycle	3 to 4 weeks — a 25 to 30% reduction within 12 months	Each week recovered per cycle across a portfolio of programs returns significant expert time and accelerates decision velocity at the stage where speed carries the highest program value.
Re-curation rate	60 to 80% of review cycles require full rebuild from raw literature	Below 20% within 12 months of governance layer deployment	Each percentage point reduction recovers expert analyst hours per quarter. At the 20% target, the majority of redundant curation cost across the portfolio is eliminated and reallocated to program-advancing activity.
Evidence reuse rate	1 to 1.5 downstream uses per evidence object before re-curation is required	3 or more distinct downstream uses per object by end of Year 1	Reuse rate above 3 means governance infrastructure cost is recovered three times per object. This is the primary ROI multiplier for board-level reporting and the metric that most directly demonstrates the compounding return on governance investment.

AI Investment Performance KPIs

KPI	Baseline (Year 0)	Target (Year 1)	Business Value Unlocked
Governed evidence grounding rate	Unknown or below 40% for most deployed tools	100% for regulatory and AI tools. 90%+ for discovery and translational tools	Determines whether AI programs produce reliable, auditable outputs. Below 90%, results cannot be validated against a shared scientific standard; every downstream decision carry unquantified risk.
AI output consistency score	Below 60% consistency on controlled query set across equivalent queries	85%+ on controlled query set within 6 months	The observable measure of AI reliability for program leads and regulators. At 85%, the program can be presented to sponsors and regulatory stakeholders producing reproducible scientific outputs.

Regulatory and Audit Readiness KPIs

KPI	Baseline (Year 0)	Target (Year 1)	Business Value Unlocked
Audit trail completeness	Below 30% of active evidence objects have a complete, exportable audit trail	100% for all published objects in the governed layer by end of Year 1	Audit trail completeness is the baseline condition for regulatory defensibility. Any evidence object without a complete trail represents a liability in a regulatory review, a due diligence process, or an internal governance audit, a liability that is substantially less costly to prevent than to remediate.
Evidence currency rate	Below 50% currency. Most associations have no defined re-review policy	At least 95% of associations current within defined thresholds by end of Year 1	Currency above 95% allows the organization to demonstrate to regulators that its evidence base reflects the current state of scientific knowledge, not the state at initial curation, which may have preceded significant trial readouts, ontology updates, or regulatory guidance revisions.

The ROI Timeline

What becomes measurable and when, from first deployment

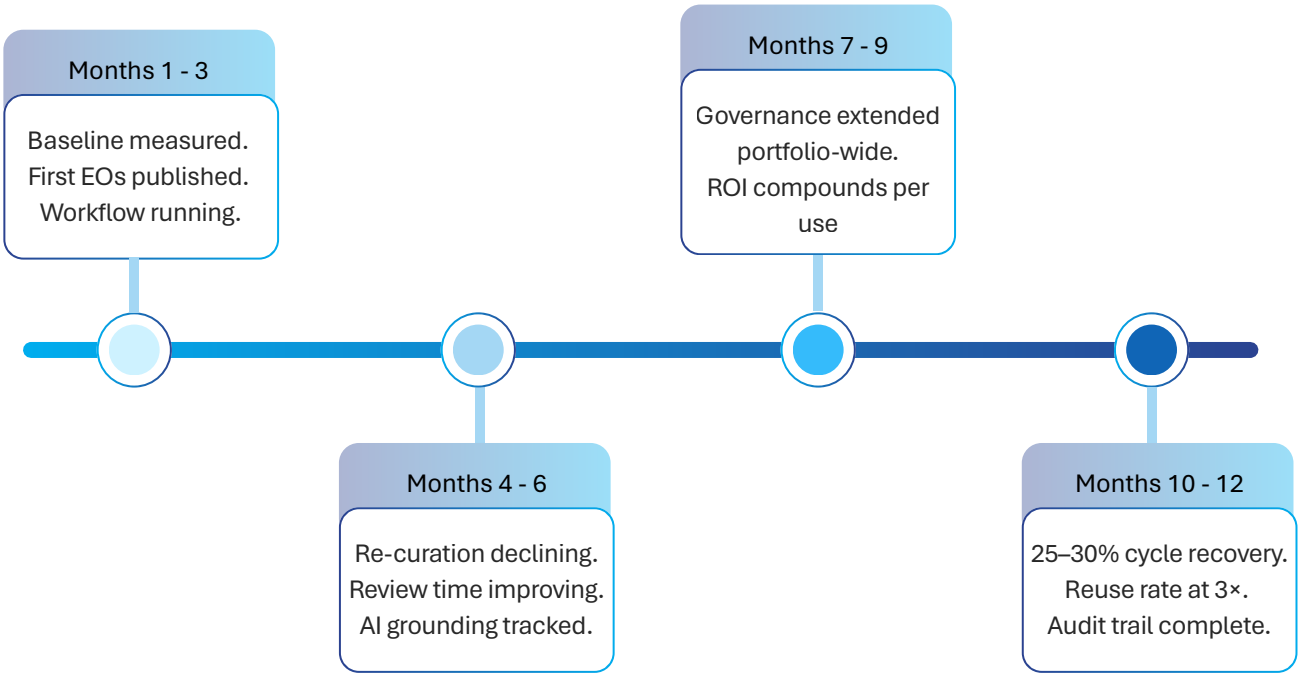

08 Leadership Decisions

Five Decisions That Determine Whether the Investment Delivers.

Evidence governance Programs underperform for a consistent and well-documented reason: the decisions that only senior leadership can make are delegated to the program team. The five decisions below are not technical choices.

They are strategic commitments, each of which is a precondition for the ROI targets in this paper to be achievable. Working through them in order produces the organizational design, measurement infrastructure, and platform clarity required to build a durable governed evidence capability.

01 Name the owner of evidence quality before the platform is selected

The most common and most expensive failure pattern in evidence governance is delegating scientific ownership to the technology team. Technology teams can build the infrastructure. They cannot decide what constitutes accepted evidence for a specific program context of use. A named Domain Steward with Program authority must be identified and empowered before any platform conversation begins. If that individual cannot be named, the Program is not ready to proceed.

02 Establish the KPI baseline within 30 days before any tooling is deployed

Re-curation rate, review cycle time, and AI grounding rate must be measured from current operations before governance infrastructure is deployed. Without a baseline, the Program cannot demonstrate ROI to the board, cannot identify where to direct initial investment, and cannot hold the governance commitment accountable to a measurable standard. The baseline measurement requires no new tooling and can be completed within 30 days using existing operational data.

03 Make governed evidence grounding a condition of AI Program funding

Every AI Program in the portfolio should be required to demonstrate at least 90% grounding in governed evidence objects before production deployment approval is granted. This decision converts the governance infrastructure from a background data Program into a hard dependency for every AI investment the organization makes. It is the single commitment most likely to ensure the governance layer is adequately resourced and sustained past the initial deployment phase.

04 Start with one indication and scale on demonstrated evidence of impact

Attempting to govern all evidence across all Programs simultaneously produces a Program that is too large to validate, too slow to demonstrate value, and too diffuse to hold accountable. Begin with the indication area where review cycle inefficiency is highest and re-curation burden is most visible. Establish the governance model, prove the KPI improvement against a measured baseline, and use that documented proof to resource and justify expansion. Governance designed to cover everything immediately covers nothing well.

05 Report governance KPIs at board level from the first quarter of deployment

Review cycle time, re-curation rate, and governed AI grounding rate belong in the standard board reporting pack from quarter one. Evidence governance failures manifest as Program failures: delayed reviews, inconsistent AI outputs, regulatory queries, submission gaps. Reporting governance KPIs at board level makes the connection between governance investment and Program performance visible, ensures the investment is resourced as a Program dependency rather than a discretionary back-office activity, and creates the accountability structure that sustains the return over the Program lifecycle.

09 Conclusion

The Evidence Is There. The Decision Is Whether to Govern It.

The capability constraint in life sciences R&D is no longer the discovery of biomarker signals. AI has addressed that dimension with demonstrable effectiveness. The constraint is the organization's ability to translate the signals already available into decision-grade evidence that can be assessed consistently, reused across programs, and presented to regulators with full provenance. That constraint is a governance problem, and it is solvable.

A governed evidence layer recovers 25 to 30% of review cycle time within 12 months, multiplies the return on every evidence curation investment by three to five times, and makes every AI Program in the portfolio more reliable and more defensible against regulatory scrutiny. It does not require a new AI model, a platform replacement, or a discovery investment.

It requires one architectural decision such as inserting a governed publication step between raw scientific content and the downstream systems that act on it, and five leadership commitments that only senior leadership can make.

The KPI framework in this paper provides the measurement structure to demonstrate return at the board level from the first quarter of deployment. The proof from regulated clinical data environments confirms that the operating model delivers at an enterprise scale. The evidence is already there. The decision is whether to govern it.

10 References

1. Javaid, H., Petrescu, C.C., Schmunk, L.J., et al. (2025). “The impact of artificial intelligence on biomarker discovery.” Open Biology / Portland Press. PMCID: PMC12802349. https://pmc.ncbi.nlm.nih.gov/articles/PMC12802349/
2. Alum, E.U. (2025). “AI-driven biomarker discovery: enhancing precision in cancer diagnosis and prognosis.” Springer Nature / NCBI. PMCID: PMC11906928. https://pmc.ncbi.nlm.nih.gov/articles/PMC11906928/
3. FDA-NIH Biomarker Working Group. (2016). BEST (Biomarkers, EndpointS, and other Tools) Resource. National Institutes of Health / NCBI Bookshelf. https://www.ncbi.nlm.nih.gov/books/NBK326791/
4. DiMasi, J.A., Grabowski, H.G., and Hansen, R.W. (2016). “Innovation in the pharmaceutical industry: New estimates of R&D costs.” Journal of Health Economics, 47, 20–33. https://doi.org/10.1016/j.jhealeco.2016.01.012
5. Mullard, A. (2016). “Parsing clinical success rates.” Nature Reviews Drug Discovery, 15, 447. https://doi.org/10.1038/nrd.2016.136
6. McKinsey & Company. (January 2025). Scaling gen AI in the life sciences industry. McKinsey Global Institute. https://www.mckinsey.com/industries/life-sciences/our-insights/scaling-gen-ai-in-the-life-sciences-industry
7. Garg, M., Karpinski, M., Matelska, D., et al. (2024). “Disease prediction with multi-omics and biomarkers empowers case–control genetic discoveries in the UK Biobank.” Nature Genetics, 56, 1821–1831. https://doi.org/10.1038/s41588-024-01898-1
8. U.S. Food and Drug Administration. (2024). Artificial Intelligence and Machine Learning (AI/ML)-Enabled Medical Devices. FDA Center for Devices and Radiological Health. https://www.fda.gov/medical-devices/software-medical-device-samd/artificial-intelligence-and-machine-learning-aiml-enabled-medical-devices

9. Barioni, C.T.S., de Barros Wandresen, R.P., Pereira, L.F., et al. (2024). "Artificial intelligence for the identification of biomarkers in cancer prevention and diagnosis: advances and perspectives." Revista Brasileira de Cancerol. https://doi.org/10.2635/2176-9745.RBC.2024v70n2.4692
10. Passaro, A., Al Bakir, M., Hamilton, E.G., et al. (2024). "Cancer biomarkers—emerging trends and clinical implications for personalized treatment." Cell, 187, 1617–1635. https://doi.org/10.1016/j.cell.2024.02.04
11. Excelra. (2023). Bridging the gap between biomarker discovery and qualification through a validity index and quantified clinical outcome associations. Excelra Biomarker Insights White Paper. https://www.excelra.com/biomarker-insights/
12. Wagner, J.A. (2008). "Strategic approach to fit-for-purpose biomarkers in drug development." Annual Review of Pharmacology and Toxicology, 48, 631–651. https://doi.org/10.1146/annurev.pharmtox.48.113006.094611
13. Altar, C., Amakye, D., Bhathena, A., et al. (2008). "A Prototypical Process for Creating Evidentiary Standards for Biomarkers and Diagnostics." Clinical Pharmacology & Therapeutics, 83(2), 368–371. https://doi.org/10.1038/sj.clpt.6100451
14. Gromova, M., Vaggelas, A., Dallmann, G., and Seimetz, D. (2020). "Biomarkers: Opportunities and Challenges for Drug Development in the Current Regulatory Landscape." Biomarker Insights, 15, 1177271920974652. https://doi.org/10.1177/1177271920974652
15. U.S. Food and Drug Administration. (2023). Guidance for Industry: Electronic Systems, Electronic Records, and Electronic Signatures in Clinical Investigations. FDA Center for Drug Evaluation and Research. https://www.fda.gov/regulatory-information/search-fda-guidance-documents/electronic-systems-electronic-records-and-electronic-signatures-clinical-investigations-questions
16. U.S. Food and Drug Administration. (2023). List of Qualified Biomarkers. FDA Biomarker Qualification Program. https://www.fda.gov/drugs/biomarker-qualification-program/list-qualified-biomarkers
17. European Medicines Agency. (2023). Opinions and letters of support on the qualification of novel methodologies for medicine development. EMA Human Regulatory Research and Development. https://www.ema.europa.eu/en/human-regulatory/research-development/scientific-advice-protocol-assistance/novel-methodologies-biomarkers/opinions-letters-support-qualification-novel-methodologies-medicine-development
18. Faculty of Medicine, Imperial College London. (2024). Biomarker Assessment, Validation, and Translation. Imperial College London Department of Surgery and Cancer. https://www.imperial.ac.uk/department-surgery-cancer/research/surgery/groups/biomarker-assessment-validation-and-translation/
19. Deloitte. (2022). Reforming diagnostics: Turning challenges into enablers. Deloitte UK Life Sciences and Health Care. https://www2.deloitte.com/content/dam/Deloitte/uk/Documents/life-sciences-health-care/deloitte-uk-future-diagnostics-companion-report-2022.pdf
20. Wishart, D.S., et al. (2021). "MarkerDB: an online database of molecular biomarkers." Nucleic Acids Research, 49(D1), D1259–D1267. https://doi.org/10.1093/nar/gkaa1067
21. Aldoseri, A., Al-Khalifa, K.N., and Hamouda, A.M. (2023). "Re-thinking data strategy and integration for artificial intelligence: concepts, opportunities, and challenges." Applied Sciences, 13(12), 7082. https://doi.org/10.3390/app13127082
22. Modak. (2024). SciBite Partnership: Biomarker Intelligence and Governed Vocabulary Pipelines. Modak Public Partner Page. https://modak.com/partners/scibite/

Connect with us

Laura P Williams
Chief AI & Data Strategy Lead
Laura.williams@modak.com

Contributors

Aarti Joshi
Murthy Gunturi

About Modak

Modak is a data engineering company specializing in life sciences and biopharma. We help pharmaceutical and biotech organizations build governed data infrastructure that translates complex, fragmented scientific evidence into decision-ready assets, across clinical trials, translational research, and regulatory workflows. Our work spans regulated data onboarding, ML-guided harmonization, provenance-tracked evidence pipelines, and cloud-agnostic data lake architecture. We work with Fortune 500 Healthcare & Life Sciences organizations to deliver the governed evidence foundations that AI programs, trial design workflows, and regulatory submissions depend on.