



Research Article

AI-Governed Knowledge Infrastructure for Predictive Compliance and Supply Resilience in Class II Medical Devices: A Case Study

Kenneth XU¹, Nick VYAS² and Markus MAU³

¹Doctorate Pan-European Studies, University of Pula, Pula, Croatia

²Marshall School of Business, University of Southern California, Los Angeles, United States,

³Doctorate Pan-European Studies, Eisenstadt, Austria

Correspondence should be addressed to: Kenneth XU; kxu124@gmail.com

Received date: 10 June 2026; Accepted date: 27 July 2026; Published date: 15 September 2026

Copyright © 2026. Kenneth XU, Nick VYAS and Markus MAU. Distributed under Creative Commons Attribution 4.0 International CC-BY 4.0

Abstract

Small and mid-sized Class II medical device manufacturers face a scale-up challenge when regulatory, quality, and supply chain information grows faster than the organization's ability to manage it. This paper examines whether governed artificial intelligence can convert fragmented records into decision-ready knowledge without weakening compliance, traceability, or human accountability. Using an applied case-study design, the research evaluates a six-month pilot at New Horizon Biotech, an anonymized Class II diabetes device manufacturer. A five-FTE AI Business Unit connected regulatory evidence, supplier records, quality documentation, and operational signals through a dual-loop model. The Regulatory Readiness Loop supported evidence retrieval, predicate comparison, submission drafting, and EU MDR GSPR mapping. The Supply Adaptation Loop supported supplier early warning, alternate-source readiness, change impact review, and landed-cost analysis. Pilot results showed a 32.3% reduction in drafting cycle time, 67.8% reduction in defect density, 88.0% reduction in evidence retrieval time, 38.8% faster lead-time recovery, zero critical stockouts, a Month 6 groundedness score of 0.94, and an 82.4% six-month ROI. The findings suggest that governed AI can support regulatory throughput and supply resilience when it is implemented as controlled knowledge infrastructure rather than autonomous decision-making.

Keywords: Governed AI; knowledge management; supply chain resilience; Class II medical devices

Introduction

Class II medical device manufacturers operate in a knowledge-intensive environment. Product design records, supplier qualifications, risk files, labeling materials, test reports, quality records, and regulatory submissions all help demonstrate that a device is safe, effective, traceable, and market ready. These records are not only administrative artifacts, they are operating evidence.

The information burden is especially visible in diabetes devices, including continuous glucose monitors, insulin pumps, and blood glucose monitoring systems. These products are not as clinically complex as many Class III implants, but they still require design controls, supplier oversight, performance evidence, labeling control, and post-market documentation. In the United States, many Class II devices follow the FDA 510(k) pathway. In the European Union, Class IIa and IIb devices require structured technical documentation under the EU Medical Device Regulation.

Large manufacturers can often respond to this burden by adding specialized teams, systems, consultants, and regional functions. Small and mid-sized manufacturers have less room to do so. As they expand across markets, regulatory, quality, and supply chain complexity may grow faster than staffing capacity. This paper defines that condition as the Scale-Up Paradox.

The Scale-Up Paradox is also a knowledge architecture problem. A firm may already have the needed evidence, but the evidence may sit across disconnected repositories. Regulatory Affairs may own submission histories. Quality may own CAPA and audit records. Supply Chain may own supplier risk data. Engineering may own the Design History File and Device Master Record. Each function may manage its own documents well, but the firm may still lack an integrated view of operational truth.

Artificial intelligence may help address this problem, but only under strong governance. In regulated manufacturing, AI cannot be treated as an ordinary productivity tool. A fluent but unsupported regulatory statement can create audit exposure. A supplier-risk score based on inconsistent records can mislead decision makers. A model that drafts without source

control can weaken the quality system it was meant to support.

This paper examines whether governed AI can serve as business information infrastructure for small and mid-sized Class II medical device manufacturers. The study uses a six-month pilot at New Horizon Biotech, an anonymized case organization. The pilot tested whether a five-FTE AI Business Unit could improve regulatory throughput and supply resilience while preserving source traceability, quality-system discipline, and human approval authority.

The paper contributes a dual-loop AI framework, pilot evidence across regulatory and supply metrics, and a practical implementation model for small and mid-sized regulated manufacturers.

Literature and Conceptual Background

Regulatory Knowledge in Class II Medical Device Operations

Medical device regulation creates a demanding knowledge environment. The FDA 510(k) pathway requires manufacturers to demonstrate substantial equivalence to a legally marketed predicate device. This makes intended use, technological characteristics, performance evidence, risk controls, and labeling central to submission quality. EU MDR technical documentation requires a different but related evidence logic, with emphasis on lifecycle documentation, clinical evaluation, post-market surveillance, and General Safety and Performance Requirements.

For firms operating across both markets, one internal evidence base may need to support multiple regulatory uses. A test report may support a 510(k) section, an EU MDR GSPR clause, and a change-control decision. A supplier qualification file may support purchasing controls, but it may also affect technical documentation if the supplier provides a critical material. This means regulatory compliance should not be treated only as a final submission task. It should be embedded into daily information management.

Knowledge Management, Information Quality, and AI

Knowledge management research emphasizes that organizational value depends on creating, storing, retrieving, sharing, and applying knowledge. In regulated manufacturing, this distinction matters. A company may possess the right supplier certificate, verification report, clinical summary, or labeling record, but still fails to use it effectively if the document cannot be found, verified, linked, and applied at the right decision point.

Information quality is therefore central to AI adoption. If supplier names are duplicated, AI may misread a sole-source risk as a diversified supply base. If outdated records remain active, AI may retrieve obsolete evidence. If metadata are missing, AI may generate a plausible but unsupported regulatory statement. For this reason, AI should be viewed as governed knowledge infrastructure, not uncontrolled automation. Its value lies in improving search, linkage, prioritization, drafting support, and decision readiness.

Supply Chain Resilience and Regulated Operations

Supply chain resilience depends on visibility, preparedness, response speed, and recovery capacity. In medical device operations, resilience cannot be separated from regulatory readiness. A supplier change may require new qualification evidence, verification testing, labeling review, or regulatory impact assessment. A missing supplier record may block both shipment and submission readiness. A regulatory delay can also freeze launch inventory and production planning.

This creates a practical gap. Knowledge management explains how firms create and apply knowledge. Supply chain resilience explains how firms absorb disruption. Regulatory research explains documentation and quality obligations. AI governance explains accountability and control. Few studies show how these streams can be connected inside a

small, regulated manufacturer. This paper addresses that gap through a dual-loop operating model.

AI Governance in Regulated Environments

AI adoption in regulated industries carries specific risks. Large language models can hallucinate unsupported claims. Users may over-trust confident outputs. Models can drift as underlying systems change. These risks are not acceptable in medical device operations, where audit trails and source traceability are essential.

The governance position used in this paper is cautious. AI may assist regulated knowledge work, but it must be bounded by controlled data, human review, source citation, approval records, and documented oversight. The focus is not AI embedded in a medical device. The focus is enterprise AI used inside the business operating system to support evidence retrieval, supplier sensing, drafting, mapping, and decision preparation.

Research Objectives and Methodology

This study uses an applied case-study design with baseline comparison. The method is appropriate because the research examines a contemporary operating model in its organizational setting. The intervention combines technology, quality systems, regulatory workflows, supplier decisions, human review, and financial analysis.

New Horizon Biotech is presented as an anonymized case organization to protect company-specific operating data while preserving the structure of the regulatory, quality, and supply chain workflows studied. The firm is suitable because it faces the Scale-Up Paradox. It has global regulatory exposure, multiple product lines, distributed operations, and limited headcount. It is large enough to face meaningful complexity, but small enough that the effect of a focused AI operating model can be observed.

Table 1. Research objectives, operational questions, and evidence sources

Objective	Operational question	Primary measures	Evidence source
Objective 1	Can regulatory and supply functions be connected through one AI-enabled operating model?	R-Loop, S-Loop, decision gates, human review controls	Framework design and pilot operating model
Objective 2	Can AI-assisted regulatory workflows reduce cycle time while maintaining quality?	Drafting cycle time, defect density, review iterations, evidence retrieval time	Baseline vs. pilot comparison
Objective 3	Can predictive supplier scoring improve supply resilience?	Lead-time recovery, sole-source exposure, stockouts, alternate-source readiness	Supplier monitoring and SRI results
Objective 4	Can the framework remain auditable and financially viable?	Groundedness, red-team testing, ROI, landed-cost simulation	Governance logs and financial analysis

The pilot focused on Western regulatory and supply chain operations, especially FDA 510(k) and EU MDR workflows. China NMPA and Australia TGA activities were excluded to reduce language, data-localization, and third-party market access confounding factors. The product scope covered an insulin pump, a continuous glucose monitoring system, and a blood glucose monitoring system.

The intervention was executed by a five-FTE AI Business Unit. The team included product ownership, regulatory systems, data operations, AI engineering, and supply chain analytics roles.

It did not replace functional owners. It supported Regulatory Affairs, Quality, Supply Chain, and Engineering by improving retrieval, risk sensing, and decision preparation.

Before the pilot, the team completed an eight-week data preparation phase. This included supplier-name standardization, duplicate vendor cleanup, regulatory document organization, and controlled repository preparation for retrieval-augmented generation. This step was essential because AI performance depended on information quality.

Table 2. Data collection design and measurement logic

Data category	Unit of observation	Example fields	Frequency	Purpose
Regulatory throughput	Submission or document package	Cycle time, errors per page, review cycles, retrieval hours	Event-based	Measures R-Loop speed and quality
Supply resilience	Supplier-week	Delay days, SEWS score, dual-sourcing status, stockouts	Weekly	Measures S-Loop response
Governance and auditability	AI output or review event	Groundedness, citation verification, red-team result, approval status	Monthly and event-based	Measures AI trust and control

Data category	Unit of observation	Example fields	Frequency	Purpose
Financial performance	Pilot cost and value stream	Labor efficiency, inventory optimization, launch acceleration, net ROI	End of pilot	Measures financial viability

The study compared baseline manual workflows against AI-assisted pilot workflows. Regulatory measures included drafting cycle time, evidence retrieval time, defect density, review iterations, and first-pass GSPR completeness. Supply measures included lead-time recovery, sole-source exposure, stockouts, and supplier early warning signals. Governance measures included groundedness, citation checks, red-team testing, and human approvals. Financial measures included pilot cost, gross value generated, net value, and six-month ROI.

The study used three threshold-based pilot propositions: at least 25% drafting cycle-time reduction without higher defect density,

improved lead-time recovery with no monitored critical stockouts, and positive six-month ROI. Because the data are operational performance metrics rather than randomized observations, the study does not claim population-level statistical inference.

Dual-Loop AI Framework

The proposed framework addresses knowledge fragmentation by connecting regulatory evidence, supplier records, quality documentation, and operational signals through two linked loops: the Regulatory Readiness Loop and the Supply Adaptation Loop.




Figure 1. Dual-Loop AI-Enabled Knowledge Management Framework

Note: The framework connects regulatory readiness and supply adaptation through a shared AI knowledge layer. AI supports retrieval, mapping, scoring, and decision preparation, while human experts retain approval authority.

Regulatory Readiness Loop

The Regulatory Readiness Loop, or R-Loop, focuses on regulatory evidence knowledge. It supports controlled evidence retrieval, predicate comparison, technical documentation preparation, GSPR mapping, and submission quality review. Its purpose is not to allow AI to

write submissions independently. Its purpose is to help regulatory experts locate, verify, reuse, and structure approved evidence more efficiently.

The R-Loop includes two main gates. Gate R1 evaluates predicate suitability and claims coherence for FDA 510(k) logic. Gate R2

evaluates EU MDR GSPR mapping. It checks whether each relevant requirement is linked to acceptable source evidence and whether open items are owned, dated, and visible. AI outputs must cite source documents, and human reviewers must approve claims before they enter formal quality or regulatory records.

Supply Adaptation Loop

The Supply Adaptation Loop, or S-Loop, focuses on supplier and operational knowledge. It supports supplier early warning, alternate-source readiness, change impact assessment, and total landed cost modeling. Its purpose is to help

supply chain and quality teams detect risk before disruption appears through missed delivery or production delay.

The S-Loop uses a Supplier Early Warning Score that combines delivery performance, quality signals, and external alerts. It also links supplier events to regulatory implications. For example, if a supplier of a critical material becomes high risk, the framework checks which products, technical files, and validation records depend on that supplier. This prevents supplier substitutions from being treated as a purely commercial decision.

Table 3. Dual-loop framework and decision gates

Loop	Gate	Primary function	AI-supported activity	Human accountability
Regulatory Readiness Loop	R1: Predicate suitability	Tests whether FDA 510(k) predicate logic is coherent	Compares intended use, technology, performance, safety, and labeling	Regulatory Affairs approves predicate strategy
Regulatory Readiness Loop	R2: GSPR completeness	Tests whether EU MDR evidence mapping is complete	Links requirements to controlled evidence and flags missing items	Regulatory Affairs and Quality approve evidence map
Supply Adaptation Loop	S1: Supplier early warning	Detects supplier risk before delivery failure	Scores delivery, quality, and external risk indicators	Supply Chain approves sourcing response
Supply Adaptation Loop	S2: Change impact	Tests whether supplier or material change affects regulatory status	Links supplier change to products, records, and filings	Quality and Regulatory Affairs approve change decision
Governance layer	CSA / validation memo	Preserves auditability of enterprise AI tools	Logs model use, source links, prompts, and approvals	QA/RA Systems Lead owns governance record

Human Accountability and AI Governance

The framework is built on human accountability. AI may retrieve evidence, suggest mappings, draft first-pass text, or flag supplier risk. It cannot approve regulatory claims, supplier changes, or quality-system decisions.

Three governance controls were used. First, all AI-generated regulatory statements required source citation. Second, reviewers followed a two-pass verification standard. One reviewer

checked technical correctness, while a second reviewer checked regulatory and quality-system suitability. Third, the team monitored groundedness, defined as the percentage of AI-generated claims supported by retrieved source documents.

The ERP and eQMS remained the formal systems of record. AI read from controlled sources, generated decision-support outputs, and returned approved artifacts through established quality-system processes.

Results and Discussion

governance, and financial performance. Table 4 summarizes the main baseline and pilot results.

Regulatory and Supply Chain Outcomes

The pilot produced measurable improvement across regulatory throughput, supply resilience,

Table 4. Baseline versus pilot performance results

Metric	Baseline	Pilot result	Improvement or outcome
Drafting cycle time	6.5 weeks	4.4 weeks	32.3% reduction
Defect density	2.8 errors/page	0.9 errors/page	67.8% reduction
Review iterations	4.2 cycles	2.5 cycles	40.4% reduction
Evidence retrieval time	12.5 hours	1.5 hours	88.0% reduction
GSPR first-pass completeness	65%	95%	30 percentage-point increase
Lead-time recovery	8.5 weeks	5.2 weeks	38.8% faster
Sole-source exposure	42% of BOM	34% of BOM	8 percentage-point reduction
Critical stockouts	2 events	0 events	100% avoidance
Median groundedness score	Not systematically measured	0.94 by Month 6	Within target range
Gross value generated	Not applicable	\$1.55 million	Positive gross value
Pilot investment	Not applicable	\$0.85 million	Cost base
Net value	Not applicable	\$0.70 million	Positive net value
Six-month ROI	Not applicable	82.4%	Positive ROI

The strongest R-Loop result was not autonomous drafting. It was faster controlled retrieval. Drafting cycle time decreased from 6.5 weeks to 4.4 weeks, defect density decreased from 2.8 errors per page to 0.9 errors per page, and evidence retrieval time declined from 12.5 hours to 1.5 hours. Regulatory writers entered review mode earlier because the system provided source-linked scaffolds that could be checked and approved.

The S-Loop also improved resilience. Lead-time recovery improved from 8.5 weeks to 5.2 weeks. Sole-source exposure declined from 42% to 34% of the bill of materials. Critical stockouts decreased from two baseline events to zero pilot events. In one case, the AI-supported model detected financial and external risk signals at a tier-2 silicone supplier before delivery failure occurred, allowing secondary-supplier qualification.




Figure 2. Baseline Versus Pilot Improvement across Key Regulatory and Supply Chain Metrics

Note: Baseline values are normalized to 100 for visual comparison because the metrics use different units. Lower values indicate better performance. Bar labels show the original measured values. Sole-source exposure declined from 42% to 34% of the bill of materials, an 8 percentage-point reduction.

Total Landed Cost Simulation

The S-Loop also created strategic value through total landed cost modeling. The model compared Asia-based manufacturing and North America near-shore manufacturing under tariff, logistics, labor, and inventory assumptions. It identified an 18.5% composite tariff threshold at which the near-shore option became cost-neutral for selected product lines.

This changed near-shoring from a general strategic discussion into a quantified decision trigger. Below the threshold, the Asia-based option remained cost-favorable due to existing production efficiencies. As tariff exposure increased, the near-shore option became more attractive due to lower logistics latency, reduced inventory carrying cost, and shorter replenishment cycles.




Figure 3. Total Landed Cost Tipping Point for Mexico-Based Manufacturing

Note: The figure compares indexed total landed cost under different tariff scenarios. The Shenzhen curve rises as tariff exposure increases, while the Mexico curve remains relatively stable. The crossing point at 18.5% marks the threshold at which Mexico-based manufacturing becomes cost-neutral relative to Shenzhen.

Governance and Groundedness Results

Governance results were central to the pilot. In regulated operations, speed alone is not success. A faster workflow that creates unsupported claims or weakens auditability is not acceptable.

By Month 6, the median groundedness score reached 0.94. This showed that AI outputs were largely tethered to controlled source evidence.

The team also used red-team tests that asked the system to summarize missing or contradictory evidence. The correct response was not creative completion. The correct response was escalation, such as identifying that source data were missing. The governance design also reduced automation bias. Reviewers were not allowed to approve claims without opening source links. This created useful friction and reinforced that AI output was a recommendation to verify, not a decision to accept.

Figure 4. Governance Guardrails for Human-in-the-Loop AI




Figure 4. Governance Guardrails for Human-in-the-Loop AI

Note: The figure shows the control stack used to keep AI-assisted regulatory and supply chain outputs auditable. AI supports retrieval, mapping, scoring, and drafting, while source citation, two-pass human review, groundedness testing, and eQMS / ERP storage preserve accountability.

Financial Performance and Pilot Propositions

The pilot generated positive six-month financial value. Gross value was estimated at \$1.55 million

against an \$850,000 investment, producing approximately \$700,000 in net value and an 82.4% ROI.

Table 5. Return on investment summary

Category	Cost, 6 months	Gross value generated	Net contribution
Software and cloud	\$350,000	Not applicable	-\$350,000
Labor, five FTEs	\$500,000	\$350,000 efficiency value	-\$150,000
Inventory optimization	Not applicable	\$450,000 capital release	+\$450,000
Revenue acceleration	Not applicable	\$750,000 new sales	+\$750,000
Total	\$850,000	\$1,550,000	+\$700,000

The results supported the three pilot propositions. Drafting cycle time exceeded the 25% improvement threshold while defect density decreased. Lead-time recovery improved, critical stockouts were avoided, sole-source exposure declined, and the pilot generated positive net value and ROI.

The broader finding is that AI value came from integration rather than isolated automation. AI did not simply make drafting faster. It connected evidence, supplier risk, quality records, and decision gates. The business value came from converting fragmented records into decision-ready knowledge.

Managerial Implications

The study offers five managerial implications for small and mid-sized regulated manufacturers.

First, data governance must be treated as ongoing operating discipline. Duplicate supplier records can produce false diversification, while outdated evidence can produce unsupported regulatory outputs. Second, predictive detection may be more realistic than capital expansion for small firms.

Third, AI governance must be built from the beginning through human review, source citation, groundedness monitoring, red-team testing, and audit logs. Fourth, regulatory and supply chain teams should not operate as separate information systems. Fifth, AI should be implemented as an overlay rather than a replacement for ERP or eQMS systems.

Table 6. Managerial implications and transfer conditions

Managerial implication	Practical meaning	Transfer condition
Treat data governance as OpEx	Fund data quality as an ongoing operating discipline	Firm has ERP or supplier data debt that affects decision quality
Use predictive detection before disruption	Act on supplier risk signals before delivery failure occurs	Firm can monitor delivery, quality, financial, and external signals
Govern AI from the beginning	Require source links, human approval, red-team testing, and audit logs	Firm operates in a regulated or audit-sensitive environment
Connect regulatory and supply data	Treat supplier events and regulatory readiness as linked	Firm has shared products, materials, and evidence dependencies
Use AI as an overlay	Preserve ERP and eQMS as systems of record	Firm lacks budget or appetite for full platform replacement

Limitations and Future Research

This study has limitations. It is based on a single-firm case, so the findings provide analytical insight rather than population-level generalizability. The pilot window was six months, which shows short-cycle improvement but not long-term durability. The scope is limited to Class II diabetes devices and Western regulatory pathways. Extension to Class III PMA devices, drug-device combinations, China NMPA pathways, or other product categories requires further study.

The pilot evaluated enterprise AI used for regulatory and supply chain operations. It did not evaluate AI embedded in medical devices or clinical decision-support tools. Results also depend on data readiness and organizational discipline.

Future research should test the framework across multiple firms, product categories, and regulatory environments. Additional research should examine model drift, cybersecurity implications, AI validation practices, and the durability of AI governance roles as firms scale. Further study is also needed on agentic AI in

regulated operations, especially where systems move from retrieval and drafting support toward controlled task execution.

Conclusion

Small and mid-sized Class II medical device manufacturers face a structural challenge. Regulatory, quality, and supply chain complexity can grow faster than the human capacity available to manage it. This creates the Scale-Up Paradox. The problem is not only insufficient staffing, it is fragmented business information.

This paper shows that governed AI can help address that problem when implemented as knowledge infrastructure rather than uncontrolled automation. The dual-loop framework tested at New Horizon Biotech connected regulatory evidence, supplier risk signals, quality records, and decision gates. The pilot improved throughput, supply recovery, documentation quality, groundedness, and financial performance. The study does not claim that AI can replace regulatory, quality, or supply chain professionals. Its conclusion is more practical. AI can help them find evidence faster, detect risk earlier, coordinate decisions more

effectively, and preserve auditability when governance is built into the operating model.

For small and mid-sized manufacturers, predictive resilience is therefore more than a technology shift, it is an operating model shift.

References

- Ahmadi, S. (2024), 'AI-driven strategies for enhancing medical device regulatory compliance,' Regulatory Focus.
- Amaral, A. and Paiva, T. (2024), 'Global regulatory challenges for medical devices: impact on innovation and market access,' Applied Sciences, 14 (20), 9304.
- Chatterjee, S. and Chaudhuri, R. (2022), 'Supply chain sustainability during turbulent environment: examining the role of firm capabilities and government regulation,' Operations Management Research, 15 (3), 1081-1095.
- Davenport, T. H. and Prusak, L. (1998), Working Knowledge: How Organizations Manage What They Know, Harvard Business School Press, Boston.
- European Commission. (2017), Regulation (EU) 2017/745 on medical devices, European Commission, Brussels.
- European Commission, Medical Device Coordination Group. (2022), MDCG 2022-21: Guidance on Periodic Safety Update Report according to Regulation (EU) 2017/745, European Commission, Brussels.
- International Medical Device Regulators Forum. (2022), Machine Learning-enabled Medical Devices: Key Terms and Definitions, IMDRF/AIMD WG/N67 Final.
- International Organization for Standardization. (2016), ISO 13485:2016 Medical devices: quality management systems: requirements for regulatory purposes, ISO, Geneva.
- KPMG. (2024), Supply chain trends 2024: the digital shake-up, KPMG International.
- McKinsey and Company. (2024), Generative AI in the pharmaceutical industry: moving from hype to reality, McKinsey and Company.
- National Academies of Sciences, Engineering, and Medicine. (2022), Building Resilience into the Nation's Medical Product Supply Chains, The National Academies Press, Washington, DC.
- National Institute of Standards and Technology. (2023), Artificial Intelligence Risk Management Framework (AI RMF 1.0), NIST AI 100-1, National Institute of Standards and Technology, Gaithersburg, MD.
- Nonaka, I. and Takeuchi, H. (1995), The Knowledge-Creating Company: How Japanese Companies Create the Dynamics of Innovation, Oxford University Press, New York.
- U.S. Food and Drug Administration. (2024), Medical devices; quality system regulation amendments, Final Rule, U.S. Food and Drug Administration.
- U.S. Food and Drug Administration. (2024), Premarket notification 510(k), U.S. Food and Drug Administration.
- Yin, R. K. (2018), Case Study Research and Applications: Design and Methods, 6th ed., SAGE Publications, Thousand Oaks.