PHARMA.AERO
WE CONNECT PHARMA

Industry Opinion Paper

A PHARMA.AERO PROJECT | SEPTEMBER 2026

AI in Pharmaceutical & MedTech Supply Chains

From Capability to Accountability:

Why the Future of AI in MedTech and Pharmaceutical Supply Chains Will Be Defined by Trust, Not Technology

Table Of Contents

I. Executive Summary

Artificial Intelligence has become a strategic priority across pharmaceutical and MedTech supply chains. Investment is accelerating and pilot programmes are multiplying, yet industry discussions remain disproportionately focused on algorithms, capability rates, and productivity metrics. The promise of AI in this sector is substantial, but so too are the conditions under which that promise can be responsibly realised. The challenge is not simply adopting advanced technologies, it is ensuring that their use aligns with the sector's stringent expectations for safety, compliance, and accountable decision-making.

Four realities are becoming increasingly apparent. First, AI generates its greatest and most durable value when it **augments human decision-making** rather than replacing human decision-makers entirely; technology amplifies expertise, it does not substitute for it.
Second, the most significant barriers to scaling AI are not found in model performance or algorithmic sophistication. They reside in **data fragmentation, governance immaturity, and regulatory uncertainty,** structural challenges that no amount of computational power can resolve on its own.
Third, **accountability for decisions** with patient safety and regulatory consequences remains non-transferable, regardless of how capable or autonomous an AI system becomes; the organisation, and the individuals within it, bear responsibility that cannot be algorithmically offloaded.
Fourth, and perhaps most critically, organisations require **structured readiness pathways**, encompassing technology, people, and governance together, before they can responsibly increase the degree of authority delegated to AI systems.

The organisations most likely to succeed over the next decade will not necessarily be those deploying the most advanced artificial intelligence. They will be those capable of integrating AI capability with regulatory compliance, human expertise, and cross-sector collaboration into a genuinely trusted operating model: one in which technology and human judgement reinforce rather than displace one another. This ambition is more demanding than simply deploying technologies, yet it is essential in a sector where failure does not merely incur financial loss but can compromise product integrity, breach regulatory obligations, and ultimately endanger patient safety.

The central strategic question, then, is no longer:

> What can AI achieve?

It has become:

> What decisions can responsibly be delegated to AI systems and under precisely what conditions?"

II. Why Pharma and MedTech Supply Chains Are Different from Typical AI Adoption Stories

Pharmaceutical and MedTech supply chains operate in environments where decisions carry direct implications as mentioned earlier. In these settings, an AI- or human-driven misjudgement can have consequences that extend well beyond operational disruption, affecting the safe delivery and appropriate use of therapies.

These supply chains are shaped by rigorous regulatory expectations that govern how products are stored, transported, documented, and traced throughout their lifecycle. Requirements related to distribution practices, traceability, data integrity, and auditability create an operating context in which efficiency improvements, while valuable, cannot be the sole basis for increasing AI involvement. Any progression toward greater AI-supported decision-making must be evaluated against a consolidated set of criteria: safeguarding patients, maintaining product quality, ensuring regulatory defensibility, and preserving transparent, explainable decision pathways.

Within this framework, the adoption of AI follows a more deliberate trajectory than in sectors where operational errors carry fewer downstream implications. The question is not whether AI can be applied, but how its use can be aligned with the sector's expectations for accountability, reliability, and traceability. As a result, pharmaceutical and MedTech supply chains provide a distinctive context: one that requires AI to enhance decision-making without diminishing the governance structures that underpin trust in the system.

III. Where AI Actually Creates Value Today

Industry discourse has a persistent tendency to leap directly to its most dramatic conclusions: autonomous supply chains, self-optimising networks, and AI agents operating with minimal human involvement.

This tendency, while understandable, blurs aspiration with current operational reality and bypasses the more instructive question of where AI is creating tangible, measurable value today.

In practice, value creation in pharmaceutical and MedTech supply chains follows a coherent and observable progression, unfolding across three distinct stages of organisational and technological maturity.

STAGE 01

INSIGHT AUGMENTATION

AI acts as an intelligence layer across fragmented data, giving teams faster visibility, natural-language access to information, and simplified investigation. It does not replace judgment; it strengthens it by enabling better, quicker, and more informed decisions. The core outcome is superior insight: the foundation for all subsequent AI value.

STAGE 02

EMBEDDED AI ANALYTICS

AI acts as an intelligence layer across fragmented data, giving teams faster visibility, natural-language access to information, and simplified investigation. It does not replace judgment; it strengthens it by enabling bWith robust, well-integrated data foundations, AI evolves into a genuine decision-support capability, strengthening functions such as demand forecasting, inventory optimisation, scenario evaluation, and supply-risk detection. At this stage, human roles shift from generating analyses to validating and acting on AI-derived insights. The outcome is not only faster decision-making but demonstrably higher decision quality in contexts where errors carry both clinical and commercial implications.etter, quicker, and more informed decisions. The core outcome is superior insight: the foundation for all subsequent AI value.

STAGE 03

BOUNDED AGENCY

With clearly defined governance parameters in place, AI systems begin to carry out routine operational actions, such as shipment monitoring, exception handling, and standardised supplier communications, without requiring approval for each individual step. Their scope of action is intentionally limited, guided by explicit rules, escalation pathways, and continuous human oversight. The benefit of this model is reduced decision latency: well-defined, low-risk tasks can be executed promptly while remaining anchored in a controlled and accountable framework.

3 DISTINCT OUTCOMES

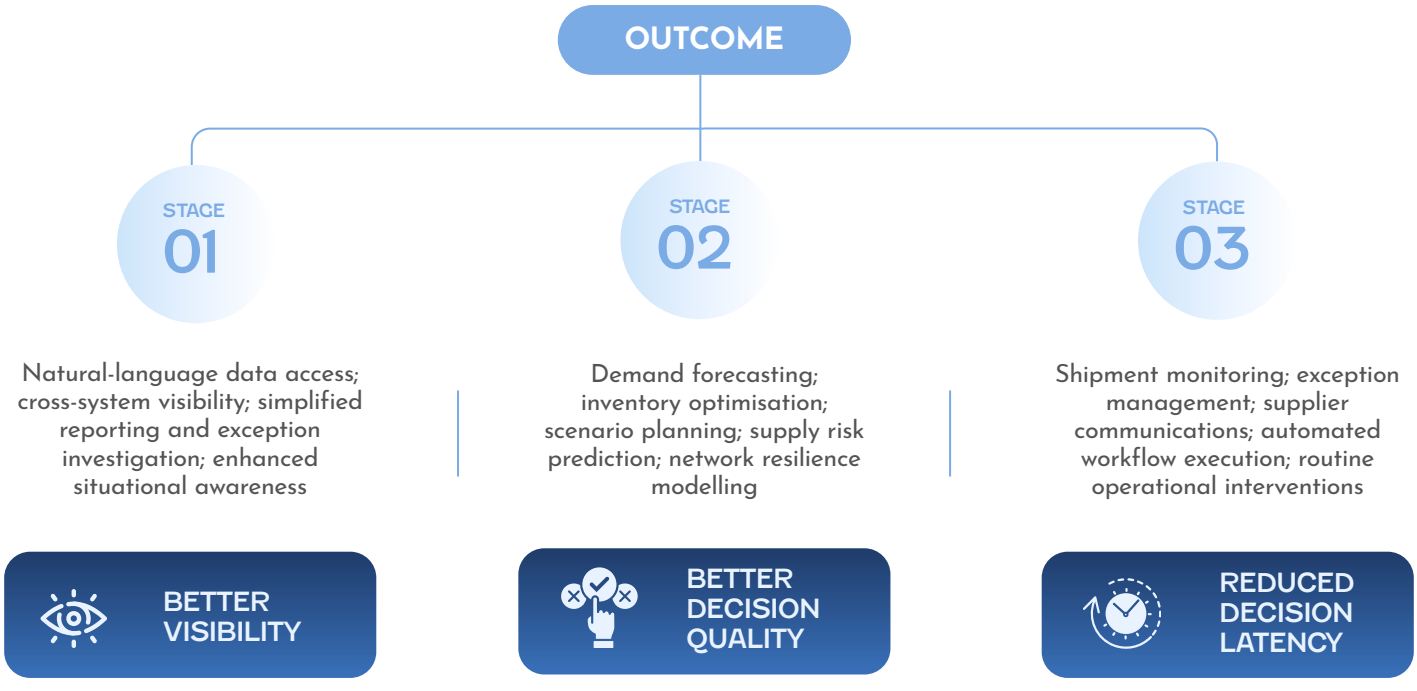

To clarify how different forms of AI contribute to value creation across pharmaceutical and MedTech supply chains, it is useful to distinguish the major capability families that underpin today's applications. Each category supports a different aspect of insight generation, decision quality, or bounded operational autonomy. The table below summarises the primary AI capability types relevant to this sector and the specific roles they play in supply chain operations.

AI CAPABILITY	PRIMARY USE IN PHARMA & MEDTECH SUPPLY CHAINS
MACHINE LEARNING	Forecasting, anomaly detection, risk scoring
DEEP LEARNING	Pattern recognition, sensor data interpretation
EXPLORATORY ANALYTICS	Baseline trend detection, distribution analysis, historical benchmarking, control tower applications
OPTIMISATION AI	Routing, scheduling, resource allocation
NATURAL LANGUAGE PROCESSING (NLP)	Regulatory scanning, supplier risk signals
COMPUTER VISION	Quality inspection, counterfeit detection
AGENTIC AI	Autonomous workflows, exception handling. Can be deterministic (rule-based) or probabilistic (learned)
LARGE LANGUAGE MODELS (LLM)	Cross-system visibility, reasoning, investigation

IV. The Real Shift: From Execution to Oversight

The most significant impact of AI on pharmaceutical and MedTech supply chains is likely to be organisational rather than technological. Historically, supply chain professionals have spent much of their time on activities that, while essential, are largely procedural: gathering information from multiple systems, reconciling inconsistencies, and carrying out routine decisions based on established rules. As AI systems increasingly support or absorb these tasks, the nature of the work changes. Professionals shift from performing manual information processing to interpreting insights, supervising system-supported actions, and ensuring that decisions remain aligned with regulatory and operational expectations.

Rather than reducing the role of supply chain teams, this evolution enables them to focus more of their time on higher-value activities and to apply their expertise more effectively.

This shift is giving rise to three operating models, each with distinct implications for governance, risk management, and the distribution of responsibility between human decision-makers and AI-enabled systems.

01.
HUMAN-IN-THE-LOOP

AI provides recommendations, but humans retain full decision authority, facilitating maximum auditability and clear accountability. This model remains essential for all decisions that directly affect product integrity, regulatory compliance, and patient safety.

02.
HUMAN-ON-THE-LOOP

AI executes predefined actions autonomously, while humans continuously supervise system performance and intervene when needed, enabling faster routine operations without removing human oversight.

03.
HUMAN-OUT-OF-THE-LOOP

AI operates independently within predefined boundaries, while humans remain responsible for establishing governance and monitoring overall system performance. This is an approach suitable only where risks are low and accountability can still be clearly maintained.

V. Why Trust, Explainability and Accountability Matter More Than Accuracy

The dominant framework through which AI systems are evaluated (model accuracy, predictive performance, benchmark scores) is, in the context of pharmaceutical supply chains, a necessary but deeply insufficient basis for deployment decisions. A system that achieves ninety-five percent predictive accuracy is an impressive technical achievement; it is not, of itself, a deployable asset in a regulated, patient-facing environment. Before any AI system can be legitimately operationalised in this sector, organisations must be able to answer a set of questions that have nothing to do with algorithmic performance: Why did the system make this particular recommendation? Can the reasoning behind that recommendation be articulated and defended during a regulatory audit? Who bears accountability if the recommendation proves incorrect or harmful? And can the decision be challenged, interrogated, and overridden by a human with appropriate authority? These are not peripheral concerns, they are the conditions under which trust is established, and trust is the precondition for sustainable AI adoption.

> *Accuracy builds confidence. Explainability builds trust. Accountability builds legitimacy. Only the third creates the conditions for sustainable adoption; and it is the third that the industry must most urgently address.*

Several risks emerge when organisations deploy AI without first resolving how decisions will be explained and justified.

RISK 1

01 — OPACITY

A central challenge is opacity: many machine-learning models, particularly deep neural networks, generate outputs without offering accessible insight into how those outputs were derived. In many commercial settings, this lack of transparency is an acceptable trade-off for predictive performance. In pharmaceutical and MedTech supply chains, however, decisions must be open to thorough examination, meaning organisations must be able to articulate not only what the system recommended, but why. Without this clarity, decisions cannot be defended to regulators, auditors, or internal quality functions, and the resulting vulnerabilities extend across compliance, legal exposure, and organisational credibility.

RISK 2

02 — ERRONEOUS OR UNFOUNDED OUTPUTS

Other risks arise when organisations deploy AI without clearly defining how outputs will be verified, and one of the most discussed is the risk of erroneous or unfounded outputs. Both humans and AI systems can produce confident but incorrect conclusions; the difference is that generative models can generate such outputs at scale, with fluency that can make inaccuracies harder to detect. In consumer contexts, occasional errors may be tolerable. In pharmaceutical and MedTech operations, where AI-supported insights may influence inventory decisions, supplier assessments, or quality-related judgements, even infrequent mistakes can have consequences that are clinically or commercially unacceptable. For this reason, organisations must maintain robust human verification mechanisms that ensure AI-generated outputs are examined, interpreted, and validated with the same rigour applied to human decision-making.

RISK 3

03 — CAPABILITY BIAS

A further risk, capability bias, is perhaps the most insidious because it operates not through system failure but through human psychology. When AI systems consistently generate accurate and useful recommendations, users naturally begin to extend their trust beyond the boundaries that those systems' actual reliability warrants. Independent professional judgement recedes; recommendations are accepted without the critical interrogation that good governance requires. The consequence is not an AI failure but a human failure enabled by AI, and one that is considerably more difficult to detect and correct than a straightforward algorithmic error. Sustaining organisational knowledge, the practical, experience-based understanding gained from "walking the walk", is essential to ensure that professionals can meaningfully interrogate and challenge AI-generated outputs. Without this retained expertise, oversight becomes procedural rather than substantive, weakening the very governance structures the sector depends on.

RISK 4

04 — DIFFUSION OF RESPONSIBILITY

The last risk concerns the way responsibility can become blurred when decisions emerge from combined human and AI inputs. In such situations, an AI system may generate a recommendation and a human may approve it, yet neither party may feel fully accountable for the outcome. This diffusion of responsibility can create uncertainty about who is answerable for consequential actions. In pharmaceutical and MedTech supply chains maintaining clear, well-defined ownership of decisions is essential to any credible governance framework.

VI. Why Most AI Programmes Struggle to Scale

There is a pattern so consistent across industries, geographies, and use cases that it has become almost axiomatic: AI pilots succeed; enterprise-scale deployments struggle. The controlled environment of the pilot, clean data curated for the purpose, engaged teams motivated by novelty, governance requirements temporarily deferred, rarely survives contact with the full complexity of organisational reality.

When AI programmes attempt to transition from proof-of-concept to production at scale, they encounter a set of structural barriers whose root causes are not technological but organisational, architectural, and relational. Understanding these barriers is the prerequisite for addressing them.

The AI Scalability Struggle

01 DATA AND TECHNOLOGY GAPS

- Fragmented system landscapes
- Data standards
- Integration gaps
- Single version of truth
- Poor-quality data
- AI problem vs data problem

02 ORGANISATIONAL MATURITY GAPS

- Organisational readiness
- Pre-AI operating models
- Workforce reconfiguration
- Change-management capability
- Business value vs unused outputs
- People and processes

03 GOVERNANCE AND COMPLIANCE GAPS

- Validation requirements
- GxP environments
- Audit expectations
- Model monitoring responsibilities
- Regulatory obligations
- Maturing regulatory framework
- Inaction risk
- Adoption velocity

04 ECOSYSTEM COLLABORATION GAPS

- Multi-party networks
- End-to-end insight
- Contractual data-sharing constraints
- Fragmentation
- Upstream/downstream constraints
- Diverging priorities
- Network-level optimisation limits
- Transparent interactions
- Supply-chain friction

VII. A Maturity Model for Responsible AI Adoption

The organisations that succeed with AI will be those that treat it as an ongoing organisational capability, not a technology project. Progress depends on continuously developing, refining, and extending three interdependent pillars that must evolve together if the overall programme is to remain coherent, compliant, and trustworthy.

01. TECHNOLOGY AND DATA FOUNDATION

The first pillar encompasses all the capabilities necessary to ensure that AI systems have access to data of sufficient quality, consistency, and generate reliable outputs. This includes investments in data quality management, system integration architecture, end-to-end supply chain visibility, and the underlying infrastructure (cloud platforms, data warehouses, connectivity standards) that makes scalable AI deployment technically possible. Without a strong foundation in this pillar, investment in the other two is premature; AI capability cannot compensate for data deficiency, and governance frameworks applied to unreliable outputs legitimise unreliable decisions.

02. HUMAN CAPABILITY

The second pillar addresses the workforce and organisational dimensions of AI adoption. It encompasses AI literacy; the degree to which supply chain professionals understand what AI systems can and cannot do, and why; as well as the decision-making skills, change readiness, and governance competence required to operate effectively in an AI-augmented environment. This pillar reflects the recognition that AI systems are only as valuable as the humans who work alongside them; a sophisticated AI capability deployed into an organisation whose workforce is unprepared to engage with it critically, to challenge its outputs when appropriate, and to maintain accountability for outcomes will generate risk rather than value.

A critical component of this capability is knowledge retention: ensuring that professionals maintain the depth of operational and regulatory understanding required to evaluate, contextualise, and, when necessary, override AI-supported recommendations. AI can augment judgement, but only organisations that preserve and cultivate human expertise can exercise accountable decision-making at scale.

03. GOVERNANCE AND COMPLIANCE

The third pillar addresses the frameworks, structures, and processes through which the organisation maintains control of its AI systems and the decisions they influence. This includes validation frameworks aligned with regulatory expectations, accountability structures that make clear who is responsible for which decisions and outcomes, mechanisms for ongoing model monitoring and performance management, and the risk management processes required to identify and respond to AI system failures or drift. In a regulated industry, this pillar is not optional or supplementary, it is the foundation upon which the legitimacy of AI-supported decisions rests, and its absence will ultimately prevent AI adoption from scaling regardless of the quality of the technology or the capability of the workforce.

These three pillars are not independent variables that can be optimised separately and subsequently integrated; they are structurally interdependent, and failure in any of these creates disproportionate risk across the whole. An organisation that deploys sophisticated AI on fragile data infrastructure will undermine the trust that its governance frameworks are designed to create.

One that advances technology and governance without investing in human capability will produce a workforce that supervises AI systems it does not adequately understand. And one that pursues capability and literacy without adequate governance will eventually encounter the regulatory and accountability barriers that make sustainable adoption impossible. Progress must be synchronised, and the organisation's maturity level should be assessed honestly and regularly against all three pillars together.

MATURITY MODEL: FIVE LEVELS OF AI MATURITY

01 FOUNDATION

The organisation establishes the fundamental preconditions for AI adoption: data visibility is achieved, key processes are standardised, and baseline system integration is in place. AI is not yet operationally deployed, but the conditions for responsible deployment are being constructed.

02 ASSISTED DECISION SUPPORT

AI-generated insights and recommendations are made available to human decision-makers, who use them to inform, but not replace, their own judgement. Human review and approval remain integral to all consequential decisions. Organisational confidence in AI outputs is built through transparency and demonstrated reliability.

03 INTEGRATED INTELLIGENCE

AI becomes embedded within planning and operational workflows rather than operating as an external advisory tool. Demand forecasting, inventory management, and supply risk assessment are systematically AI-enhanced. Governance frameworks are formalised and validated. The organisation's decision architecture begins to be redesigned around AI capability.

04 CONTROLLED AUTONOMY

Selected, well-defined categories of operational decision are delegated to AI systems for autonomous execution, subject to human supervision and clearly defined escalation mechanisms. The boundaries of permitted autonomy are explicit, regularly reviewed, and validated against regulatory requirements. Human-on-the-loop oversight is the dominant operating model for these decisions.

05 AUTONOMOUS EXECUTION UNDER GOVERNANCE

AI systems execute defined categories of activity independently, within a comprehensive and continuously monitored governance architecture. Human oversight is systemic rather than transactional. The organisation has demonstrated, and can evidence to regulators, that its AI systems operate reliably, explainably, and within defined risk tolerances across their permitted scope of activity.

The governing principle of this maturity model is one that admits no convenient exceptions: organisations should increase the autonomy delegated to AI systems only when technology readiness, human capability, and governance maturity have advanced together to a level that can sustain and account for that autonomy.

VIII. Strategic Reframes for Supply Chain Leaders

The most consequential barriers to AI adoption in this sector are not primarily technical; they are conceptual. The mental models that supply chain leaders apply to AI, inherited from earlier waves of capability and digital transformation, are frequently misaligned with the realities and demands of responsible AI deployment in a regulated, patient-facing industry. Five specific reframes are proposed, not as incremental adjustments to existing thinking, but as substantive and deliberate shifts in the questions being asked and the priorities being set.

REFRAME 1

FROM "DECIDING TO USE AI"

TO "USING AI TO DECIDE"

The capability framing treats AI primarily as a mechanism for displacing human effort, and it leads organisations towards use cases selected on the basis of process volume rather than strategic significance or governance appropriateness. The delegation framing is more rigorous and more honest: it requires organisations to evaluate each decision domain against criteria of risk, reversibility, regulatory obligation, and accountability before determining whether, and under what governance conditions, that domain can responsibly be handed to an AI system. This is a harder question to answer, but it is the right one and the organisations that answer it with discipline will build more defensible and more durable AI programmes than those that pursue automation for its own sake.

REFRAME 2

FROM "WORKFORCE DECREASE"

TO "COGNITIVE AUGMENTATION"

The framing of AI as a mechanism for reducing headcount is not only strategically shortsighted, it is counterproductive in an environment where the quality of human oversight determines the trustworthiness of AI-assisted decisions. The more accurate and more valuable framing positions AI as a system that extends the cognitive reach of expert professionals: enabling them to process more information, evaluate more scenarios, identify more risks, and exercise better-informed judgement at greater scale. The goal is not fewer people but better decisions; and organisations that design their AI strategies around this objective will develop the human capability and governance maturity that responsible AI adoption requires, rather than eroding the very workforce they need to sustain it.

REFRAME 3

FROM "GOVERNANCE FIGHT"

TO "GOVERNANCE FIRST"

The temptation to deploy AI systems rapidly and to address governance as a subsequent refinement is understandable in a competitive environment where speed to capability appears to confer advantage. In pharmaceutical and MedTech supply chains, however, this sequencing is not merely inadvisable, it is ultimately self-defeating. Organisations that scale AI without first establishing robust governance frameworks will encounter trust failures, compliance barriers, and regulatory scrutiny at the precise point where their AI programmes are most visible and most consequential. Governance is not the overhead that slows AI adoption; it is the foundation that makes sustainable AI adoption possible.

REFRAME 4

FROM "INTEGRATING ADVANCED AI"

TO "ADVANCING DATA INTEGRATION FIRST"

The organisational pressure to deploy advanced AI capabilities (generative models, autonomous agents, predictive analytics) can create a dangerous inversion of the correct sequence of investment. Most pharmaceutical organisations do not, in honest assessment, have an AI problem. They have a data problem: inconsistent standards, fragmented systems, and quality gaps that undermine the reliability of any intelligence built upon them. Advanced AI deployed on unreliable data does not produce advanced intelligence; it produces unreliable outputs at scale, and with sufficient fluency to be convincing. The unglamorous but essential work of data architecture, integration, and quality management must precede investment in the AI capabilities that depend upon it.

REFRAME 5

FROM "PROTECTING BOUNDARIES"

TO "CREATING NETWORK INTELLIGENCE"

The instinct to treat data sharing and inter-organisational coordination primarily through the lens of information security risk and contractual obligation leads to a posture of defensive minimalism that forecloses the most significant value creation opportunities available in this sector. The organisations that will define the next generation of pharmaceutical supply chain performance are those that develop the capability to share appropriate information securely, to coordinate decisions at network level, and to orchestrate AI across the boundaries of individual enterprises in a manner that generates intelligence no single organisation could produce alone. Ecosystem collaboration, approached strategically and governed rigorously, is not a concession to external pressure, it is the next durable source of competitive advantage in this industry.

IX. Conclusions

Across the document, a consistent theme emerges: meaningful AI adoption in pharmaceutical and MedTech supply chains is less about technological breakthroughs and more about the disciplined integration of capability, governance, and human expertise. The sector's operating environment, defined by stringent regulatory expectations, complex multi-party networks, and decisions with direct patient implications, requires AI to be introduced through structured, accountable, and progressively bounded forms of support rather than abrupt shifts toward autonomy.

The three stages of AI capability outlined in this paper illustrate how value is created when organisations advance deliberately: first by improving visibility, then by strengthening decision-support functions, and only thereafter by enabling tightly governed forms of autonomous execution. Each stage demands corresponding maturity in data quality, oversight mechanisms, and professional roles. As AI absorbs procedural tasks, the work of supply chain teams becomes more interpretive and supervisory, reinforcing the need for clarity in decision ownership and explainability. The risks examined (opacity, misalignment, responsibility diffusion, and ecosystem fragmentation) underscore that AI does not simplify governance; it makes governance more central.

What ultimately distinguishes responsible adopters is not the sophistication of their models but the coherence of their operating frameworks. Organisations that invest in robust data foundations, transparent decision pathways, well-defined accountability structures, and constructive collaboration with partners will be positioned to use AI in ways that strengthen both operational performance and regulatory defensibility. Those that treat AI as an organisational capability, supported by people, processes, and governance, rather than a discrete technology initiative will be better equipped to scale it safely and sustainably.

In this sector, the value of AI will be realised not through the pursuit of autonomy for its own sake, but through the creation of systems that enhance the reliability, traceability, and resilience of supply chain decisions. The purpose is clear: to support the safe, timely, and dependable movement of products on which patients rely. Organisations that anchor their AI strategies in this purpose will not only meet the expectations placed upon them, they will shape a more trustworthy and capable supply chain for the future.

X. Acknowledgements

The AI in Pharmaceutical and MedTech Supply Chains project was developed, coordinated, and managed by Pharma.Aero.

The project was led by MSD, a strategic member of Pharma.Aero, with research conducted by three Master students from the Antwerp Management School (AMS), following the master of Global Supply Chain. An external expert also contributed valuable insights and supported the review of the project's findings.

This consortium reflects the strength of industry–academic collaboration, demonstrating how diverse expertise can be combined to produce rigorous, practice-relevant outcomes.

PROJECT CONTRIBUTORS AND AUTHORS:

- Ruud van der Geer, Assoc. Director Supply Chain Management EMEA, MSD
- Benedict Wiedermann, student at the AMS
- Himanshu Singh, student of the AMS
- Diem Nguyen Ngoc, student of the AMS
- Frank Van Gelder, Secretary General, Pharma.Aero / Coach for AMS Master students
- Sara Van Lerberghe, Project Manager, Pharma.Aero
- Philippe Ducarme, Project Expert, External Consultant

References

- Abramova, R., Filatova, M., Sazonava, M., & Serikova, K. (2026). Ethical and practical aspects of AI implementation in the pharmaceutical industry. International Journal of Pharmaceutical and Healthcare Marketing, 20(1), 238–268. https://doi.org/10.1108/IJPHM-10-2024-0103

- Agarwal, A., & Nene, M. J. (2025). A five-layer framework for AI governance: Integrating regulation, standards, and certification. Transforming Government: People, Process and Policy, 19(3), 535–555. https://doi.org/10.1108/TG-03-2025-0065

- Al-Hourani, S., & Weraikat, D. (2025). A systematic review of artificial intelligence (AI) and machine learning (ML) in pharmaceutical supply chain (PSC) resilience: Current trends and future directions. Sustainability, 17(14), 6591.

- Anjomshoae, A., Jacquet, A., Dragicevic, A., & Alahmad, W. (2025). Delay management in pharmaceutical cold chain logistics: A systematic literature review. Proceedings of the Eastern Asia Society for Transportation Studies, 15. https://easts.info/online/proceedings/vol.15/pdf/D_PP3842.pdf

- Arrieta, A. B., Díaz-Rodríguez, N., Del Ser, J., Bennetot, A., Tabik, S., Barbado, A., Garcia, S., Gil-Lopez, S., Molina, D., Larysz, R., & Herrera, F. (2020). Explainable artificial intelligence (XAI): Concepts, taxonomies, opportunities and challenges toward responsible AI. Information Fusion, 58, 82–115. https://doi.org/10.1016/j.inffus.2019.12.012

- Badejo, O., & Ierapetritou, M. (2024). Enhancing pharmaceutical supply chain resilience: A multi-objective study with disruption management. Computers & Chemical Engineering, 188, 108769. https://doi.org/10.1016/j.compchemeng.2024.108769

- Borghoff, U. M., Bottoni, P., & Pareschi, R. (2025). Human-artificial interaction in the age of agentic AI: A system-theoretical approach. Frontiers in Human-Dynamics, 7, 1579166. https://doi.org/10.3389/fhumd.2025.1579166

- Chander, B., John, C., Warrier, L., & Gopalakrishnan, K. (2025). Toward trustworthy artificial intelligence (TAI) in the context of explainability and robustness. ACM Computing Surveys, 57(6), 1–49

- EMA. (2017). Good distribution practice. European Medicines Agency. https://www.ema.europa.eu/en/human-regulatory-overview/post-authorisation/compliance-post-authorisation/good-distribution-practice

- European Commission. (2024). Artificial Intelligence Act. http://data.europa.eu/eli/reg/2024/1689/oj

- Ferreira, R. M. F. D. (2025). Piloting a maturity model for responsible artificial intelligence: A Portuguese case study. Journal of Responsible Technology, 21, 100117. https://doi.org/10.1016/j.jrt.2025.100117

- Gill, A. (2025). Pharmaceutical supply chains: Risks, challenges and strategic response. Journal of Applied Business and Economics, 27(2), 91. https://doi.org/10.33423/jabe.v27i2.7581

- Gough, P. (2026). The evolving role of the qualified person in an AI-driven pharma compliance landscape. Drug Discovery Today, 31(3), 104649. https://doi.org/10.1016/j.drudis.2026.104649

References

- Gstettner, S., Pathak, A., Burke, D., Grönlund, F., & Mallet, T. (2024). How GenAI reimagines supply chain management. BCG. https://www.bcg.com/publications/2024/genai-reimagines-supply-chain
- Hansen, H. F., Jensen, R., & Møller, C. (2024). Understanding artificial intelligence diffusion through an AI capability maturity model. Information Systems Frontiers, 26, 2147–2163. https://doi.org/10.1007/s10796-024-10528-4
- Jayakumar, J. (2025). Cloud compliance intelligence framework for GxP validation in life sciences. International Journal of Applied Mathematics, 38(11). https://doi.org/10.12732/ijam.v38i11s.1214
- Kushwah, S., & Dave, N. (2025). AI hallucination and strategies to overcome: Enhancing human-AI interaction. AIMV 2025, 1–6. https://doi.org/10.1109/AIMV66517.2025.11203756
- Leslie, D. (2019). Understanding artificial intelligence ethics and safety. The Alan Turing Institute. https://doi.org/10.5281/zenodo.3240529
- Manoharan, G., Shameem, A., Sawant, P. D., Goke, N., Pankajam, A., & Panicker, A. (2025). Human-AI collaboration for upskilling workforce in the industry 5.0 era. WCCEST 2025, 1–6. https://doi.org/10.1109/WCCEST66994.2025.11389885
- McKinsey & Company. (2025). The state of AI in 2025: Agents, innovation, and transformation. https://www.mckinsey.com/capabilities/quantumblack/our-insights/the-state-of-ai
- Mosqueira-Rey, E., Hernández-Pereira, E., Alonso-Ríos, D., Bobes-Ascarán, J., & Fernández-Leal, Á. (2023). Human-in-the-loop machine learning: A state of the art. Artificial Intelligence Review, 56(4), 3005–3054. https://doi.org/10.1007/s10462-022-10246-w
- Mylrea, M., & Robinson, N. (2023). AI trust framework and maturity model: Improving security, ethics and trust in AI. IEEE Security & Privacy, 21(4), 56–64. https://doi.org/10.53889/citj.v1i1.198
- Overhaul. (2026). The future of cold chain logistics whitepaper. https://www.overhaul.com/whitepapers/future-of-cold-chain-logistics-whitepaper
- Panigrahy, S. (2025). Human-AI interaction in logistics: Augmenting human capabilities for supply chain excellence. https://doi.org/10.5281/ZENODO.15861979
- Pasas-Farmer, S., & Jain, R. (2025). From discovery to delivery: Governance of AI in the pharmaceutical industry. Green Analytical Chemistry, 13. https://doi.org/10.1016/j.greeac.2025.100268
- Sadiq, R. B., Al-Sarray, A., & Mohammed, H. J. (2021). Artificial intelligence maturity model: A systematic literature review. PeerJ Computer Science, 7, e661. https://doi.org/10.7717/peerj-cs.661
- Xu, Z., Jain, S., & Kankanhalli, M. (2025). Hallucination is inevitable: An innate limitation of large language models. https://doi.org/10.48550/arXiv.2401.11817
- Zhu, C., Xin, J., & Trinh, T. K. (2025). Data quality challenges and governance frameworks for AI implementation in supply chain management. https://pinnaclepubs.com/index.php/PAPPS

Industry Opinion Paper

AI in Pharmaceutical & MedTech Supply Chains

From Capability to Accountability:

Why the Future of AI in MedTech and Pharmaceutical Supply Chains Will Be Defined by Trust, Not Technology

SEPTEMBER 2026

PHARMA.AERO
WE CONNECT PHARMA