



Manufacturing & Service Operations Management

Publication details, including instructions for authors and subscription information:
<http://pubsonline.informs.org>

OM Forum—Supply Chain Management in the AI Era: A Vision Statement from the Operations Management Community

Maxime C. Cohen, Tinglong Dai, Georgia Perakis, Narendra Agrawal, Gad Allon, Robert N. Boute, Gérard P. Cachon, Zhe Chen, Morris Cohen, Rares Cristian, Vinayak Deshpande, Francis de Véricourt, Jan C. Fransoo, Joren Gijbrecchts, Pavithra Harsha, Ming Hu, Pinar Keskinocak, Caleb Kwon, Hau Lee, Sheng Liu, Konstantina Mellou, Ishai Menache, Jason Miller, Serguei Netessine, Tava Lennon Olsen, Jeevan Pathuri, Robert Peels, Yongzhi Qi, Ananth Raman, Anne Robinson, Zuo-Jun Max Shen, Masha Shunko, David Simchi-Levi, Hannah Smalley, Jing-Sheng Song, Jayashankar M. Swaminathan, Christopher Tang, Sridhar Tayur, Maxi Udenio, Jan Van Mieghem, Yuqian Xu, Dennis Zhang

To cite this article:

Maxime C. Cohen, Tinglong Dai, Georgia Perakis, Narendra Agrawal, Gad Allon, Robert N. Boute, Gérard P. Cachon, Zhe Chen, Morris Cohen, Rares Cristian, Vinayak Deshpande, Francis de Véricourt, Jan C. Fransoo, Joren Gijbrecchts, Pavithra Harsha, Ming Hu, Pinar Keskinocak, Caleb Kwon, Hau Lee, Sheng Liu, Konstantina Mellou, Ishai Menache, Jason Miller, Serguei Netessine, Tava Lennon Olsen, Jeevan Pathuri, Robert Peels, Yongzhi Qi, Ananth Raman, Anne Robinson, Zuo-Jun Max Shen, Masha Shunko, David Simchi-Levi, Hannah Smalley, Jing-Sheng Song, Jayashankar M. Swaminathan, Christopher Tang, Sridhar Tayur, Maxi Udenio, Jan Van Mieghem, Yuqian Xu, Dennis Zhang (2026) OM Forum—Supply Chain Management in the AI Era: A Vision Statement from the Operations Management Community. *Manufacturing & Service Operations Management* 28(3):687-705. <https://doi.org/10.1287/msom.2025.1065>

Full terms and conditions of use: <https://pubsonline.informs.org/Publications/Librarians-Portal/PubsOnLine-Terms-and-Conditions>

This article may be used only for the purposes of research, teaching, and/or private study. Commercial use or systematic downloading (by robots or other automatic processes) is prohibited without explicit Publisher approval, unless otherwise noted. For more information, contact permissions@informs.org.

The Publisher does not warrant or guarantee the article's accuracy, completeness, merchantability, fitness for a particular purpose, or non-infringement. Descriptions of, or references to, products or publications, or inclusion of an advertisement in this article, neither constitutes nor implies a guarantee, endorsement, or support of claims made of that product, publication, or service.

Copyright © 2026, INFORMS

Please scroll down for article—it is on subsequent pages



With 12,500 members from nearly 90 countries, INFORMS is the largest international association of operations research (O.R.) and analytics professionals and students. INFORMS provides unique networking and learning

opportunities for individual professionals, and organizations of all types and sizes, to better understand and use O.R. and analytics tools and methods to transform strategic visions and achieve better outcomes. For more information on INFORMS, its publications, membership, or meetings visit <http://www.informs.org>

OM Forum—Supply Chain Management in the AI Era: A Vision Statement from the Operations Management Community

Maxime C. Cohen,^{a,*} Tinglong Dai,^{b,*} Georgia Perakis,^{c,*} Narendra Agrawal,^d Gad Allon,^e Robert N. Boute,^{f,g} Gérard P. Cachon,^e Zhe Chen,^h Morris Cohen,^e Rares Cristian,^c Vinayak Deshpande,ⁱ Francis de Véricourt,^j Jan C. Fransoo,^k Joren Gijsbrechts,^l Pavithra Harsha,^m Ming Hu,ⁿ Pinar Keskinocak,^o Caleb Kwon,^p Hau Lee,^q Sheng Liu,ⁿ Konstantina Mellou,^r Ishai Menache,^r Jason Miller,^s Serguei Netessine,^e Tava Lennon Olsen,^t Jeevan Pathuri,^f Robert Peels,^u Yongzhi Qi,^h Ananth Raman,^v Anne Robinson,^w Zuo-Jun Max Shen,^x Masha Shunko,^y David Simchi-Levi,^c Hannah Smalley,^o Jing-Sheng Song,^z Jayashankar M. Swaminathan,ⁱ Christopher Tang,^{aa} Sridhar Tayur,^{bb} Maxi Udenio,^g Jan Van Mieghem,^{cc} Yuqian Xu,ⁱ Dennis Zhang^{dd}

^aDesautels Faculty of Management, McGill University, Montreal, Quebec H3A 0G4, Canada; ^bCarey Business School, Johns Hopkins University, Baltimore, Maryland 21202; ^cSloan School of Management, Massachusetts Institute of Technology, Cambridge, Massachusetts 02139; ^dSanta Clara University, Santa Clara, California 95053; ^eUniversity of Pennsylvania, Philadelphia, Pennsylvania 19104; ^fVlerick Business School, 1210 Brussels, Belgium; ^gKU Leuven, 3000 Leuven, Belgium; ^hJD.com, Beijing 101111, China; ⁱUniversity of North Carolina at Chapel Hill, Chapel Hill, North Carolina 27599; ^jEuropean School of Management and Technology Berlin, 10178 Berlin, Germany; ^kTilburg University, 5037 AB Tilburg, Netherlands; ^lRamon Llull University (Esade), 08034 Barcelona, Spain; ^mIBM, Armonk, New York 10504; ⁿUniversity of Toronto, Toronto, Ontario M5S 1A1, Canada; ^oGeorgia Institute of Technology, Atlanta, Georgia 30332; ^pUniversity of Texas at Austin, Austin, Texas 78712; ^qStanford University, Stanford, California 94305; ^rMicrosoft, Redmond, Washington 98052; ^sMichigan State University, East Lansing, Michigan 48824; ^tUniversity of Melbourne, Parkville, Victoria 3010, Australia; ^uFlostock, 5251 CC Vlijmen, Netherlands; ^vHarvard University, Cambridge, Massachusetts 02138; ^wRobinson Insights, Redhill, Surrey RH2 9AR, Great Britain; ^xThe University of Hong Kong, Hong Kong; ^yUniversity of Washington, Seattle, Washington 98195; ^zDuke University, Durham, North Carolina 27708; ^{aa}University of California, Los Angeles, Los Angeles, California 90095; ^{bb}Carnegie Mellon University, Pittsburgh, Pennsylvania 15213; ^{cc}Northwestern University, Evanston, Illinois 60208; ^{dd}Washington University in St. Louis, St. Louis, Missouri 63130

*Corresponding authors

Contact: maxccohen@gmail.com, <https://orcid.org/0000-0002-2474-3875> (MaC); dai@jhu.edu, <https://orcid.org/0000-0001-9248-5153> (TD); georgiap@mit.edu, <https://orcid.org/0000-0002-0888-9030> (GP); nagrawal@scu.edu, <https://orcid.org/0000-0002-0822-3335> (NA); gadallon@wharton.upenn.edu, <https://orcid.org/0000-0002-4815-6667> (GA); r-boute@kellogg.northwestern.edu (RNB); cachon@wharton.upenn.edu (GPC); chenzhe19@jd.com (ZC); cohen@wharton.upenn.edu, <https://orcid.org/0000-0002-8644-8820> (MoC); raesc@mit.edu (RC); vinayak_deshpande@kenan-flagler.unc.edu, <https://orcid.org/0000-0002-6681-028X> (VD); francis.devericourt@esmt.org, <https://orcid.org/0000-0003-2832-0738> (FdV); jan.fransoo@tilburguniversity.edu, <https://orcid.org/0000-0001-7220-0851> (JCF); joren.gijsbrechts@esade.edu, <https://orcid.org/0000-0002-0846-6855> (JG); pharsha@us.ibm.com, <https://orcid.org/0000-0002-6049-7739> (PH); ming.hu@rotman.utoronto.ca, <https://orcid.org/0000-0003-0900-7631> (MH); pinar@isye.gatech.edu, <https://orcid.org/0000-0003-2686-546X> (PK); caleb.kwon@mccombs.utexas.edu (CK); haulee@stanford.edu (HL); shg.liu@rotman.utoronto.ca (SL); konstantina.mellou@microsoft.com (KM); ishai@microsoft.com (IM); mill2831@broad.msu.edu (JM); netessin@wharton.upenn.edu, <https://orcid.org/0000-0002-3587-3894> (SN); t.olsen@mbs.edu, <https://orcid.org/0000-0002-5477-6280> (TLO); jpathuri@microsoft.com (JP); robert.peels@flostock.nl (RP); qiyongzhi@jd.com, <https://orcid.org/0009-0004-1434-379X> (YQ); araman@hbs.edu (AnaR); anne.robinson@stanfordalumni.org (AnnR); maxshen@hku.hk, <https://orcid.org/0000-0003-4538-8312> (Z-JMS); mshunko@uw.edu, <https://orcid.org/0000-0003-4864-0798> (MS); dslevi@mit.edu (DS-L); hannahsmalley@gatech.edu (HS); jssong@duke.edu (J-SS); jay_swaminathan@kenan-flagler.unc.edu (JMS); chris.tang@anderson.ucla.edu, <https://orcid.org/0000-0001-9597-7620> (CT); stayur@andrew.cmu.edu, <https://orcid.org/0000-0002-8008-400X> (ST); maxi.udenio@kuleuven.be (MU); vanmieghem@kellogg.northwestern.edu, <https://orcid.org/0000-0002-8954-7613> (JVM); yuqian_xu@kenan-flagler.unc.edu, <https://orcid.org/0000-0001-5733-0412> (YX); denniszhang@wustl.edu, <https://orcid.org/0000-0002-4544-775X> (DZ)

Received: December 5, 2025

Revised: January 27, 2026

Accepted: February 18, 2026

Published Online in Articles in Advance:
March 26, 2026

Note: Because of the atypically large author team and the collaborative nature of this project, coauthorship on this paper should not be considered a conflict of interest with respect to reviewing subsequent submissions, participating in editorial decisions, or providing promotion and tenure evaluations.

<https://doi.org/10.1287/msom.2025.1065>

Copyright: © 2026 INFORMS

Abstract. *Problem definition:* Artificial intelligence (AI) is rapidly transforming the research and practice of supply chain management. Yet its impact depends on how effectively it is integrated with the theories, methods, and fundamental principles of operations management (OM), which must also evolve to account for the informational, incentive, and institutional changes brought by AI. The OM community has an important role and responsibility to lead in shaping not only how AI transforms supply chains but also how the supply chains that enable AI are designed to be sustainable, resilient, and equitable. *Methodology/results:* This vision statement organizes the discussion around five layers of the interaction between AI and supply chain management: intelligence, execution, strategy, human, and infrastructure. It synthesizes recent research and industry practice to show how AI enhances forecasting, planning, decision making, risk management, and human–machine collaboration and also examines the supply chains that support AI. Finally, it highlights persistent challenges in data quality, model integration, governance, and workforce adaptation. *Managerial implications:* Realizing AI’s promise in supply chain management requires reliable data and infrastructure, integration of learning and optimization, transparent and explainable decision systems, and a long-term commitment to human–AI collaboration. Together, these elements form the foundation for resilient, adaptive, and trustworthy supply chains in the AI era.

Keywords: supply chain management • operations management • artificial intelligence • machine learning • optimization • resilience • human–AI collaboration

1. Introduction

Supply chain management (SCM) has deep historical roots. Long before the term existed, ancient civilizations developed extensive trade networks that moved goods across vast distances, and merchants in the medieval era organized increasingly complex flows of materials and products across regions. The Industrial Revolution amplified this complexity substantially, creating factory systems, rail networks, and global distribution channels that required unprecedented coordination. Yet it was not until the 1980s that “supply chain management” emerged as a formal discipline (Cohen 2024). It was then that scholars and practitioners recognized that these coordination challenges demanded not only operational experience but also systematic frameworks. The field developed rigorous approaches to inventory control, production planning, and logistics optimization, building expertise rooted in operations research, systems thinking, and empirical analyses.

The field of artificial intelligence (AI) was formalized in the 1950s when researchers such as Herbert Simon, Allen Newell, and John McCarthy sought to create machines capable of intelligent behavior, namely, systems that could reason, learn, and solve problems previously requiring human cognition (Simon 1987). AI’s history has been tumultuous: early optimism about symbolic reasoning and expert systems gave way to “AI winters” when AI researchers’ high hopes from the 1950s and 1960s were very far from being realized by the 1970s (Lighthill 1973). AI resurged in the mid-to-late 1990s and 2000s driven by breakthroughs such as Deep Blue (Campbell et al. 2002), machine learning (ML) replacing rule-based systems, and advances in deep learning. More recent developments in generative AI and large language models (LLMs) mark yet another leap.

These two fields have now converged with remarkable force. AI is becoming increasingly central to how companies sense demand, allocate resources, route shipments, and respond to disruptions: ML improves forecasting accuracy, reinforcement learning (RL) automates inventory decisions, computer vision enables warehouse robotics, and LLMs promise to make advanced planning tools more accessible. This convergence presents both significant potential and responsibility. Supply chains are complex systems in which decisions cascade across organizations, geographies, and time horizons. Thus, the challenge is not simply technical but integrative: how do we combine what AI does well (pattern recognition, rapid adaptation, processing vast data streams, granular dynamic prediction) with what operations management (OM) provides (structural understanding, constraint logic, rigorous modeling, strategic trade-offs)?

Meaningful integration demands collaboration. AI that remains anchored in loss minimization misses the core driver of supply chain performance: decision quality. At the same time, supply chain models that fail to exploit AI’s capacity to encode rich, nonparametric dynamics risk omitting the very patterns that shape real-world outcomes.

The OM community has an important role and responsibility to lead in shaping how AI transforms supply chains (Dai and Swaminathan 2026). Our discipline’s core insight is that effective coordination requires understanding objectives, constraints, and trade-offs. This feature becomes even more critical as AI expands computational possibilities: algorithmic power must be channeled through sound operational logic. Demand forecasts should guide inventory decisions, but they are not a decision rule by themselves. Inventory policies must also reflect constraints, service-level targets, lead times, and the firm’s risk posture. Likewise, optimization models can recommend actions under stated assumptions, but managers still decide what objectives matter, check whether the assumptions fit the current situation, and make exceptions when conditions change. A natural next question is what it would mean for supply chains to “be better” in the AI era. Even if AI sharply reduces forecast error, planning friction, and execution noise, the field does not run out of problems and challenges. The center of gravity shifts toward governance, accountability, and the human and economic consequences of AI-enabled coordination. Building resilient AI-enabled supply chains requires integrating AI’s adaptive intelligence with OM’s structural understanding of how supply chains actually work amid real-world complexities.

We build on these themes by structuring our discussion around five layers through which AI interacts with supply chain management: intelligence (Section 2), execution (Section 3), strategy (Section 4), human (Section 5), and infrastructure (Section 6). By “intelligence,” we mean the methods that turn data into predictions, representations, and recommendations. By “execution,” we mean the operational decisions and workflows that translate these outputs into actions at scale, including planning and control decisions that must be updated as conditions evolve. By “strategy,” we mean longer horizon choices about network design, sourcing posture, and competitive positioning. The “human” layer focuses on how people interact with AI systems, including trust, accountability, and skill formation. Finally, the “infrastructure” layer concerns the technical and physical foundations that make AI possible, including data pipelines, compute, and the supply chains that support them. We refer to AI methods in three concrete families: ML for prediction and pattern recognition, RL for sequential decision making under feedback, and LLMs for language-based interfaces and workflow

support. This layered approach provides a coherent lens for understanding where AI already delivers value, where bottlenecks persist, and where the OM community is uniquely positioned to guide responsible and effective implementation. Section 7 provides sector-specific case studies; Section 8 discusses readiness, implementation, and governance considerations; and Section 9 concludes with implications for researchers, practitioners, and educators.

This vision statement emerged through an extensive collaborative process that included 42 OM researchers, industry practitioners, and technology leaders who contributed to the book *AI in Supply Chains: Perspectives from Global Thought Leaders* (Cohen and Dai 2026). With the support of the broader OM community, we decided to partner again to write this article to synthesize and communicate our findings to a wider academic and practitioner audience. Rather than presenting a singular perspective, this article combines insights from diverse expertise areas: from theoretical foundations in optimization and simulation to hands-on implementation experiences at major corporations and logistics providers. Each contributor brought deep domain knowledge; some address how RL can enhance inventory control, identifying its practical limitations, others examine human–AI collaboration dynamics in high-stakes service environments, and others explore infrastructure bottlenecks in AI’s own supply chain or share implementation lessons learned through years of trial and refinement. This breadth reflects our conviction that technology alone cannot drive AI transformation. This transformation requires simultaneous attention to algorithms, infrastructure, efficiency, and ethics as well as automation and human capability development. By presenting this unified vision, our goal is for the OM community to guide practitioners toward implementations that are sound, effective, and responsibly governed. Ultimately, we hope that this collective perspective will stimulate further research, ground public discourse in technical and organizational realities, and support decision makers seeking to harness AI in ways that enhance resilience, productivity, and social welfare.

2. The Intelligence Layer: From Predictive to Generative AI

In this section, we provide an overview of the intelligence layer, describing how predictive and generative AI techniques—including ML, RL, and LLMs—enhance forecasting, simulation, and optimization to support supply chain decisions. By integrating AI with OM models, the intelligence layer forms the analytical backbone of AI-enabled supply chains and sets the stage for subsequent discussions on decision-focused learning and human interaction.

2.1. Integrating ML, Simulation, and Optimization

Whereas AI has greatly enhanced our ability to solve SCM problems (or, more broadly, OM problems), SCM knowledge forms the foundational structure for effective decision making, which can be enhanced but should not be replaced by AI. First, SCM principles can guide data collection and help identify relevant variables to record, such as on-hand inventory or inventory position, which is critical to forming high-quality databases for ML training, model simulation, and optimization. Second, supply chain operations models encode critical domain knowledge and network structures that do not need to be relearned from data. For example, inventory systems follow a natural rhythm of replenishment cycles, and this would govern any learning process.

ML, which serves as an adaptive intelligence layer, adds value on top of these structures by capturing complex patterns and adapting to dynamic conditions. Specifically, ML is adept at uncovering new relationships among many variables from data, such as modeling how demand depends on promotions, weather, seasonality, and customer profiles (Cohen et al. 2022, Lei et al. 2024). Advanced demand prediction models can rely on hundreds of dynamic variables coming from both internal and external data sets. These models can be frequently retrained to capture subtle variations in the underlying market environment and can even take advantage of domain adaptation techniques to learn across domains and effectively transfer knowledge (Tarighat et al. 2025). Moreover, advanced RL techniques enable the learning of complicated and nonparametric policy functions in response to varying environments. Finally, anomaly detection methods can be trained on vast amounts of data to detect early signals of rare adversarial events, such as panic buying and supply chain disruptions (Adulyasak et al. 2024). The learned relationships and functions bring flexibility and adaptiveness to classic OM models and tools, including simulation and optimization, whereas OM knowledge and critical trade-offs ensure that decisions follow the physical rules of the operational system with interpretability.

By combining the structural properties of OM models with the adaptive intelligence of ML techniques, we can design decision systems that are not only structurally sound but also well-suited to the complexities of real-world operations (Hu and Liu 2026). In essence, OM domain knowledge defines the skeleton of the problem by articulating the objectives, constraints, and trade-offs, whereas AI works within that framework to learn and adapt. For instance, this integration can be formalized through a semiparametric regression framework with deep neural networks, in which OM structural properties encode the parametric backbone and ML supplies nonparametric flexibility.

Recent work applies this framework to applications in three foundational OM domains: inventory control, service operations, and consumer choice modeling for pricing and revenue management (Hu et al. 2025). For small data, large-scale decision problems, we can first pretrain a designed transformer model on large-scale, domain-informed synthetic data that encodes OM knowledge and structural features of the decision environment and then fine-tune the model with a few real observations (Zhang et al. 2026).

2.2. Decision-Focused AI in Supply Chains

Beyond structural integration, another critical dimension is how we train these AI systems to align predictions with decisions. Contemporary decision making in supply chains and enterprise operations has long relied on a sequential pipeline: predictive models forecast uncertain quantities such as demand, and optimization routines then prescribe actions, such as pricing decisions (e.g., Cohen et al. 2017). This predict-then-optimize approach is extensively studied and widely adopted, particularly as advances in ML have driven a focus on predictive accuracy. Despite these advances, a central challenge remains. Improvements in prediction accuracy do not always yield proportional improvements in decision outcomes because optimization depends on all predicted variables in a highly interconnected and nonlinear fashion. Moreover, because of the often discrete nature of optimization problems, individually small prediction errors can result in fundamental shifts in the final decision space. Thus, separating forecasting from optimization can lead to inefficiencies and misalignment with actual operational goals.

Decision-focused AI aims to address this gap by embedding optimization objectives directly into model training. Rather than minimizing standard prediction errors, models are trained with loss functions defined in terms of downstream decision costs; see, for example, Amos and Kolter (2017), Donti et al. (2021), Elmachtoub and Grigas (2022), and Cristian et al. (2023, 2025). This ensures that learning aligns with the ultimate goal of making better decisions. Applications such as inventory allocation clearly demonstrate the advantages: by learning demand forecasts together with the corresponding optimization task, decision-focused methods can reduce costs, often requiring only point forecasts, avoiding the complexity of full distributional modeling.

The paradigm also extends beyond learning. The emergence of LLMs and agentic AI provides an additional pathway: enabling practitioners to specify and execute optimization problems through natural language. By combining decision-aware training with accessible interfaces, future systems can bridge predictive analytics, optimization, and human interaction, democratizing access to advanced decision-making tools.

However, significant challenges remain. Training decision-focused models is computationally intensive, and they are often narrowly tailored to specific tasks. Additionally, interpretability is often reduced as optimized point forecasts may not correspond to the true underlying distributions. Extending these methods to multistage decision problems, scaling them to large supply chains, and improving generalization across tasks are important areas for future research.

Overall, decision-focused AI represents a shift in perspective: aligning predictive modeling with prescriptive optimization to improve operational outcomes, leveraging new AI interfaces to broaden accessibility and practical impact.

2.3. Large-Language Decision Agents

The democratization potential mentioned earlier becomes concrete when we consider how LLMs can transform day-to-day SCM at a global scale. Whereas the previous section addressed how models are trained, LLMs address a different bottleneck: the gap between sophisticated optimization tools and the practitioners who use them. SCM requires addressing a variety of complex decision-making challenges from sourcing strategies to planning and execution. Over the last few decades, advances in computation and information technologies have enabled the transition from manual, intuition- and experience-based decision making to more automated and data-driven decisions using a variety of tools that apply optimization techniques. At a high level, these techniques use mathematical methods to improve decision making.

Yet business planners and executives still need to spend considerable time and effort to (i) understand and explain the recommendations coming out of these technologies, (ii) analyze various scenarios and examine what-if questions, and (iii) update the mathematical models used in these tools to reflect current business environments. Addressing these challenges often requires data science teams and/or technology providers to explain results or implement necessary model changes, and this can significantly slow decision making.

LLMs offer a solution to these bottlenecks. Recent work shows how LLMs can democratize supply chain technology by facilitating the understanding of tools' outcomes and enabling direct interaction with supply chain systems without a human in the loop (Li et al. 2023, Menache et al. 2025). These advances can substantially reduce the time to decision from days or weeks to minutes, dramatically increasing planners' and executives' productivity, efficiency, and impact (Simchi-Levi et al. 2025).

But how can LLMs address these types of challenges accurately and efficiently? An important aspect of supply chain planning tools is the use of mathematical

models, typically optimization methods, to generate a recommendation (i.e., a supply chain plan), and these optimization methods have their own language. Indeed, many optimization tasks are written in the form of mathematical programs that translate the structure of the supply chain and business requirements into a mathematical model. Introducing LLMs into this environment raises an important concern: hallucinations and, more broadly, the reliability of LLM-generated formulations. Even advanced models can occasionally produce confident but incorrect suggestions, such as omitting a key constraint or mischaracterizing a linear problem. Another difficulty is the stochastic nature of LLMs, which can produce different outputs for the same input. This does not change the definition of optimality, but it can complicate how optimal decisions are identified and verified in practice, especially when the model generates multiple plausible responses and the decision logic is not fully transparent. In practice, these risks can be mitigated through guardrails that constrain LLM outputs and ensure consistency with the underlying optimization logic. The key insight is not to replace these mathematical models with LLMs but rather to use them in tandem by diligently prompting and documenting all the governing rules. Menache et al. (2025) describe the framework; the implementation at Microsoft Cloud Supply Chain's demand forecasting and fulfillment decisions; and the impact on speed, accuracy, and productivity of decision makers. Additional examples in procurement illustrate the wide spectrum of decisions that can be impacted by this new technology.

Before closing this section, we distinguish between an LLM and an agent. An LLM is a model that generates and transforms text, whereas an agent can be defined as a broader system that may use an LLM as one component to plan steps, call tools, and take actions under defined constraints and oversight. As LLMs are increasingly used as interfaces to enterprise systems, standards for connecting models to external context and tools are emerging. One example is the model context protocol, which aims to standardize secure access to data and tool calls, making it easier to connect LLM-based assistants and agents to enterprise data sources and operational systems.

3. The Execution Layer: Planning, Inventory, and Automation at Scale

The execution layer translates insights from the intelligence layer into concrete actions across the supply chain. It encompasses supply chain planning, inventory control, and the automation of warehouse and transportation operations. By focusing on how decisions are operationalized at scale, this section explains how optimal machine learning (OML) bypasses traditional forecasting to link data directly to planning

decisions, how RL can automate inventory policies, and how AI-powered robotics and internet of things (IoT) enable adaptive fulfillment and logistics.

3.1. Reimagining Supply Chain Planning Using ML: A Road Map to Agility and Resilience

Whereas the intelligence layer addresses what decisions to make, the execution layer determines how those decisions are operationalized in practice. We now discuss a new paradigm—OML—to transform supply chain planning in the face of global volatility and disruption (Agrawal et al. 2024). Traditional planning methods rely heavily on demand forecasting, which often fails because of data fragmentation, misaligned incentives, and limited responsiveness. The application of ML to supply chain planning typically focuses on using the available extended data to improve forecasts. However, even if perfect forecasts without forecast errors were available, they would not be sufficient to make good supply chain decisions. These decisions are impacted by complex resource constraints, customer-specific business requirements, differences in profitability by customer segments and by regions, and both local and global service-level requirements.

The OML framework bypasses forecasting by directly linking granular supply and demand data to planning decisions through an ML-enabled decision engine, digital twin simulations, and an end-to-end data architecture. The OML methodology optimizes the relationship between key data inputs, supply chain planning decisions, and key financial and operational performance indicators. It utilizes a technology infrastructure that enables efficient access and storage of relevant data along with computing resources and advanced optimization techniques that can solve large-scale problems. It also enables fast and effective sales and operations planning processes.

Two Fortune 150 case studies, both conducted by AD3 Analytics, demonstrate OML's ability to significantly improve service levels and reduce inventory costs (see Cohen et al. 2026, section 5). The first is a semiconductor equipment manufacturer responsible for delivering contracted, warranty, and on-demand support services to customers who purchase their products. They operate a complex global supply chain that distributes more than 50,000 stock keeping units to support, in a cost-effective and customer-focused manner, thousands of service contracts and warranties (Agrawal et al. 2024). The second is the supply chain for an advanced consumer electronics product that is manufactured by a contract manufacturer in Asia, shipped to the company's primary distribution center in the United States, and sold through physical as well as e-commerce channels such as Walmart, Best Buy, and Amazon.

In contrast to conventional ML tools, OML generates interpretable and actionable decisions and supports dynamic scenario analysis to build resilience and agility. OML offers a scalable, adaptable, and transparent solution to address future high-impact supply chain challenges, which cannot be predicted but can be anticipated through the development of optimized scenario-specific strategies.

3.2. RL-Powered Inventory Control

Whereas OML demonstrates how ML can optimize planning decisions at scale, another class of AI techniques, RL, offers a complementary approach to inventory control through adaptive learning. Specifically, RL can be seen as an ML approach in which an agent learns decision policies through trial and error, guided by rewards in a dynamic environment (Sutton and Barto 2018, Powell 2022). These ideas sit alongside a long OM/operations research tradition in inventory control, in which simple parametric policies such as base-stock (order-up-to) rules remain workhorses in practice, and learning-based methods are most compelling when the state is high-dimensional, constraints are complex, or the environment is nonstationary. Modern RL techniques appeal to inventory managers because they can handle high-dimensional states and actions, automatically extract features, and adapt to changing distributions of demand and supply. However, these methods face practical barriers: training requires vast amounts of data and extensive simulation, reward shaping must reflect business objectives, and model complexity can hinder interpretability and trust.

We expect deep reinforcement learning (DRL) to play a complementary role in inventory control: automating stable, repetitive replenishment tasks and leaving oversight, exception handling, and strategic trade-offs to human planners (Boute and Van Mieghem 2021). Today and in the near future, successful DRL implementation still requires human expertise in supply chain management and ML with the skills to diagnose convergence issues, guide retraining, and ensure performance under change. Generative AI can further reduce this expertise requirement by automating tuning, explaining policies, and using agentic reasoning with targeted tools to solve complex supply chain problems. With the right blend of human judgment and AI capabilities, DRL can become a trusted partner in building faster, smarter, and more resilient supply chains.

3.3. AI-Powered Automation

Beyond algorithmic decision making for inventory and planning, AI is also transforming the physical execution of supply chain operations through intelligent automation and by embedding intelligence directly

into the flow of goods and services. As supply chains face demand volatility, labor shortages, and disruptions, fulfillment systems must become more adaptive, decentralized, and responsive. Robotics, IoT, and additive manufacturing are no longer auxiliary tools; they are execution agents that learn from data, adjust in real time, and act without waiting for manual intervention.

Robotic systems in modern warehouses already demonstrate this shift. At e-retailers such as Amazon and JD.com, fleets of mobile robots adjust routes and tasks on the fly to match changing order volumes or machine availability (Dresser 2025). These systems operationalize upstream forecasts and capacity plans into continuous, coordinated actions. BMW's humanoid collaborative robots (cobots) and Walmart's generative AI tools illustrate how execution agility emerges when automation is directly linked to planning intelligence (Song 2026).

IoT provides the sensing backbone for adaptive control. Continuous streams of location, condition, and usage data allow fulfillment networks to preempt delays, reroute flows, and synchronize across sites. AI analytics applied to these signals enable predictive maintenance, dynamic replenishment, and condition-based inventory strategies that replace static rules with context-aware, adaptive responses. This is a modern instantiation of a long OM tradition: state-dependent control policies grounded in Markov decision processes and stochastic dynamic programming. What is new in the AI era is less the logic of adapting actions to states, but more the ability to infer richer state variables from high-frequency sensor data and to update decision rules at scale as conditions evolve at a fast pace.

Additive manufacturing further enhances responsiveness by supporting localized, on-demand production of critical parts. When combined with AI-based diagnostics, 3D printing systems can trigger repairs automatically and reduce downtime, especially in distributed service networks in which lead time is costly. This evolution does not eliminate human involvement but redefines it. Automation relieves workers of repetitive tasks, enabling them to oversee intelligent systems, resolve anomalies, and coordinate across automated nodes. As execution autonomy grows, firms must prioritize transparency, governance, and human-machine trust.

For OM researchers, these developments open fertile ground. Studying how intelligent systems operate under uncertainty, balance responsiveness with control, and integrate human oversight will require a new set of models and methods. Embedding ML, simulation, digital experimentation, and behavioral analysis into the OM toolkit is essential to understanding and guiding the next generation of execution systems.

4. The Strategy Layer: AI and Supply Chain Risk and Resilience

The strategy layer examines how AI influences the design and resilience of supply chains. It introduces the possibility of structural shifts driven by technology, assesses how AI can improve resilience while also introducing new risks, and discusses the organizational and technical frictions that accompany AI deployment. This opening sets the context for evaluating the balance between opportunities and challenges in AI-driven supply chains.

4.1. Strategic Impact and Structural Shifts

The capabilities described thus far primarily enhance performance within existing supply chain structures. A natural question, then, is whether AI can fundamentally reshape how supply chains are designed and managed rather than merely improving their operation at the margin. By definition, genuine technology-enabled structural shifts are rare: they occur only when a fundamental constraint is relaxed—when an activity that was previously prohibitively costly or technologically infeasible becomes feasible—and when the associated gains from doing so are large. Paved roads and standardized shipping containers are canonical examples of technologies that induced such structural change in supply chains. To date, AI has arguably not done so. However, it has the potential to do so, and that potential may materialize in ways that are neither immediate nor easily anticipated.

A widely recognized channel through which AI may affect supply chains is its potential to improve forecasts of where and when products are needed. However, optimism about this channel may be misplaced for two reasons. First, the bulk of supply chain costs arise from the labor and energy required to store and transport goods rather than from forecast errors. Second, physical constraints limit the benefits of incremental gains in forecast accuracy: for example, whether predicted demand is 63.5 units or 65.3 units is of little consequence when products are shipped in cases of 24 units.

AI-enabled autonomous vehicles hold immense potential. Labor represents one of the most significant and constrained resources in supply chains: not only costly but also limited by physical endurance and scheduling restrictions. Advances in autonomous mobility could, therefore, reshape where and how often key supply chain activities occur, fundamentally altering the geography and cadence of operational tasks. At the same time, any large-scale redesign driven by autonomous mobility will be shaped by regulatory and societal constraints—particularly around safety standards, labor policy, and public acceptance—which may slow or redirect adoption. These constraints point to the governance considerations explored in Section 8.

More speculatively, AI may change how supply chains are managed. In practice, a relatively small set of specialists designs and maintains the planning and execution systems that standardize many recurring decisions across the supply chain even though the network itself often evolves through market interactions, partnerships, and sourcing relationships. This centralized design of decision rules and systems makes it hard to reflect the full variety of local circumstances. In the same spirit of using LLMs to democratize supply chain technology (see Section 2.3), one can imagine pushing some decisions “back to the field” in a structured way so that context-specific choices are made closer to where information is generated. For obvious cost and scalability reasons, those decision makers cannot be human in every case, but AI agents could potentially serve this role under appropriate governance and safeguards. Could robust, dynamic, AI-enabled, more decentralized management structures become part of the future of supply chains?

4.2. Balancing AI’s Resilience Benefits and Emerging Risks in SCM

Whereas the previous section explored AI’s potential for structural transformation, a more immediate concern is how AI affects supply chain resilience and risk. AI has significantly enhanced supply chain resilience by improving sensing capabilities and accelerating response times (Cohen and Tang 2024). Technologies such as predictive analytics, ML, and real-time monitoring enable firms to anticipate disruptions and automate mitigation strategies. For example, Resilinc, a California-based software company specializing in supply chain risk management, developed a purchase order delay prediction model that allows firms to proactively reroute carriers and minimize supply chain disruptions.

These AI-enabled sensing and response capabilities are valuable across the spectrum of supply chain risks, namely, supply, product, production, and logistics. Predictive analytics can identify supplier anomalies (supply risks), ML detects product quality issues (product risks), IoT sensors monitor equipment health (production risks), and blockchain enhances traceability in transportation (logistics risks). Companies such as IBM and DHL use AI to optimize inventory and warehouse operations, respectively, showcasing the breadth of AI-enabled risk mitigation. These sensing and response capabilities draw on a mix of tools; some are AI-based and others are not. IoT and blockchain expand visibility and traceability, whereas AI methods (e.g., ML risk scoring coupled with optimization) can help translate these data streams into proactive warnings and better mitigation decisions.

However, AI adoption introduces new risks. Algorithmic bias can distort decisions when training data are incomplete or skewed. As AI systems become

more interconnected, the potential for cybersecurity vulnerabilities grows, increasing the risk of unauthorized access to sensitive data. Excessive automation may also reduce human oversight, raising concerns about accountability and transparency. For example, whereas AI-driven fraud detection systems, such as those deployed in large consumer goods firms, have proven effective at identifying anomalies and reducing losses, they must be monitored to minimize false positives and ensure fairness.

To balance these benefits and risks, firms should integrate AI with traditional risk management strategies, such as maintaining backup capacity and diversifying suppliers. Addressing challenges such as data quality, system integration, and organizational silos is also essential. Failed implementations because of poor data or incompatible systems highlight the need for robust infrastructure and cross-functional collaboration.

Ultimately, whereas AI strengthens supply chain resilience, its deployment must be guided by ethical safeguards and strategic oversight to ensure sustainable and trustworthy risk management.

4.3. Deployment Frictions and Design Safeguards

The path from AI's risk management potential to realized value is far from automatic. The story of AI in supply chains is as much about friction as it is about promise. Start-ups such as Recurrency, Hexight, Cinch, and The Rounds highlight how quickly AI can deliver measurable improvements, including shrinking inventory buffers, cutting logistics costs, or reducing waste. Yet scaling these gains is rarely straightforward. Integrating modern AI systems with legacy enterprise resource planning software, siloed data, and workflows built on "tribal knowledge" creates persistent bottlenecks. For many firms, the hardest part is not training the algorithm but stitching it into a complex operational fabric.

Beyond technical challenges, organizational frictions loom large. Employees accustomed to manual decision making often see AI as opaque or threatening, creating resistance that slows adoption. Change management is not a side task but central to unlocking value. In one example, firms using AI agents to automate quoting and order entry discovered that success required rethinking roles, retraining staff, and building trust in the system's recommendations. Without these adjustments, pilots stall in "proof-of-concept purgatory," never delivering enterprise-wide impact. Some of these challenges apply to technology adoption broadly, but AI adds a distinct operational burden. Models can degrade silently as data distributions shift, so adoption requires ongoing monitoring, retraining, and clear ownership of model risk. In practice, the hardest governance questions are often not "can we deploy?" but

"who is accountable when recommendations drift, and how quickly can we detect and correct it?"

There is also the temptation to over-automate. Automated forecasting, replenishment, and routing work well in steady states but falter when unexpected disruptions hit. The COVID-19 pandemic made it clear that "black swan" events expose the brittleness of systems that lack human judgment. When models trained on historical patterns face conditions outside their data, they risk amplifying errors rather than mitigating them.

Designing and incorporating safeguards helps transform these frictions into resilience. Incremental rollouts allow firms to stress-test AI alongside existing processes before replacing them. Human-in-the-loop systems ensure that managers can override or reinterpret recommendations when anomalies arise. Transparent models and clear communication foster organizational buy-in.

Ultimately, AI's role is not to replace people in supply chains but to augment them and make them more productive and efficient. Companies that pair algorithmic speed with human adaptability will not only scale faster but will also withstand disruption with greater confidence.

5. The Human Layer: Human–AI Collaboration

The human layer explores the interplay between humans and AI in supply chains. It considers how cognitive complementarity and collaboration can elevate decision making, examines the dynamics of human–AI interactions within organizations and between supply chain partners, and uses workforce scheduling as a concrete example. This introduction frames the importance of human insight and oversight in AI-enabled operations.

5.1. Human–AI Cognitive Complementarity

AI's value ultimately depends on how humans and algorithms work together. AI is increasingly deployed in high-value service domains such as healthcare, finance, and law, in which decision errors carry significant costs and consequences. In these settings, AI's greatest potential lies not in replacing human expertise but in complementing it. Machines excel at detecting statistical patterns in data, whereas humans contribute cognitive flexibility, contextual insight, and the ability to reframe problems in new or uncertain environments (Cukier et al. 2021).

Yet incorporating AI into these services raises challenges that research is only beginning to uncover. Analytical models grounded in rational inattention show that, although algorithmic input enhances accuracy, it can also shift the distribution of errors and increase cognitive burden under certain conditions

(Boyaci et al. 2024). In medical diagnostics, for example, AI may reduce false negatives but also increase false positives (Stoffel et al. 2018), generating inefficiencies in downstream processes. In congested systems, in which human agents operate under time pressure, additional capacity may be required to realize AI's benefits, sometimes rendering adoption unprofitable (Legros and de Véricourt 2025).

Opacity is another defining feature: high-performing AI systems rarely explain their reasoning. This raises the question of how much decision makers should trust machine outputs. Whereas models often perform well in training and testing, they frequently struggle with empirical generalizability in practical settings (Lebovitz et al. 2022). Faced with such uncertainty, humans may under-rely on AI, overriding it unnecessarily (Dietvorst et al. 2018), or over-rely, deferring to it without sufficient scrutiny (Vasconcelos et al. 2023).

Over time, service providers may learn how well AI performs (and when) and reduce such misjudgments. In high-stakes contexts, however, verification bias constrains learning as reliability can often only be assessed when follow-up actions occur, such as a biopsy after a suspected tumor (Pepe 2003). This bias may lead providers to adopt underperforming systems or reject accurate ones, sometimes arbitrarily (de Véricourt and Gurkan 2023).

Explainable AI is often proposed as a remedy, but its effects are mixed, and it is still an active research area. Low-quality explanations reduce cognitive burden without improving accuracy, whereas high-quality ones can mitigate overreliance and improve accuracy though sometimes at the cost of greater effort and lower efficiency (Boyaci et al. 2025). Finally, managerial oversight can itself foster overreliance: experiments show that agents are more likely to be blamed by managers for rejecting AI even when correct, incentivizing overreliance that undermines the complementarity human–AI systems are meant to achieve (Nikpayam et al. 2024).

5.2. Human–AI Interactions in AI-Enabled Supply Chains

AI-enabled supply chains demand more than technological sophistication; they require effective human–AI interactions in which humans and algorithms coadapt, share decision authority, and reshape supply chain operations. These interactions can be internally focused (employees) or externally focused (customers, suppliers). As firms advance AI capabilities, reskilling is critical: humans must learn to interpret AI outputs, manage autonomous systems, and adapt to evolving tools in copilot mode. Humans contribute oversight, contextual judgment, and training, whereas AI delivers speed, predictive insight, and consistency. As Swaminathan

et al. (1998) anticipates, the gains may stem from both augmentation and substitution of activities.

For internally focused systems, key challenges include training and capacity development. Whereas less-experienced employees may benefit from real-time AI guidance, evidence also indicates that more experienced workers can realize equal or even greater gains when AI complements their existing expertise (Knight et al. 2026). Without structured development, such gains risk plateauing into dependency rather than deeper expertise (Brynjolfsson et al. 2025). Continuous feedback loops are especially promising: real-time feedback improves worker decision making, supplying data that enables faster model retraining. Over time, feedback shifts from corrective to developmental, compounding both human and algorithmic learning. Yet careful design is essential because over-monitoring raises privacy concerns, and opaque algorithms undermine trust.

Externally focused systems raise cultural challenges as customers and suppliers interact directly with AI interfaces. Anthropomorphism and identity disclosure can improve transparency and engagement (Luo et al. 2019, Xu et al. 2026). Perceptions of fairness in algorithmic assignment or evaluation strongly affect morale and retention (Bai et al. 2022). Cultural alignment, thus, determines whether AI is embraced as a partner or resisted as a threat.

In contemporary agentic AI systems, it has become increasingly common to pair AI agents that act in the world with separate AI judges that evaluate, critique, and guide the decisions of these agents (Zhuge et al. 2024). In these multiagent systems, task-oriented agents autonomously plan and execute actions, whereas judge models provide oversight, such as scoring solution quality, detecting safety violations, or arbitrating between competing agent proposals. This division of labor aims to improve reliability and alignment at scale because automated judges can be run far more frequently than human reviewers while still enforcing explicit criteria or policies. Ultimately, future systems will likely include AI agents, AI judges, and expert humans who collaborate to maximize efficiency and accuracy.

In summary, AI-enabled supply chains represent not a technological upgrade but a sociotechnical transformation. Success depends on balancing augmentation with autonomy, embedding transparency and fairness, and cultivating trust through everyday human–AI interactions. More research is needed on this important, constantly evolving topic.

5.3. Workforce Scheduling Algorithms and the Potential Role of AI

The principles of human–AI interaction become concrete when we examine specific applications. Retail labor scheduling represents a large-scale stochastic optimization problem with multiple constraints and

competing objectives. The optimization problem is stochastic because demand at the hourly level is hard to predict, and labor supply, which is typically characterized by high turnover, absenteeism, and tardiness, is also uncertain (Kwon and Raman 2023, Apaolaza et al. 2025). Feasible schedules have to satisfy numerous operational (e.g., when individual workers are available to work) and regulatory constraints (Van den Bergh et al. 2013, Pinedo 2016). Retailers also have to trade off multiple objectives, such as cost, customer service, and perceived employee fairness in generating optimal schedules.

Most large retailers use workforce management (WFM) systems to schedule and manage their labor. The global WFM market, valued at US\$8–9 billion annually, includes established providers such as UKG, Zebra Reflexis, and Ceridian Dayforce. WFM systems help retailers plan their labor schedule; track attendance; and generate reports for payroll processing, regulatory compliance, and performance analytics (Grand View Research 2023, Marian and Xing 2024, Xing et al. 2024).

AI matters here because it can substantially improve key inputs to these optimization cores (e.g., finer-grained demand and absence predictions) and can lower the human–computation interface cost of building, explaining, and adjusting schedules. Empirical research examining large-scale retail scheduling data reveals the power of combining human judgment with algorithms. Store managers frequently override schedules generated by algorithms in WFM systems with evidence suggesting these interventions are associated with improved labor productivity (Kwon et al. 2024, 2025). Managers leverage private information on local conditions and individual employee capabilities that many currently implemented algorithms cannot access, particularly for newly hired workers. They also expand the feasible scheduling solution space by distinguishing between hard and soft constraints with many final schedules violating system-recorded availability constraints.

The empirical findings underscore how AI can be used to improve retailer performance, customer experience, and employee equity. Recognizing this opportunity, many WFM systems have now incorporated ML approaches in their forecasting algorithms. AI-powered WFM tools can also analyze historical attendance patterns to predict absence and tardiness risk. Natural language interfaces for employee scheduling can help capture employee preferences and better communicate and explain generated schedules.

6. The Infrastructure Layer: Supply Chains Behind AI

The infrastructure layer addresses the physical foundations that enable AI adoption, such as semiconductors,

data centers, and energy grids. It explores dynamic interdependencies between AI demand and the supply chains for computing hardware and electricity, identifies material and infrastructure bottlenecks, and outlines the economic and policy stakes of scaling the infrastructure needed to support AI-driven operations.

6.1. Dynamic Interdependencies

The exponential growth in AI adoption has exposed critical interdependencies between AI compute demand and the underlying supply chains for semiconductors and electricity. These interdependencies create feedback loops that amplify volatility and risk, particularly through the bullwhip effect—by which demand shocks propagate upstream, distorting orders, capacity investments, and lead times (Lee et al. 1997, Fransoo and Udenio 2021)—and the pork cycle, characterized by cyclical fluctuations in supply and prices (Parker and Shonkwiler 2014, Peels et al. 2025).

At the core of this dynamic is the semiconductor supply chain, in which advanced AI chips (e.g., NVIDIA's Blackwell architecture) are produced in extremely complex factories (e.g., TSMC's most advanced fabs). This production requires cutting-edge lithography (e.g., ASML's Extreme UltraViolet machines) and advanced packaging. Recent lead times for semiconductor production are about six months, and capacity expansion is constrained by highly specialized tooling suppliers whose own lead times can exceed 12 months (Weise 2024, Winkler et al. 2025). When demand surges (e.g., as seen post-2022 with generative AI), chip orders spike, but supply lags, creating shortages and inflated backlogs. Once capacity catches up, the risk of overcorrection looms. This can result in excess inventory and price volatility as observed in past supply chain disruptions.

The electricity supply chain introduces another layer of complexity. Data centers, clustered near a limited number of internet hubs because of latency requirements (e.g., Northern Virginia's Data Center Alley), face grid access delays of two to four years because of transformer shortages and renewable energy integration challenges. Unlike semiconductors, energy infrastructure is decentralized and geographically constrained, making grid capacity a potential bottleneck. Hyperscalers may mitigate this by diversifying data center locations, but local energy shortages could still delay AI deployment, especially in high-demand regions.

Pork cycle dynamics further destabilize the system. Capacity investments, driven by inflated demand forecasts, often overshoot, leading to cyclical overcapacity (e.g., TSMC's 2026–2027 expansions). If AI demand growth slows (e.g., because of algorithmic efficiency gains, slower than anticipated AI adoption, or grid constraints), overcapacity could trigger price collapses, repeating the cycle.

Firms must carefully align chip and energy investments, diversify supply chains, and adopt flexible capacity planning. Policymakers should accelerate grid expansion and renewable energy integration to prevent AI's growth from stalling because of infrastructure constraints.

6.2. Material and Infrastructure Bottlenecks

These interdependencies rest on a physical foundation that is itself vulnerable. AI cannot scale without a robust, reliable supply chain. Behind every LLM or generative AI application lies a vast upstream network of raw materials, semiconductor manufacturing, hardware components, data centers, cooling systems, and energy infrastructure. This “invisible backbone” is what makes AI possible, but it is also riddled with fragilities. If supply chains falter, AI progress stalls.

Several risks dominate. At the material level, silicon, rare earth elements, and critical minerals such as cobalt and lithium are geographically concentrated and politically sensitive, making them vulnerable to export restrictions; conflict; and environmental, social, and governance-driven supply constraints. Semiconductor manufacturing adds another choke point: a handful of players, most notably TSMC and ASML, control irreplaceable stages of chip production, exposing AI growth to concentration risk, intellectual property theft, and environmental costs. On the hardware side, counterfeit components and cybersecurity vulnerabilities create operational threats. Building and running AI data centers compounds the challenge. Copper and fiber-optic cables, cooling systems, and specialized fluorinated coolants are subject to shortages and geopolitical frictions. One of the most pressing issues is energy: AI data centers already consume 1%–2% of global electricity with forecasts of up to 20% by 2030, straining grids and undermining sustainability goals. Without careful planning, power shortages, blackouts, or emissions backsliding could throttle AI adoption.

The main takeaway is stark: AI is only as strong as the supply chain that underpins it. To ensure resilience, policymakers and industry leaders must act on multiple fronts, including diversifying suppliers, enforcing antitrust to avoid choke points, embedding sustainability into design and manufacturing, and investing in renewable energy and recycling infrastructure. Life cycle planning, from capacity expansion to end-of-life hardware recovery, is essential. AI's promise cannot be realized unless its physical foundation is secured, diversified, and future-proofed.

Despite rapid progress in AI methods and widespread experimentation, a clear and consistent picture of their incremental value over traditional OM approaches remains mixed in many settings. In forecasting, for example, Bonilla et al. (2023) find that ML models do not consistently outperform classical statistical approaches

even for multivariate and nonlinear time series. This raises a basic question: when performance gains are modest, do they justify the substantially higher computational, financial, and environmental costs of deploying AI at scale?

The resource footprint is not limited to training. Large-scale inference and continuous retraining can also carry substantial ongoing costs, including energy use that depends on when and how inference is run. Petropoulos et al. (2025), for example, estimate that a large retailer's shift to ML-based forecasting can generate more than 100,000 tons of additional CO₂ emissions annually. These costs motivate OM research directions that treat compute as a constrained resource, including carbon-aware scheduling of inference (e.g., shifting workloads toward periods with greater renewable availability) and models that incorporate dynamic pricing of computational resources as energy availability varies over time.

Finally, the benefit–cost trade-off must account for reliability and the cost of errors. As AI systems become more capable yet still fail unpredictably, the cost of effective oversight can be material and, in some settings, prohibitive (Bastani and Cachon 2025, Dai and Taylor 2025). A more complete assessment, therefore, weighs not only accuracy improvements but also inference and storage costs, the risk and cost of errors, and whether incremental gains translate into meaningful operational outcomes.

6.3. Economic and Policy Stakes

The scale of these infrastructure challenges is reflected in recent investment patterns. AI computing requires an extensive ecosystem of enablers (Jacobides et al. 2021). The supply chain for AI concerns the procurement, production, and distribution activities needed to support AI computing, including importing advanced graphics processing units, manufacturing AI servers, building data centers, and generating electrical power, among others. For example, imports of computers, computer accessories, peripherals, and parts in the first seven months of 2025 were 58% higher than the same period in 2024 and 113% higher than the same period in 2023. Data center construction has also surged with the nominal value of construction of data centers built in the first seven months of 2025 being 36% higher than the same period in 2024 and 115% higher than the same period in 2023. Concomitantly, firms are importing more electrical equipment, and domestic plants are producing the most electrical equipment in two decades to meet the electricity demands of these data centers, which are straining existing infrastructure (Weise 2024).

Investment in physical enablers of AI is not unique to the United States and is also occurring in Europe (Licht and Wohlrabe 2024, Kantour 2025) and China

(Lundvall and Rikap 2022, Bloomberg Businessweek 2025). Given that widespread embedding of AI into organizational work processes is just beginning (Bonney et al. 2024), further investment in the physical AI infrastructure is likely as capital expenditures by Alphabet, Amazon, Meta, and Microsoft continue to rise (Winkler et al. 2025). This setting may create a useful strategic research event, in Merton's (1987) nomenclature, to study a wide range of operations and supply chain phenomena, including capital investment, international trade, vertical integration, and contracting, among others. For example, careful study of the relationship dynamics between Microsoft and OpenAI (Carr and Bass 2025, Soni 2025) may spur refinement of theories concerning transaction cost economics (Ketokivi and Mahoney 2020) and economic property rights (Holzhacker et al. 2025).

7. Industry and Sector Spotlights

To ground the discussion in practical contexts, this section highlights how AI is deployed across diverse examples of supply chains. From e-commerce and humanitarian relief to agriculture, healthcare, and pricing, it draws lessons from industry-specific innovations and challenges, illustrating how sectoral experiences inform broader principles for AI-enabled supply chains.

7.1. AI and E-Commerce Supply Chain

AI's impact varies significantly across industries and is shaped by both sector-specific challenges and operational contexts. E-commerce supply chains, with their data abundance and rapid feedback loops, offer valuable lessons in AI deployment. JD.com, one of China's major e-retailers, illustrates four critical success factors that extend beyond e-commerce to other domains (Lee et al. 2026).

First, organizations need the right setup. AI requires extensive data, and JD.com has built a complex system with cleansed consumer purchase and supplier replenishment data. JD.com recognizes the importance of advanced analytics and has formed a strong team of experts in data science and operations research. Management was willing to experiment with new methods and to continuously explore, test, improve, and refine, similar to the Six Sigma improvement cycle.

Second, AI needs to work in conjunction with complex optimization techniques. In short, success requires tight collaboration among computer scientists and operations researchers. Management articulates business problems and objectives, and AI-enabled workflows can help formulate optimization models, set constraints, gather relevant data, solve the optimization problem, and interpret results. Many steps that previously required significant technical support can now be streamlined with AI.

Third, AI can be used to explain sophisticated algorithms and improve managerial acceptance. At JD.com, demand forecasting requires complex statistical algorithms coupled with ML. To increase adoption, the company used AI to explain and interpret forecasts, allowing decision makers to understand the results before deployment at scale. This approach helped realize the benefits of advanced algorithms, improving adoption and trust.

Fourth, integration with physical operations is essential. Robots and autonomous vehicles execute tasks on the warehouse floor or in last-mile delivery (and sometimes even in physical stores), and they need to interact seamlessly with the planning systems. This integration ensures that AI decisions translate into real-world execution without friction.

The above example of JD.com is not unique. Similar advances and innovations are actively pursued at many other large e-commerce retailers (e.g., Alibaba, Amazon, Shopify, and Walmart); retail chains in grocery, fashion, and electronics; and many technology solution providers.

7.2. AI and Humanitarian Supply Chains

Humanitarian supply chains (HSCs) provide critical goods, services, and personnel during disasters such as earthquakes, hurricanes, wars, and pandemics. HSCs must handle unpredictable demand, time-critical delivery, supply scarcity, and complex stakeholder coordination (Ergun et al. 2023). Stockpiles of critical supplies must be strategically located and sized. Traditional OM models provide rigorous frameworks for facility location and inventory replenishment. However, slow demand updates, lead times, and uncertain supply lines challenge decision making. AI can augment conventional methods by processing large amounts of information in real time to help anticipate needs and allocate resources.

AI has been used for disaster risk assessment, early warning, and the dynamic allocation of relief resources. Recent advances show how AI can strengthen earthquake early warning and rapid detection, for example, by extracting weak signals and improving event detection from streaming seismic data. At the same time, reliable earthquake prediction—forecasting the precise time, location, and magnitude of future events—remains an open scientific challenge. We focus here on detection and early warning applications, not on claiming that earthquakes can be predicted reliably in advance; consistent with the literature (e.g., Mousavi et al. 2025), we view AI-based methods as part of ongoing efforts to improve warning systems. More broadly, as Ergun et al. (2023) discuss, learning-based models have been explored for earthquake-related forecasting tasks using historical and geological information, but their scope and reliability vary by setting and should

not be interpreted as solving earthquake prediction. During disasters, urgent demand requires flexible procurement or prepositioned stockpiles (Smalley et al. 2015), and secondary outbreaks can follow infrastructure damage (e.g., cholera after an earthquake or severe flooding); in these contexts, AI can support stockpile distribution across vulnerable subpopulations over geography and time through data-driven forecasting and real-time resource allocation.

The global COVID-19 pandemic underscored the need for a coordinated response amid uncertainty. Vaccines were in high demand, but distribution faced challenges such as supply availability, cold-chain logistics, scheduling constraints, and geographic disparities in uptake (Kim et al. 2022). To be effective, HSCs require coordination, navigating trade-offs in reach and efficacy, accounting for prioritization and fairness (e.g., most vulnerable first), and potential human behaviors (e.g., vaccine hesitancy). AI models can help address these many interconnected factors from early warning systems to interactive tools for crisis management.

Across these scenarios, lives depend on a fast, effective humanitarian response. AI can enable real-time decision making that accounts for changing conditions and system vulnerabilities, ultimately strengthening the resilience and responsiveness of HSCs.

7.3. AI and Agricultural Supply Chains

Agricultural supply chains represent a third potential context for AI deployment. Perhaps because of its foundational role in human survival, agriculture has historically been a driver of technological innovation rather than merely a beneficiary. From irrigation systems in ancient civilizations, through the first industrial revolution, to automated harvesters and biotechnology, this sector has continuously evolved in response to both necessity and opportunity. It is, therefore, unsurprising that agriculture has become an active frontier for the application of AI (Olsen 2026).

AI applications in agricultural supply chains can be understood through a value-chain lens. Upstream, at the production stage, ML—and more specifically, deep learning models supported by improved computer vision—are increasingly used alongside predictive and prescriptive analytics and robotics. These tools enable more accurate crop yield prediction; pest detection; and site-specific interventions such as precision weeding, planting, and the optimization of fertilizer, pesticide, and irrigation use. Whereas precision agriculture has been an active research area for more than 50 years, AI has significantly improved the sophistication and precision of the models used in this area.

Midstream, in postharvest handling and logistics, sensor networks and geospatial data, combined with computer vision, supervised learning, and time series forecasting, play a central role. Such technologies

support grading and quality assessment; predict and optimize storage life; and improve aggregation, routing, and transportation decisions, thereby reducing spoilage and inefficiencies in highly perishable supply chains.

Downstream, at the consumer-facing end of the chain, natural language processing, recommender systems, blockchain, and predictive analytics enhance demand forecasting and dynamic pricing, enabling greater personalization of products. These tools also strengthen traceability and transparency, linking consumer preferences back to farm-level practices and improving food safety.

Taken together, these applications show that AI is not simply about efficiency gains. By enabling more accurate yield forecasts, robust responses to weather shocks, and smarter cold-chain management, AI can build resilience and improve the sustainability of agricultural supply chains.

7.4. AI and Healthcare Delivery Supply Chains

Healthcare delivery supply chains differ from traditional product supply chains because they orchestrate patient flows, diagnostic capacity, and specialist time under stringent safety and equity constraints (Betcheva et al. 2021). AI's near-term impact is most visible in screening and triage, in which autonomous and assistive systems shift detection upstream, reconfigure referral queues, and change the mix of scarce resources (e.g., specialist minutes, imaging slots). For example, autonomous AI for diabetic eye exams increases real-world specialist clinic productivity in a cluster-randomized trial (Abramoff et al. 2023). From a system perspective, AI-enabled eye screening can be cost-effective when patient volume is sufficiently high (Ahmed et al. 2025).

Realizing these gains requires aligning incentives, redesigning care pathways, and measuring value in terms of clinical outcomes and throughput, not just model accuracy (Hunt et al. 2026). When AI acts as a gatekeeper, it can reduce low-value demand reaching specialists but risks deferring care for intermediate-risk patients if thresholds are misaligned with downstream capacity. When AI serves as a second opinion, it can improve accuracy and learning, but may add delay and cognitive burden if not properly integrated into workflow (Dai and Abramoff 2023). Choosing between these roles should depend on local congestion, follow-up compliance, and inequality of access: parameters that OM models can estimate and optimize.

The familiar supply chain coordination logic means that payment models shape uptake. Under fee for service, providers may face weak incentives to invest in AI that substitutes for billable specialist work even when the system value is positive. Under value-based arrangements, savings from avoided complications and improved follow-up can support reimbursement

for autonomous AI services and infrastructure (Abramoff et al. 2024). Designing contracts that share savings, reward timely follow-up, and penalize avoidable false negatives can align AI use with social value.

Implementation requires robust workflow engineering, not only for screening and triage but also for documentation and coding workflows. Ambient AI scribes and related tools can reduce documentation burden yet intensify coding incentives, so guardrails around clinician authorship, transparency, and payer oversight are needed to ensure that AI-enabled supply chains improve access and outcomes rather than fueling a coding arms race.

7.5. AI and Supply Chain Pricing

Pricing has long been one of the most powerful levers for influencing supply chain performance. In the AI era, pricing is no longer a downstream commercial activity isolated from supply, logistics, or operations; it has become a dynamic, data-driven coordination mechanism that directly shapes demand patterns, capacity utilization, inventory flow, and overall network resilience (Cohen 2026b). AI-driven pricing for both products and services—whether in retail, mobility, e-commerce, logistics, or subscription-based models—now interacts tightly with forecasting, replenishment, fulfillment, and resource allocation decisions.

Historically, most firms relied on simple rules such as cost-plus pricing, periodic markdown schedules, or manually updated fare tables. AI has transformed this landscape by enabling granular, personalized, context-dependent, and continuously updated price recommendations. ML models synthesize large volumes of transactional data, customer attributes, market conditions, competitor actions, and supply chain constraints to prescribe prices that reflect real-time conditions. In summary, firms increasingly move beyond static, rule-based approaches toward dynamic, data-driven, real-time adaptive strategies, allowing prices to adjust within seconds when demand surges, inventory becomes scarce, or fulfillment capacity tightens. This paradigm shift fundamentally alters how supply chains operate. Prices become signals that balance the system: higher prices during scarcity throttle demand and protect service levels, whereas lower prices help clear excess inventory, smooth demand across periods, or utilize slack capacity.

AI-enabled pricing is most powerful when tightly integrated with the intelligence and execution layers of supply chains. Decision-focused learning, RL, and digital twins described earlier in this article provide the structural backbone needed for pricing models to translate predictions into operationally feasible decisions. For example, inventory-aware pricing adjusts discounts and markdowns based not only on demand forecasts but also on replenishment lead times, service-level commitments, and capacity constraints. A second

example is capacity- and congestion-aware pricing (e.g., surge pricing in ride-hailing or peak-load pricing in parcel networks) that uses prices to shape demand to available operational capacity. Finally, a third example is end-to-end decision engines that combine demand signals and supply chain constraints to synchronize pricing with replenishment, routing, staffing, and assortment decisions. In this sense, pricing acts as a real-time interface between the market and the supply chain, converting operational bottlenecks into optimized economic outcomes rather than service failures.

Despite its potential, AI-enabled pricing creates several risks that need to be carefully mitigated. First, algorithmic tacit collusion can emerge when independently trained pricing algorithms learn parallel high-price strategies even without explicit coordination, creating challenges for antitrust oversight. Second, fairness and privacy concerns arise because AI-based price personalization can increase efficiency but may inadvertently create discriminatory outcomes or erode customer trust, especially when based on correlated or sensitive attributes. Third, complex opaque models may become difficult for practitioners to interpret, increasing the need for transparency, governance, and human oversight. Fourth, overreactive price updates may amplify demand variability and volatility, intensifying bullwhip effects if not coordinated with inventory planning.

Generative AI will further transform supply chain pricing by making advanced pricing capabilities accessible to nontechnical users (Cohen 2026a). LLM-based assistants can help planners test pricing scenarios, examine elasticity patterns, or evaluate promotion impacts via natural language prompts. As these tools integrate with optimization engines and digital twins, pricing agents will increasingly automate routine decisions as humans guide strategic trade-offs. Generative AI can also increase access to and reduce technical barriers for market research (Brand et al. 2023), data augmentation, and data scraping (Cohen and Hage-Youssef 2025), all of which can support pricing decisions.

AI-driven pricing has become a central mechanism for coordinating supply and demand across complex networks. Realizing its full value requires not only robust models but also responsible governance, transparency, and alignment with operational constraints and societal expectations.

8. Readiness, Implementation, and Governance

We believe that AI will be transformative in global supply chains only if its rollout is grounded in pragmatism, long-term commitment, and deep organizational change rather than hype-driven, plug-and-play promises. Despite the significant opportunities that AI

brings, companies face two persistent categories of bottlenecks that we refer to—metaphorically—as technical debt and organizational debt. By “debt,” we mean accumulated gaps that build up when new systems are adopted faster than the underlying capabilities needed to support them. Technical debt includes outdated infrastructure, legacy information technology (IT) architectures, inconsistent or dirty data, and fragmented systems. Organizational debt reflects skills gaps, siloed processes, misaligned incentives, and cultural resistance. These bottlenecks interact: technical limitations impede effective use of AI tools, whereas organizational frictions prevent the adoption and scaling of even well-designed systems. Without addressing both, AI pilots often remain isolated experiments rather than scaling across the value chain (Dai and Tayur 2022).

The path forward requires addressing structural, cultural, and infrastructural barriers in parallel. This includes cleaning, structuring, and standardizing data; upgrading fragmented IT and legacy systems; breaking down organizational silos; and aligning incentives both within and across enterprises. Many valuable AI supply chain use cases depend on cross-enterprise data sharing, which requires trust, incentives, and collaborative agreements between partners. With these foundational challenges in mind, success begins with aligning AI initiatives to core supply chain priorities such as resilience, automation, sustainability, and cost efficiency. A clear road map anchored in leadership sponsorship and enterprise-wide resource commitments prevents AI efforts from fizzling out as one-off pilots.

However, strategic alignment alone is not sufficient. AI is only as strong as the data and systems upon which it relies. Predictive AI depends on consistent, high-quality inputs and a stable business environment. Generative AI may hallucinate, and agentic AI remains far from capable of autonomous high-stakes decisions in messy, multiparty supply chains. To mitigate risks, companies must upgrade core platforms to cloud-based solutions, deploy application programming interfaces, and invest in modern digital architectures that enable better data access and integration of data-driven insights. At the same time, instilling data governance practices—such as strong data quality, master data management, and pervasive stewardship—helps ensure that predictive and prescriptive models operate effectively across the entire supply chain.

AI transformation is not just a technical challenge but also a human one. Frontline workers, from planners to logistics managers, must understand, trust, and embed AI tools into daily routines. Training, cross-functional collaboration, and retraining programs help build adoption, reduce skepticism, and reinforce AI as

an augmentation tool rather than a replacement threat. Building on this human-centered approach, companies must also sustain internal and external trust by safeguarding against algorithmic bias, fairness risks, and unintended consequences. Regular ethics audits ensure accountability, explainability, and transparency in how AI models influence supply chain decisions.

Ultimately, scaling AI is a cultural transformation. Leaders must reduce siloed decision making, create incentives that reward collaboration, and address skepticism directly. Communicating a clear vision, securing buy-in across all management levels, and involving frontline employees builds the trust and alignment necessary for sustained adoption. AI in supply chains should begin with targeted, complementary applications, not full autonomy. Early focus areas include demand forecasting, documentation, visibility enhancements, and repetitive operational tasks. From there, organizations can adopt a phased rollout that starts by laying a robust digital foundation of integrated, reliable data systems. Next, they must restructure decision-making processes to reduce politicization and silos and then retrain frontline workers and incentivize adoption. Finally, they should pilot in contained contexts, measure return on investment, and scale gradually.

AI’s transformative potential in global supply chains will not come from quick-fix automation but from disciplined execution, cultural change, and sustained leadership commitment. Companies that modernize their data and systems, invest in talent, maintain ethical guardrails, and manage change deliberately will unlock AI’s long-term ability to reshape how supply chains function, pragmatically and powerfully, over the long arc of time.

9. Conclusion

The convergence of AI and SCM marks a fundamental reimagining of how goods and services flow through our global economy. Achieving this transformation requires balance. Throughout this article, a central theme emerges: AI delivers the greatest value when it operates within the structures that OM provides. Forecasting algorithms become meaningful only when embedded in decision-focused frameworks, and warehouse robots enhance productivity when integrated into well-designed fulfillment systems. Success stems from pairing AI’s capacity to learn and adapt with OM’s structural understanding of how supply chains function in practice. A related theme is that progress depends on addressing both the technical debt embedded in legacy systems and the organizational debt reflected in skills, incentives, and processes. These bottlenecks are distinct yet mutually reinforcing, and overcoming them together is essential to realizing AI’s full impact on supply chains.

This perspective builds on decades of OM work in statistical and time-series forecasting, judgmental forecasting, and planning models. ML extends and solidifies that foundation when new data sources become available, nonlinear patterns emerge, or granular heterogeneity materially changes decision quality, but the core lesson remains decision-first rather than prediction-first.

Even if AI sharply reduces routine supply chain frictions, including forecast errors, coordination failures, and much of day-to-day operational uncertainty, the field's frontier would shift from better optimization to institutional and organizational design. "Be better" does not mean cost and speed alone but, rather, improving a vector of objectives (e.g., service, resilience, sustainability, and fairness) under explicit constraints. And even short of that limiting case, physical supply chains remain far from fully autonomous because many critical aspects of the operational state are not reliably captured in data, and any agentic system can only act on what is observed and measured. The central questions then become governance questions: who chooses the objective function; who is accountable when an autonomous policy fails; and how can firms audit, bound, and contest algorithmic decisions that are increasingly opaque, adaptive, and distributed across organizational and national borders? In a global supply chain, this is not a technical footnote but a core design challenge: aligning AI decision tools with heterogeneous regulatory regimes, cultural expectations, and risk tolerances, preventing bias, gaming, and brittle reliance on data patterns that may not survive regime shifts.

At the same time, the binding constraint becomes organizational and economic rather than computational. When AI substitutes for routine planning and coordination, it reallocates tasks, decision rights, and surplus, reshaping roles, wages, career ladders, and ultimately product demand, and this then feeds back into sourcing choices, contracting, and supplier stability; the efficiency dividend is real, but its consequences depend on incentive alignment and how quickly workers transition into complementary roles that require judgment, negotiation, and system stewardship. Nor does near-perfect coordination in normal times buy robustness in abnormal times: tail risks, cascading failures, cyber sabotage, and geopolitics turn resilience into a first order design problem, demanding architectures that preserve optionality, maintain meaningful human override, and remain stable under adversarial pressures and rare shocks. In this sense, AI does not end OM/SCM; it elevates it, redirecting the field's next big contributions toward (i) governing autonomous systems, (ii) redesigning organizations and workforce transitions so that gains are broadly realized, and (iii) building resilience by design when

the world refuses to be deterministic and when we need to be ready for future black swan events.

A call to researchers: The research community faces both exciting opportunities and serious responsibilities. We need new models that capture the human–AI collaboration dynamic. The field needs rigorous studies of deployment frictions, organizational change, and long-term societal impacts of AI-enabled supply chains. Critically, interdisciplinary collaboration is essential. Computer scientists must partner with OM researchers to build decision systems that are both technically robust and operationally sound. Behavioral scientists must help us understand how workers interact with AI tools, and ethicists must guide us toward fair and transparent implementations.

A call to practitioners: Industry leaders need to start pragmatically and scale thoughtfully. They need to resist the temptation of hype-driven deployments. AI will not fix broken processes or poor data. Instead, they should view AI as a long-term capability that requires sustained commitment, continuous learning, and a willingness to adapt. The companies that thrive are those that combine AI with robust operations, invest in both technology and their people, and innovate, maintaining ethical principles.

A call to educators: Academic institutions must prepare the next generation for a world in which supply chain professionals work alongside intelligent systems. Curricula should integrate AI literacy with traditional OM foundations, teaching students not just how algorithms work but when to use them, how to interpret their outputs, and when to override their recommendations. We need graduates who can connect technical expertise with practical business execution, apply optimization theory, and drive organizational change.

Moving forward, several critical challenges require attention. The infrastructure underlying AI itself (semi-conductors, data centers, energy grids, etc.) faces bottlenecks that could constrain adoption. Additionally, the environmental footprint of AI systems, such as energy consumption and electronic waste production, must be addressed if AI-enabled supply chains are to be truly sustainable. At the same time, opportunities abound. Emerging technologies such as LLMs and multi-AI agent systems promise to democratize access to a wide variety of optimization and analytics tools, allowing smaller firms to benefit from capabilities once reserved to large enterprises.

Beyond developing ways to reduce AI's energy consumption (as discussed before), it is crucial to consider when deploying computationally intensive AI is truly warranted. In our concluding outlook, we note that the marginal benefits of AI should be measured against its marginal costs. Not every supply chain problem justifies a resource-hungry AI solution if a leaner analytic approach yields similar outcomes. By

acknowledging this trade-off, we underscore the importance of using AI judiciously, namely, applying it when it adds clear value and pursuing simpler solutions when they suffice as part of a sustainable and responsible supply chain strategy.

This vision statement emerges from extensive collaboration across research, industry, and technology communities. Yet it represents a beginning, not an endpoint. The transformation of supply chains through AI will unfold over years, bringing challenges and opportunities that we cannot yet foresee. Its success will depend on how we—researchers, practitioners, and educators—work together to ensure that AI’s power is guided by operational wisdom, ethical responsibility, and human purpose. The future of SCM will be shaped not merely by those who deploy AI, but by those who understand its limits and design for its responsible use. The OM community has both the insight and the obligation to lead that transformation, building supply chains that are not only more efficient and resilient but also more intelligent, equitable, and sustainable.

Acknowledgments

The authors are grateful for the constructive and thoughtful feedback from the department editor, Professor Wedad Elmaghraby, and an anonymous associate editor; from speakers and attendees of the Second Washington Operations Workshop: AI in Supply Chains; and from Saif Benjaafar, Warren Powell, and Mohan Sodhi.

References

- Abramoff MD, Dai T, Zou J (2024) Scaling adoption of medical artificial intelligence: Reimbursement from value-based care and fee-for-service perspectives. *NEJM AI* 1(5):A1pc2400083.
- Abramoff MD, Whitestone N, Patnaik JL, Rich E, Ahmed M, Husain L, Hassan MY, et al. (2023) Autonomous artificial intelligence increases real-world specialist clinic productivity in a cluster-randomized trial. *NPJ Digital Medicine* 6(1):184.
- Adulyasak Y, Benomar O, Chouaoui A, Cohen MC, Khern-Am-Nuai W (2024) Using AI to detect panic buying and improve products distribution amid pandemic. *AI Soc.* 39(4):2099–2128.
- Agrawal N, Cohen MA, Deshpande R, Deshpande V (2024) How machine learning will transform supply chain management. *Harvard Bus. Rev.* 103(3–4):128–137.
- Ahmed M, Dai T, Channa R, Abramoff MD, Lehmann HP, Wolf RM (2025) Cost-effectiveness of AI for pediatric diabetic eye exams from a health system perspective. *NPJ Digital Medicine* 8(1):3.
- Amos B, Kolter JZ (2017) OptNet: Differentiable optimization as a layer in neural networks. *Proc. 34th Internat. Conf. Machine Learning*, vol. 70 (JMLR.org), 136–145.
- Apaolaza B, Gallino S, Kwon C (2025) What makes scheduling “responsible”? Evidence from 280 million shifts across 20 retailers. Working paper, The Wharton School, University of Pennsylvania, Philadelphia.
- Bai B, Dai H, Zhang DJ, Zhang F, Hu H (2022) The impacts of algorithmic work assignment on fairness perceptions and productivity: Evidence from field experiments. *Manufacturing Service Oper. Management* 24(6):3060–3078.
- Bastani H, Cachon GP (2025) The human-AI contracting paradox. Preprint, submitted December 29, <https://doi.org/10.2139/ssrn.5962739>.
- Betcheva L, Erhun F, Jiang H (2021) OM forum—Supply chain thinking in healthcare: Lessons and outlooks. *Manufacturing Service Oper. Management* 23(6):1333–1353.
- Bloomberg Businessweek (2025) DeepSeek’s “tech madman” founder is threatening us dominance in AI race. Online (May 13), <https://www.bloomberg.com/news/features/2025-05-13/deepseek-races-after-chatgpt-as-china-s-ai-industry-soars>.
- Bonilla ME Jr, McDonald J, Toth T, Sadler B (2023) Traditional vs machine learning approaches: A comparison of time series modeling methods. *SMU Data Sci. Rev.* 7(2):2.
- Bonney K, Breaux C, Buffington C, Dinlersoz E, Foster LS, Goldschlag N, Savage K (2024) Tracking firm use of AI in real time: A snapshot from the business trends and outlook survey. NBER Working Paper No. 32319, National Bureau of Economic Research, Cambridge, MA.
- Boute RL, Van Mieghem JA (2021) Digital operations: Autonomous automation and the smart execution of work. *Management Bus. Rev.* 1(1):177–186.
- Boyaci T, Canyakmaz C, de Véricourt F (2024) Human and machine: The impact of machine input on decision making under cognitive limitations. *Management Sci.* 70(2):1258–1275.
- Boyaci T, de Véricourt F, Canyakmaz C (2025) Beyond the black box: Unraveling the role of explainability in human–AI collaboration. Preprint, submitted January 28, <https://doi.org/10.2139/ssrn.5045415>.
- Brand J, Israeli A, Ngwe D (2023) Using LLMs for market research. Working paper, Harvard Business School Marketing Unit, Boston.
- Brynjolfsson E, Li D, Raymond LR (2025) Generative AI at work. *Quart. J. Econom.* 140(2):889–942.
- Campbell M, Hoane AJ, Hsu F (2002) Deep blue. *Artificial Intelligence* 134(1–2):57–83.
- Carr A, Bass D (2025) Microsoft’s CEO on how AI will remake every company, including his. *Bloomberg Online* (May 5), <https://www.bloomberg.com/news/features/2025-05-15/microsoft-ceo-satyarnadella-on-his-ai-efforts-and-openai-partnership>.
- Cohen MA (2024) Supply chain research: Overview of joint work. Tang CS, ed. *Responsible and Sustainable Operations: The New Frontier* (Springer Nature, Cham, Switzerland), 3–8.
- Cohen MC (2026a) How to use generative AI for pricing. *MIT Sloan Management Rev.* (January 29), <https://sloanreview.mit.edu/article/how-to-use-generative-ai-for-pricing/>.
- Cohen MC (2026) *Pricing in the Age of AI* (MIT Press, Cambridge, MA).
- Cohen MC, Dai T, eds. (2026) *AI in Supply Chains: Perspectives from Global Thought Leaders*, Springer Series in Supply Chain Management, vol. 27 (Springer, Cham, Switzerland).
- Cohen MC, Hage-Youssef E (2025) Generative AI for data scraping. Preprint, submitted July 23, <https://doi.org/10.2139/ssrn.5353923>.
- Cohen MC, Tang C (2024) The role of AI in developing resilient supply chains. *Georgetown J. Internat. Affairs* 25(1).
- Cohen MA, Agrawal N, Deshpande V (2026) Reimagining supply chain planning using machine learning: A roadmap to agility and resilience. Cohen MC, Dai T, eds. *AI in Supply Chains: Perspectives from Global Thought Leaders*, Springer Series in Supply Chain Management, vol. 27 (Springer, Cham, Switzerland), 149–166.
- Cohen MC, Gras PE, Pentecoste A, Zhang R (2022) *Demand Prediction in Retail: A Practical Guide to Leverage Data and Predictive Analytics* (Springer, Cham, Switzerland).
- Cohen MC, Leung NHZ, Panchamgam K, Perakis G, Smith A (2017) The impact of linear optimization on promotion planning. *Oper. Res.* 65(2):446–468.
- Cristian R, Harsha P, Perakis G, Quanz B (2025) Efficient end-to-end learning for decision-making: A meta-optimization approach. Working paper, Massachusetts Institute of Technology, Cambridge, MA.

- Cristian R, Harsha P, Perakis G, Quanz BL, Spantidakis I (2023) End-to-end learning for optimization via constraint-enforcing approximators. Working paper, Massachusetts Institute of Technology, Cambridge, MA.
- Cukier K, Mayer-Schönberger V, de Véricourt F (2021) *Framers: Human Advantage in an Age of Technology and Turmoil* (Dutton, New York).
- Dai T, Abramoff MD (2023) Incorporating artificial intelligence into healthcare workflows: Models and insights. *INFORMS TutORials in Operations Research* (INFORMS, Catonsville, MD), 133–155.
- Dai T, Swaminathan JM (2026) Artificial intelligence and operations: A foundational framework of emerging research and practice. *Production Oper. Management* Forthcoming.
- Dai T, Taylor T (2025) Designing enterprise AI systems: Hallucination, creativity, and moral hazard, Preprint, submitted January 5, <https://doi.org/10.2139/ssrn.5996714>.
- Dai T, Tayur S (2022) Designing AI-augmented healthcare delivery systems for physician buy-in and patient acceptance. *Production Oper. Management* 31(12):4443–4451.
- de Véricourt F, Gurkan H (2023) Is your machine better than you? You may never know. *Management Sci.* 72(1):558–574.
- Dietvorst BJ, Simmons JP, Massey C (2018) Overcoming algorithm aversion: People will use algorithms if they can (even slightly) modify them. *Management Sci.* 64(3):1155–1170.
- Donti PL, Rolnick D, Kolter JZ (2021) DC3: A learning method for optimization with hard constraints. *Internat. Conf. Learn. Representations*.
- Dresser S (2025) Amazon launches a new AI foundation model to power its robotic fleet and deploys its 1 millionth robot. *Amazon News Online* (June 30), <https://www.aboutamazon.com/news/operations/amazon-million-robots-ai-foundation-model>.
- Elmachtoub AN, Grigas P (2022) Smart “predict, then optimize”. *Management Sci.* 68(1):9–26.
- Ergun O, Hopp WJ, Keskinocak P (2023) A structured overview of insights and opportunities for enhancing supply chain resilience. *IIE Trans.* 55(1):57–74.
- Fransoo J, Udenio M (2021) The bullwhip effect. Vickerman R, ed. *International Encyclopedia of Transportation*, vol. 3 (Elsevier, Amsterdam), 130–135.
- Grand View Research (2023) Workforce management market size, share & trends analysis report (2023–2030). Industry report, Grand View Research, San Francisco.
- Holzhaecker M, Loch H, Miller J, Scott A (2025) Property rights restrictions and self-employed workers: Evidence from for-hire owner-operators in US trucking. *Production Oper. Management* 34(6):1306–1325.
- Hu M, Liu S (2026) Coupling artificial intelligence and supply chain operations. Cohen MC, Dai T, eds. *AI in Supply Chains: Perspectives from Global Thought Leaders*, Springer Series in Supply Chain Management, vol. 27 (Springer, Cham, Switzerland), 123–134.
- Hu M, Liu X, Shen G (2025) AI as flesh, OM as bone: Semiparametric regression with deep neural networks. Preprint, submitted September 8, <https://doi.org/10.2139/ssrn.5460214>.
- Hunt MS, Dai T, Abramoff MD (2026) Evaluating commercial multimodal AI for diabetic eye screening and implications for an alternative regulatory pathway. *NPJ Digital Medicine* 9(1):42.
- Jacobides MG, Brusoni S, Candelon F (2021) The evolutionary dynamics of the artificial intelligence ecosystem. *Strategy Sci.* 6(4):412–435.
- Kantour N (2025) Amazon to invest around \$54 billion in UK to support innovation, AI push. *The Wall Street Journal Online* (June 24), <https://www.wsj.com/tech/amazon-to-invest-around-54-billion-in-u-k-to-support-innovation-ai-push-493ba55a>.
- Ketokivi M, Mahoney JT (2020) Transaction cost economics as a theory of supply chain efficiency. *Production Oper. Management* 29(4):1011–1031.
- Kim D, Keskinocak P, Pekgun P, Yildirim I (2022) The balancing role of distribution speed against varying efficacy levels of covid-19 vaccines under variants. *Sci. Rep.* 12(1):7493.
- Knight B, Mitrofanov D, Netessine S (2026) Human-algorithm collaboration in gig work: The role of experience, skill level, and task complexity. *Inform. Systems Res.*, ePub ahead of print February 9, <https://doi.org/10.1287/isre.2024.1664>.
- Kwon C, Raman A (2023) The effects of inconsistent work schedules on employee lateness and absenteeism. Preprint, submitted November 2, <https://doi.org/10.2139/ssrn.4592763>.
- Kwon C, Moreno A, Raman A (2025) The impact of input inaccuracy on leveraging AI tools: Evidence from algorithmic labor scheduling. Preprint, submitted November 20, <https://doi.org/10.2139/ssrn.4602747>.
- Kwon C, Raman A, Tamayo J (2024) Human–computer interactions in demand forecasting and labor scheduling decisions. Preprint, submitted April 15, <https://doi.org/10.2139/ssrn.4773773>.
- Lebovitz S, Lifshitz-Assaf H, Levina N (2022) To engage or not to engage with AI for critical judgments: How professionals deal with opacity when using AI for medical diagnosis. *Organ. Sci.* 33(1):126–148.
- Lee HL, Padmanabhan V, Whang S (1997) Information distortion in a supply chain: The bullwhip effect. *Management Sci.* 43(4):546–558.
- Lee HL, Chen Z, Qi Y, Shen ZJM (2026) AI-driven supply chain transformation. Cohen MC, Dai T, eds. *AI in Supply Chains: Perspectives from Global Thought Leaders*, Springer Series in Supply Chain Management, vol. 27 (Springer, Cham, Switzerland), 223–244.
- Legros B, de Véricourt F (2025) Human–AI interaction in congested service systems: When it improves performance, and when it doesn’t. Preprint, submitted September 4, <https://doi.org/10.2139/ssrn.5401875>.
- Lei D, Qi Y, Liu S, Geng D, Zhang J, Hu H, Shen ZJM (2024) Pooling and boosting for demand prediction in retail: A transfer learning approach. Preprint, submitted June 26, <https://doi.org/10.2139/ssrn.4490516>.
- Li B, Mellou K, Zhang B, Pathuri J, Menache I (2023) Large language models for supply chain optimization. Preprint, submitted July 8, <https://arxiv.org/abs/2307.03875>.
- Licht T, Wohlrabe K (2024) AI adoption among German firms. Preprint, submitted December 5, <https://doi.org/10.2139/ssrn.5045243>.
- Lighthill J (1973) Artificial intelligence: A general survey. The Lighthill Report commissioned by the UK Science Research Council.
- Lundvall B, Rikap C (2022) China’s catching-up in artificial intelligence seen as a co-evolution of corporate and national innovation systems. *Res. Policy* 51(1):104395.
- Luo X, Tong S, Fang Z, Qu Z (2019) Machines vs. humans: The impact of artificial intelligence chatbot disclosure on customer purchases. *Marketing Sci.* 38(6):937–947.
- Marian K, Xing J (2024) Market guide for retail workforce management applications. Industry report, Gartner report (Document ID: G00800951), Gartner, Inc., Stamford, CT.
- Menache I, Pathuri J, Simchi-Levi D, Linton T (2025) How generative AI improves supply chain management. *Harvard Bus. Rev.* 104(1–2):86–95.
- Merton RK (1987) Three fragments from a sociologist’s notebooks: Establishing the phenomenon, specified ignorance, and strategic research materials. *Annual Rev. Sociol.* 13(1):1–29.
- Mousavi SM, Cattania C, Beroza GC (2025) The pursuit of reliable earthquake forecasting. *Phys. Today* 78(8):42–50.
- Nikpayam H, Kremer M, de Véricourt F (2024) When delegating AI-assisted decisions drives AI over-reliance. Working paper, ESMT, Berlin.
- Olsen TL (2026) Artificial intelligence in agricultural supply chains: Practice and prospects. Cohen MC, Dai T, eds. *AI in Supply Chains: Perspectives from Global Thought Leaders*, Springer Series

- in Supply Chain Management, vol. 27 (Springer, Cham, Switzerland), 255–266.
- Parker PS, Shonkwiler JS (2014) On the centenary of the German hog cycle: New findings. *Eur. Rev. Agricultural Econom.* 41(1): 47–61.
- Peels R, Udenio M, Fransoo JC, Hoste R (2025) Breeding delays determine the wavelength of the pork cycle. Working paper, Flostock, Vlijmen, Netherlands.
- Pepe MS (2003) *The Statistical Evaluation of Medical Tests for Classification and Prediction* (Oxford University Press, Oxford, UK).
- Petropoulos F, Grushka-Cockayne Y, Siemsen E, Spiliotis E (2025) Wielding Occam's razor: Fast and frugal retail forecasting. *J. Oper. Res. Soc.* 76(8):1564–1583.
- Pinedo ML (2016) *Scheduling: Theory, Algorithms, and Systems*, 5th ed. (Springer, New York).
- Powell WB (2022) *Reinforcement Learning and Stochastic Optimization: A Unified Framework for Sequential Decisions* (Wiley, Hoboken, NJ).
- Simchi-Levi D, Dai T, Menache I, Wu MX (2025) Democratizing optimization with generative AI. Preprint, submitted September 24, <https://doi.org/10.2139/ssrn.5511218>.
- Simon HA (1987) Two heads are better than one: The collaboration between AI and OR. *Interfaces* 17(4):8–15.
- Smalley HK, Keskinocak P, Swann J, Hinman A (2015) Optimized oral cholera vaccine distribution strategies to minimize disease incidence: A mixed integer programming model and analysis of a Bangladesh scenario. *Vaccine* 33(46):6218–6223.
- Song JJ (2026) Reshaping supply chains through AI-empowered automation. Cohen MC, Dai T, eds. *AI in Supply Chains: Perspectives from Global Thought Leaders*, Springer Series in Supply Chain Management, vol. 27 (Springer, Cham, Switzerland), 167–176.
- Soni A (2025) Microsoft's AI edge under scrutiny as OpenAI turns to rivals for cloud services. *Reuters Online* (July 29), <https://www.reuters.com/business/microsofts-ai-edge-under-scrutiny-openai-turns-rivals-cloud-services-2025-07-29/>.
- Stoffel E, Becker AS, Wurnig MC, Marcon M, Ghafoor S, Berger N, Boss A (2018) Distinction between phyllodes tumor and fibroadenoma in breast ultrasound using deep learning image analysis. *Eur. J. Radiology Open* 5:165–170.
- Sutton RS, Barto AG (2018) *Reinforcement Learning: An Introduction*, 2nd ed. (The MIT Press, Cambridge, MA).
- Swaminathan JM, Smith SF, Sadeh NM (1998) Modeling supply chain dynamics: A multiagent approach. *Decision Sci.* 29(3):607–632.
- Tarighat N, Cohen MC, Clark JJ (2025) Domain adaptation for retail demand prediction. *IEEE Access* 13:146267–146294.
- Van den Bergh J, Beliën J, De Bruecker P, Demeulemeester E, De Boeck L (2013) Personnel scheduling: A literature review. *Eur. J. Oper. Res.* 226(2):367–385.
- Vasconcelos H, Jörke M, Grunde-McLaughlin M, Gerstenberg T, Bernstein MS, Krishna R (2023) Explanations can reduce overreliance on AI systems during decision-making. *Proc. ACM Human-Comput. Interaction* 7(CSCW1), Article 129.
- Weise K (2024) AI data centers are creating a new kind of electricity demand. *The New York Times Online* (December 25), <https://www.nytimes.com/2024/12/25/technology/ai-data-centers-electricians.html>.
- Winkler R, Rattner N, Herrera S (2025) Big tech's \$400 billion AI spending spree just got Wall Street's blessing. *The Wall Street Journal Online* (July 31), <https://www.wsj.com/tech/ai/tech-ai-spending-company-valuations-7b92104b>.
- Xing J, Grinter S, Hanscome R, Marian K, Chandra R, Bobo D, Chouksey A (2024) Market guide for workforce management applications. Gartner report (Document ID: G00800951), Gartner, Inc., Stamford, CT.
- Xu Y, Dai H, Yan W (2026) Identity disclosure and anthropomorphism in voice chatbot design: A field experiment. *Management Sci.* 72(1):223–241.
- Zhang Z, Han J, Hu M, Peng Y (2026) LLM-inspired pretrain-then-finetune for small-data, large-scale optimization. Preprint, submitted February 20, <https://doi.org/10.2139/ssrn.6169687>.
- Zhuge M, Zhao C, Ashley D, Wang W, Khizbullin D, Xiong Y, Liu Z, et al. (2024) Agent-as-a-judge: Evaluate agents with agents. Preprint, submitted October 14, <https://arxiv.org/abs/2410.10934>.