

Review

Resilient and Intelligent Supply Chains: Advances and Challenges in AI-Driven Optimization and Forecasting

Alina Itu

Automation and Information Technology, Transilvania University of Brasov, Mihai Viteazu nr. 5, 5000174 Brasov, Romania; alina.itu@unitbv.ro

Abstract

Supply chains are increasingly exposed to compounding disruptions, volatile demand, and sustainability constraints, which challenge optimization approaches designed for stable operating conditions. This review synthesizes recent advances in supply chain optimization with a focus on the integration of artificial intelligence and operations research in decision-making. The paper examines three major capability layers: prescriptive optimization for planning and resource allocation, predictive modeling for demand and risk anticipation, and digitalized execution through simulation and digital twin environments. Across these layers, the analysis shows that hybrid AI-OR architectures tend to outperform isolated methods in settings characterized by high demand volatility, multi-echelon complexity, and disruption exposure, by combining predictive adaptability with constraint-aware decision quality. The review also highlights a strategic shift from single-objective efficiency toward multi-objective performance that jointly manages cost, service, resilience, and environmental impact. From an implementation perspective, the evidence indicates that measurable industrial gains depend less on algorithm novelty alone and more on system-level integration, data governance, and cross-functional deployment. Key research gaps remain in benchmark standardization, explainability, uncertainty-aware optimization, and long-horizon validation under disruption. The paper concludes that the next generation of supply chain optimization will be defined by continuously learning, human-supervised decision ecosystems that remain robust under uncertainty while delivering operational and sustainability outcomes.

Keywords: supply chain resilience; supply chain optimization; demand forecasting; deep learning; digital twins

1. Introduction

In today's intensely globalized and interconnected economy, the design and management of efficient, resilient supply chains have become cornerstones of competitive advantage. Supply chain optimization—encompassing the end-to-end coordination of sourcing, production, inventory, transportation, and distribution—seeks to deliver the right product, in the right quantity, at the right place and time, while minimizing total system cost [1]. Closely intertwined with this challenge is demand forecasting: the discipline of predicting future customer requirements to inform procurement, production, and logistics decisions. Together, these disciplines form the foundation for balancing supply and demand, managing risk, and achieving operational excellence in both stable and turbulent environments.



Received: 11 April 2026

Revised: 18 April 2026

Accepted: 23 April 2026

Published: 28 April 2026

Copyright: © 2026 by the author.

Licensee MDPI, Basel, Switzerland.

This article is an open access article distributed under the terms and conditions of the [Creative Commons Attribution \(CC BY\)](https://creativecommons.org/licenses/by/4.0/) license.

Historically, both supply chain optimization and demand forecasting have relied heavily on operations research and statistical methods—including linear and mixed-integer programming, simulation, and time-series models such as AutoRegressive Integrated Moving Average (ARIMA) [2]. As supply networks became more complex and customer expectations for flexibility increased, traditional approaches revealed critical limitations in managing high dimensionality, rapid change, and uncertainty at scale. The data revolution has transformed the landscape, with information now streaming in real time from point-of-sale systems, sensors, Internet of Things (IoT) devices, and external markets, enabling analytics-powered decision-making previously impossible at scale [3].

A transformative paradigm shift is underway with the increasing adoption of artificial intelligence (AI) and, more notably, neural networks and deep learning. AI-driven methods—from classical machine learning to advanced architectures such as Long Short-Term Memory (LSTM) networks, transformers, and reinforcement learning—are enhancing both demand forecast precision and optimization sophistication [4,5]. Deep neural networks are particularly impactful in capturing complex, nonlinear relationships and extracting patterns from large, heterogeneous datasets, making them well-suited for high-dimensional supply chain problems [6]. These developments have transformed supply chains into adaptive, self-learning ecosystems capable of withstanding market disruptions while remaining responsive to dynamic customer demands [7]. A recent comprehensive survey confirms that AI technologies—including machine learning, natural language processing, and generative AI—are reshaping logistics, risk management, and demand forecasting across the full spectrum of supply chain processes [8].

Digital supply chain twins—computational representations that mirror real-time network states—exemplify this integration of AI with advanced optimization. These systems enable organizations to simulate scenarios, anticipate disruptions, and optimize decisions in real time across multi-echelon networks [9]. However, such technological advances remain challenging to implement at scale. Integrating AI-based solutions into legacy IT architectures, ensuring data quality and security, fostering interpretability and trust in automated recommendations, and maintaining robust performance under extreme volatility remain critical hurdles [7,10].

This review paper critically examines three major, interrelated challenges defining the current landscape of supply chain optimization and demand forecasting:

- **Modeling and optimizing highly complex, multi-echelon supply networks.** Modern supply chains span multiple echelons—including suppliers, manufacturers, warehouses, distributors, and retailers—operating across diverse geographies and time zones. Capturing and optimizing such multidimensional networks presents a formidable challenge. The proliferation of product variants, just-in-time practices, and customized services further complicates optimization, compounded by constraints related to capacity, lead times, perishability, and cross-border logistics [1]. Achieving robust solutions requires integrating stochastic modeling, mixed-integer programming, network design, and simulation alongside real-time optimization strategies. Recent success with reinforcement learning approaches demonstrates notable progress in automating inventory and logistics decisions at scale [11].
- **Achieving robust and adaptive demand forecasting amidst volatility and external shocks.** Accurate demand forecasting is challenged by volatile consumer preferences, seasonality, promotional activities, and external shocks such as pandemics or geopolitical crises. Traditional models often fail under rapid environmental change or when historical patterns lose predictive power. Contemporary approaches blend statistical techniques, machine learning, and deep learning methods, leveraging high-frequency data from IoT sensors, social media signals, and contextual

market information to adaptively learn evolving demand patterns [12]. Emerging neural architectures—particularly transformer-based models and recurrent neural networks—show substantial promise for capturing temporal dependencies at multiple scales, though persistent challenges remain around model selection, explainability, and overfitting in highly dynamic settings.

- **Harnessing AI and deep learning for next-generation, integrated supply chain management.** Digital transformation is reshaping supply chain management, with big data analytics and AI enabling real-time monitoring, predictive insights, and intelligent decision automation at unprecedented scale. Successful deployment poses integration challenges across fragmented IT landscapes, raises interpretability and cybersecurity concerns, and tests algorithm scalability in mission-critical applications [3]. Building trust and seamless collaboration between human decision-makers and intelligent systems is a critical frontier for future research.

This paper adopts a critical narrative review approach, a choice motivated by its cross-disciplinary scope spanning operations research, machine learning, and digital twin engineering, communities with limited terminological overlap where a single systematic search query would be neither practical nor representative. The review focuses on 2018–2025 publications in peer-reviewed journals, supplemented by foundational references that remain actively cited and by technical reports documenting deployment evidence. Sources were selected for methodological state-of-the-art standing, availability of quantitative industrial evidence, or explicit connection across two or more of the three capability layers examined. The resulting corpus prioritizes integrative breadth over exhaustive coverage of any single subfield; domain-specific systematic reviews are cited within each section for readers seeking deeper coverage.

To ensure terminological precision across the interdisciplinary scope of this review, we define in Table 1 the core concepts used throughout the paper. These definitions draw on standard usage in the operations research and supply chain management literatures and are intended to disambiguate terms that are sometimes used interchangeably in broader discourse.

Table 1. Working Definitions of Core Concepts.

Concept	Working Definition as Used in This Review
Robustness	The ability of a plan or policy to maintain acceptable performance across a predefined set of scenarios without requiring real-time modification [13,14]
Resilience	The capacity of a supply chain to absorb disruption, adapt its structure or operations during the disruption, and recover to a target performance level within an acceptable time horizon [15,16]
Viability	The ability of a supply chain to survive and sustain operations under prolonged, severe disruptions that exceed the design assumptions of robustness and resilience mechanisms [17,18]
Adaptivity	The capacity of a decision system to modify its parameters, policies, or structure in response to observed changes in the operating environment, typically through feedback or learning mechanisms [4,11]
Digitalization	The integration of digital technologies, including sensors, cloud platforms, analytics, and AI, into supply chain planning and execution processes to enable data-driven, continuous decision support [7,19]

Table 1. *Cont.*

Concept	Working Definition as Used in This Review
Integration	The architectural and organizational coupling of planning, forecasting, and execution functions into a coherent decision system with shared data, aligned objectives, and coordinated workflows [10,20]
Intelligence (as applied to supply chains)	The embedding of automated sensing, inference, and decision-recommendation capabilities into operational processes, typically through AI/ML components that augment or partially replace manual decision-making [3,11]

By organizing our review around these focal challenges, we synthesize classical and emerging methods, highlight industrial impact of state-of-the-art solutions, and outline promising areas for further innovation.

2. Modeling and Optimizing Highly Complex, Multi-Echelon Supply Networks

2.1. Nature and Significance of Multi-Echelon Supply Chains

Multi-echelon supply chains are structured as interdependent layers of suppliers, production facilities, distribution centers, and downstream demand nodes, where decisions in one echelon propagate across the entire network. In such systems, optimization is no longer a purely local exercise; procurement, production, inventory, and transportation must be coordinated as a coupled control problem under uncertainty. The practical difficulty lies in the simultaneous presence of structural complexity (large, geographically distributed networks) and operational complexity (multi-product portfolios, heterogeneous lead times, capacity limits, perishability, and contractual constraints) [9,13].

Recent resilience and disruption research has shown that disturbances do not remain isolated at the point of origin. Instead, disruption effects amplify and propagate through network topology, policy parameters, and synchronization delays, producing ripple effects that alter lead-time distributions, backlog trajectories, and service performance well beyond the initially affected node [15,16]. This insight has materially changed the objective function of network planning. Classical cost minimization remains necessary, but is increasingly insufficient unless augmented with explicit resilience constructs such as recovery speed, viability under prolonged stress, and stability of service levels during compound disruptions [13,14]. Network modeling perspectives—including complex network analysis and agent-based approaches—provide formal tools for studying how disruption propagates through supply chain topology and for quantifying structural resilience [21,22], while updated risk management frameworks now incorporate ESG and IT dimensions alongside traditional supplier risk categories [23]. In this context, multi-echelon optimization is best viewed as a dynamic balance among efficiency, robustness, and adaptability rather than as a static deterministic problem.

To structure the methodological discussion that follows, it is important to recognize that supply chain decisions operate at distinct horizons, strategic, tactical, and operational, each with different problem structures, data requirements, and solution tolerability. Table 2 maps these decision levels to their characteristic optimization problems and the method families best suited to each. While many real-world planning processes span multiple levels simultaneously, the methodological strengths and limitations discussed in Sections 2.2.1–2.2.3 are most clearly understood when referenced against this hierarchy.

Table 2. Decision-Level Taxonomy for Supply Chain Optimization.

Decision Level	Typical Horizon	Representative Problems	Dominant Methods	Key References
Strategic	1–5 years	Facility location, network design, sourcing structure, capacity investment	Mixed-integer programming, stochastic network design	[1,2,15]
Tactical	Weeks–months	Inventory positioning, production planning, S&OP, supplier allocation	Stochastic programming, simulation-optimization, scenario analysis	[13–15]
Operational	Hours–days	Replenishment, order dispatching, dynamic routing, adaptive inventory control	Reinforcement learning, heuristics, real-time forecasting pipelines	[4–6,11]

A conceptual distinction often blurred in the supply chain optimization literature, and one that has direct methodological consequences, is the type of uncertainty being addressed. At least six categories deserve differentiation. (i) Demand noise refers to stationary stochastic fluctuations around a predictable baseline and is well handled by safety-stock calculus, stochastic programming, and classical forecasting models [2,24]. (ii) Structural breaks, abrupt, persistent regime shifts in demand or supply patterns, invalidate stationarity assumptions and require adaptive or change-detection mechanisms such as regime-switching models or continuously retrained ML pipelines [5,25]. (iii) Supplier and logistics disruptions are discrete, low-frequency, high-impact events that call for scenario-based optimization, network redundancy design, and stress-testing through simulation [13,15,16]. (iv) Model misspecification arises when the functional form or parameter assumptions of an optimization or forecasting model diverge from the true data-generating process; it is mitigated by ensemble methods, robust optimization, and validation against out-of-distribution scenarios. (v) Data latency and quality degradation affect all methods but are particularly damaging to real-time RL and digital twin approaches, whose performance depends on timely and accurate state observation [4,11]. (vi) Distribution shift, the gradual or sudden divergence between training-time and deployment-time data distributions, is the primary failure mode of purely data-driven approaches and motivates transfer learning, domain adaptation, and hybrid architectures that anchor learned models with structural constraints [5,6]. Recognizing these distinctions is important because the case for methodological hybridization is not generic: each uncertainty type has a natural methodological counterpart, and effective system design matches method to uncertainty rather than deploying complexity uniformly.

2.2. Recent Methodological Developments

2.2.1. Mathematical Programming, Stochastic Modeling, and Decomposition

Mathematical programming methods are most directly applicable at the strategic and tactical decision levels, where the problem structure is well-defined, constraints are relatively stable, and the planning horizon justifies the computational investment of exact or near-exact solvers. Mixed-integer formulations are still the dominant paradigm for facility-location decisions, sourcing allocation, capacity positioning, and flow design because they provide strong constraint expressiveness and transparent managerial interpretation [1,2]. However, the scalability limits of deterministic formulations become apparent in high-dimensional, uncertain environments, where demand and supply volatility violate stationarity assumptions.

To address this gap, recent work has increasingly embedded uncertainty through stochastic and resilience-oriented Operations Research (OR) models that explicitly account for disruption exposure and post-disruption system behavior [13,15]. In practice, this

means moving from single-plan optimization toward scenario-conditioned plan families and stress-tested feasible regions. The primary tradeoff is computational: richer uncertainty representations improve realism but increase solution-space complexity. As a result, decomposition strategies and hybrid solver workflows are not optional engineering choices but enabling mechanisms for tractability in real-scale networks [14]. The contemporary pattern in industry and research is therefore an architecture in which exact/near-exact optimization provides structural rigor, while additional modeling layers absorb uncertainty and temporal dynamics.

The practical validity of mathematical programming formulations rests on several assumptions that become limiting in volatile environments. Deterministic models require stable parameterization of costs, capacities, and lead times, an assumption that degrades rapidly under disruption. Stochastic extensions relax this but depend critically on the quality of scenario construction: if the scenario set underrepresents tail events or misspecifies correlation structures across supply nodes, the resulting solutions may be robust to the wrong uncertainties. Decomposition methods, while essential for tractability, introduce coordination loss between subproblems, and the optimality gap introduced by decomposition is rarely quantified in applied settings. These assumptions do not invalidate the methods, but they define the envelope within which their solutions remain reliable.

2.2.2. Simulation and Digital-Twin-Centered Analysis

Simulation and digital twin methods operate primarily at the tactical level, where their value lies in evaluating policy alternatives and stress-testing plans under disruption scenarios before operational commitment, though they increasingly support strategic network design as well. Discrete-event simulation is particularly valuable when policy behavior depends on queueing, delays, nonlinear replenishment dynamics, and time-varying disturbances that are difficult to encode in closed-form optimization [9,15]. Its major contribution is explanatory power: it allows planners to observe how local interventions reshape global trajectories of service levels, inventory, and transportation utilization over time.

Digital twins extend this capability by integrating network representation, operational parameters, and scenario orchestration into a unified computational environment linked to real planning processes [7,9,10]. Conceptually, digital twins narrow the gap between “model for analysis” and “model for operations.” Practically, they enable repeated stress testing of policies before implementation and facilitate contingency design based on quantified KPI impacts. This has direct implications for multi-echelon optimization: instead of selecting a plan solely on nominal optimality, firms can select plans on conditional resilience performance under plausible disruption paths. The resulting workflow is iterative and evidence-driven: optimize, simulate, compare, and revise.

The value of simulation and digital twin approaches depends on the behavioral realism of the underlying model. Discrete-event simulation requires explicit specification of process logic, queue disciplines, replenishment triggers, and delay distributions, parameters that are often estimated from historical data and may not transfer to novel disruption regimes. Digital twins further assume continuous synchronization between the virtual and physical systems; when model fidelity degrades due to unobserved parameter drift or missing feedback loops, scenario outputs can become misleadingly precise. A related limitation is that simulation provides diagnostic insight and policy comparison but does not, by itself, generate optimal solutions, it must be coupled with an optimization or search mechanism to move from evaluation to prescription.

2.2.3. Learning-Based Optimization and Hybrid AI-OR Pipelines

Learning-based methods, particularly reinforcement learning and deep sequence models, are most naturally suited to operational-level decisions where action frequency is high, feedback is rapid, and environment dynamics are nonstationary. RL is particularly well-suited to inventory and replenishment control because its temporal-difference learning mechanism is designed to handle delayed rewards, a structural feature of multi-echelon systems where ordering actions today manifest as service-level and cost outcomes only after lead times spanning days to weeks. Value-based methods (e.g., deep Q-networks) learn state-action value functions that implicitly discount future consequences, making them effective when the action space is discrete (e.g., order/no-order, reorder-point selection). Policy-based and actor-critic methods extend this to continuous action spaces (e.g., order quantities) and are better suited to high-dimensional multi-echelon control where joint actions across locations must be coordinated simultaneously [4,11]. RL approaches can learn adaptive inventory and replenishment policies from interaction data and are particularly attractive for multi-echelon control with stochastic demand and delayed feedback [4,11]. In multi-node, multi-material supply chains such as civil aircraft manufacturing, multi-agent reinforcement learning with POMDP formulations has demonstrated approximately 45% efficiency improvement over single-agent baselines [26]. In parallel, deep time-series methods (including transformer architectures) have improved demand signal extraction in high-dimensional settings, supporting better downstream optimization inputs [5,6].

Nevertheless, pure AI pipelines rarely satisfy industrial requirements on their own. The main constraints remain model governance, stability under distribution shift, and explainability for mission-critical decisions [5,11]. Consequently, in multi-echelon settings with nonstationary demand and frequent disruption exposure, the most promising direction appears to be hybridization rather than replacement: optimization modules enforce feasibility and hard business constraints; Machine Learning (ML)/Deep Learning (DL) modules improve prediction and pattern extraction; RL modules adapt policy parameters under changing conditions. Evidence from recent inventory-control studies further suggests that coupling learning with stochastic programming can improve operational performance while preserving decision discipline [4]. This synthesis is increasingly recognized as the practical endpoint of methodological evolution in complex multi-echelon environments.

Reinforcement learning methods require careful state-space design: if the state representation omits decision-relevant variables (e.g., supplier lead-time variability, in-transit inventory) or includes noisy features, learned policies may converge to locally effective but globally suboptimal behavior. Training stability is a further practical constraint, RL agents in multi-echelon settings are sensitive to reward shaping, discount factors, and exploration strategy, and small design choices can produce qualitatively different policies. Deep sequence models for forecasting, in turn, assume that the temporal patterns learned from training data remain valid during deployment; under structural breaks such as demand regime shifts or logistics network reconfiguration, model performance can deteriorate abruptly unless retraining or adaptation mechanisms are in place. These dependencies underscore why hybrid architectures are attractive: they allow each component to operate within its assumption envelope while other components compensate for its failure modes.

Table 3 compares the primary strength, primary limitation as well as the typical planning role of the methodological approaches.

Three cross-cutting trade-offs emerge from this comparative view. First, a persistent tension exists between feasibility assurance and adaptive responsiveness: methods with strong constraint enforcement (mixed-integer programming, decomposition) offer decision discipline but react slowly to environmental change, while learning-based methods adapt rapidly but lack built-in guarantees that solutions remain operationally feasible. Second,

calibration burden appears as a recurring limitation across nearly all families, scenario construction in stochastic OR, behavioral parameterization in simulation, state-space design in RL, and training-data curation in deep learning all require substantial domain expertise and ongoing maintenance. Third, governance complexity increases monotonically as methods move from transparent mathematical programs toward opaque learned policies, creating a practical ceiling on adoption in regulated or mission-critical settings. These recurring patterns reinforce the argument for hybrid architectures not as a matter of methodological preference, but as a structural response to the fact that no single family resolves all three trade-offs simultaneously.

Table 3. Comparative perspective on methodological families for multi-echelon optimization.

Method Family	Primary Decision Level	Primary Strength	Primary Limitation	Typical Planning Role	Key References
Mixed-integer and network optimization	Strategic/Tactical	Strong feasibility control, transparent constraints	Computational burden under large uncertainty sets	Structural design, capacity, allocation	[1,2]
Stochastic and resilience-oriented OR	Tactical	Explicit treatment of uncertainty and disruption propagation	Scenario/model calibration complexity	Risk-aware tactical planning	[13,15,16]
Decomposition and hybrid solve strategies	Strategic/Tactical	Scalability for large networks	Coordination loss if decomposition is poorly coupled	Large-scale multi-echelon solve workflows	[14]
Discrete-event simulation	Tactical	Dynamic system behavior and policy diagnostics	No intrinsic optimality guarantee	Stress testing and policy comparison	[9,15]
Digital twins	Tactical/Strategic	Integration of model, scenarios, and operational decision support	Data integration and governance burden	Continuous resilience monitoring and what-if planning	[7,9,10]
RL-based policy optimization	Operational	Adaptive control in stochastic, sequential contexts	Training stability and explainability constraints	Replenishment and inventory policy adaptation	[4,11]
Deep sequence forecasting (transformer-based)	Operational	Captures nonlinear long-range temporal dependencies	Data intensity and drift sensitivity	Forecast layer feeding optimization	[5,6]
Hybrid AI-OR architectures	Cross-level	Balances adaptivity and feasibility	Higher implementation complexity	Closed-loop planning and replanning	[4,5]

2.3. Industrial Impact

The industrial impact of these methodological advances is best characterized as a transition from isolated optimization projects to integrated planning ecosystems. In earlier implementations, companies often deployed single-point optimizers for network design or inventory settings and revisited assumptions infrequently. Contemporary deployments are increasingly continuous and layered: optimization engines generate feasible plans, simulation environments stress-test plan robustness, and AI modules improve forecast

quality or policy responsiveness [7,9,11]. This shift is particularly visible in sectors exposed to demand volatility and disruption cascades, where planning value depends as much on recovery behavior as on nominal efficiency.

From a comparative standpoint, three patterns are consistently observed. First, OR-centric methods remain strongest for governance, feasibility, and auditability of structural decisions. Second, simulation-centric methods provide superior visibility into disruption mechanics, tail effects, and Key Performance Indicator (KPI) sensitivity under compound shocks. Third, AI-centric methods can increase adaptivity and speed, but their net benefit depends heavily on data quality, monitoring, and explainability controls. Based on the pattern observed across the reviewed studies, we interpret this as convergence toward hybrid pipelines in which each method family addresses the failure modes of the others [4,5,11,13]. This convergence has strategic implications: available evidence suggests that organizations operating in disruption-prone, multi-echelon environments benefit from institutionalizing hybrid planning loops, which can improve service continuity under stress while preserving cost discipline over longer horizons.

2.3.1. Ready-Made Tools and Deployment Patterns

Commercial planning platforms now reflect this convergence in their product architectures. anyLogistix emphasizes combined optimization and simulation for network design, inventory planning, and risk analysis, with scenario-oriented workflows suitable for stress testing [27]. SAP IBP provides an integrated cloud planning environment spanning demand planning, multilevel supply planning, S&OP, and simulation features intended to coordinate planning horizons in one process backbone [20]. Blue Yonder's supply chain planning suite similarly combines optimization and real-time scenario modeling with AI/ML-assisted planning components [28]. o9 positions its platform around a unified data model, mixed solver stack (heuristics, LP, MIP), and scenario-driven planning with high automation intent [29].

In deployment terms, these tools are rarely selected only on algorithmic sophistication. Practical selection criteria include integration burden with Enterprise Resource Planning (ERP)/control-tower landscapes, solver openness for enterprise-specific constraints, model transparency for governance, and the ability to run cross-functional what-if analysis at planning cadence. A common implementation pattern is phased adoption: firms begin with a planning module that addresses an urgent bottleneck (typically demand-supply balancing or inventory), then expand toward scenario orchestration and end-to-end synchronization as data maturity improves.

2.3.2. Representative Case Study

A representative and quantitatively rich example is the digital-twin simulation study of food retail supply-chain resilience by Burgos and Ivanov [9]. The modeled network includes 10 product categories, 28 supermarket locations across five countries, 30 suppliers, and three distribution centers, and it evaluates distinct and combined disruption scenarios reflecting pandemic-era conditions. The scenario design includes demand surges of +75%, +10%, and +35% across phases, temporary supplier shutdowns, and transport bottlenecks [9].

The quantitative findings are directly relevant to multi-echelon optimization policy design. Relative to baseline, mean lead time increased from 0.10 days to 3.112 days in the combined disruption scenario. Late orders rose from 0 (baseline) to 3949 under demand shock and to 9067 under combined disruptions, corresponding to roughly 8% of orders in the latter case. Service performance deteriorated accordingly: ELT service level moved from 1.00 in baseline conditions to 0.87 in the demand-shock scenario and 0.80 in the combined scenario. Backlog effects were also substantial, with backlog reaching 2.78% of total demand

in the demand-shock case and 2.67% under combined disruptions [9]. Logistics pressure became visible through a rise in total shipped vehicles from 39,697 to 49,387.

These results support two core inferences for this review. First, compound disruptions produce non-additive effects; service deterioration and delay inflation are not linearly predictable from single-shock analysis. Second, planning resilience requires coordinated methods: optimization for structural feasibility, simulation for dynamic stress diagnostics, and adaptive control mechanisms for recovery-phase policy tuning [4,7,9,11]. In our assessment, this case exemplifies why methodological pluralism is not merely an academic preference but an operational necessity in multi-echelon systems exposed to compound disruptions.

A note on representativeness is warranted. The Burgos and Ivanov case [9] models a specific food retail network under pandemic-era disruption scenarios with carefully calibrated parameters. While the quantitative findings are internally valid within the simulation design, their transferability to other sectors, network topologies, or disruption types should not be assumed without recalibration. Simulation-based case studies inherently reflect the modeler's assumptions about demand distributions, capacity constraints, and recovery dynamics; results are therefore conditional on these design choices. More broadly, the absence of standardized multi-echelon benchmarks in the supply chain resilience literature means that most published case studies, including this one, function as structured illustrations rather than as generalizable empirical evidence.

3. Achieving Robust and Adaptive Demand Forecasting Amid Volatility and External Shocks

3.1. Strategic Role and Technical Complexity of Demand Forecasting

Demand forecasting is a core decision input for inventory control, replenishment, production scheduling, and service-level management. In contemporary supply chains, its role has expanded from periodic planning support to continuous operational control, especially in environments characterized by short product life cycles, high promotion intensity, omnichannel substitution effects, and exogenous shocks. Under such conditions, forecast quality is defined not only by point accuracy but also by robustness to structural breaks, responsiveness to regime shifts, and usefulness for downstream decision optimization.

A central challenge is that modern demand processes are increasingly nonstationary. Historical regularities can be rapidly invalidated by policy interventions, macroeconomic shocks, supply constraints, or abrupt channel migration. This weakens purely static model classes and motivates adaptive, probabilistic, and multi-resolution pipelines. Recent forecasting research and practice therefore emphasize hybrid architectures that combine statistical baselines, machine learning predictors, and operational calibration layers to preserve both predictive power and deployability [2,5,6,25]. In this setting, forecasting must be evaluated as part of a decision system rather than as an isolated modeling task.

Critically, the relationship between forecast quality and supply chain performance is not direct, it is mediated by the decision rules and system parameters that consume the forecast. Forecast bias, for instance, does not simply reduce accuracy; positive bias systematically inflates safety stocks and working capital, while negative bias increases stockout frequency and emergency procurement costs. Forecast variance, even when unbiased, amplifies replenishment instability through the bullwhip mechanism, particularly in multi-echelon systems with fixed reorder intervals. Calibration, the alignment between predicted uncertainty and realized outcome distributions, determines whether probabilistic forecasts translate into appropriate service-level targets or into either chronic over-protection or under-protection of inventory positions. Horizon sensitivity matters because models that perform well at short lead times may degrade substantially at the tactical horizons where

production commitment and supplier allocation decisions are made. These distinctions imply that evaluating forecast methods solely on point-accuracy metrics such as MAE or MAPE can be misleading from an operational standpoint: a model with marginally lower accuracy but better calibration and lower bias may produce superior inventory and service outcomes when embedded in a replenishment policy. The practical consequence is that forecast evaluation in supply chain contexts should be decision-consistent, assessed against the operational cost or service-level impact of the decisions the forecast informs, not against statistical loss functions in isolation [2,24].

3.2. Recent Methodological Developments

3.2.1. Statistical, Hierarchical, and Probabilistic Forecasting

Statistical forecasting remains foundational, but current developments emphasize structured combination, hierarchical coherence, and uncertainty-aware outputs. Hierarchical and aggregation-aware approaches are especially important in supply chains, where decisions are made simultaneously at Stock Keeping Unit (SKU)-store, category, regional, and network levels. Evidence from recent supply-chain-focused reviews confirms that aggregation strategy materially affects forecast quality and should be chosen in relation to the intended decision layer rather than by convention [30].

Probabilistic framing is equally critical. Inventory and service policies depend on risk distributions, not only on expected values. Consequently, modern operational pipelines increasingly require quantile or distributional forecasts that can be mapped to stockout and overstock costs. This transition aligns with broader forecasting evidence showing that model ranking under one metric may not transfer to other decision criteria, reinforcing the need for decision-consistent evaluation frameworks [2,24].

3.2.2. Machine Learning and Deep Learning for Demand Signals

The second major development is the move from local univariate models to global learning architectures trained across large panels of related series. Deep learning models such as Temporal Fusion Transformers (TFT) and efficient transformer variants were developed to handle mixed covariate structures, long-range dependencies, and multi-horizon outputs in practical forecasting settings [6,25]. Hybrid architectures combining BiLSTM and NARX networks have also demonstrated strong performance in supply chain demand forecasting tasks [31]. Their value in supply-chain applications lies in cross-series generalization, where information from stable demand segments can improve predictions for sparse or volatile items.

At the same time, interpretability and computational efficiency remain binding constraints. Transformer research has therefore evolved toward architectures that reduce memory complexity and support operational-scale inference while preserving performance [6]. Survey evidence further indicates that transformer-based methods now cover a broad set of forecasting regimes, but performance depends strongly on data curation, horizon design, and decomposition strategy [5]. In short, deep models are not universally dominant; they are most effective when embedded in disciplined feature, validation, and governance workflows.

3.2.3. Competition-Driven and Large-Scale Retail Forecasting Advances

A third methodological stream comes from large-scale benchmarking in retail forecasting. M5 (Makridakis 5, the fifth edition of the Makridakis forecasting competition)-related studies have accelerated progress in global modeling, ensembling, and hierarchy-aware evaluation. Importantly, work assessing the representativeness of M5 data suggests only small discrepancies relative to additional real retailer datasets, supporting cautious transferability of methodological findings to broader retail contexts [32]. Complementary evidence

from M5-oriented model-development studies shows that carefully engineered global ensembles can rank within the top 1% of competition submissions, illustrating the practical strength of scalable cross-learning pipelines [33].

These developments are relevant for industrial forecasting because they operationalize three principles: shared learning across many related series, explicit hierarchy handling, and metric-aware optimization under retail-specific loss functions. The practical implication is that high-performing demand forecasting is increasingly a systems engineering problem involving data architecture, model governance, and continuous recalibration.

Table 4 summarizes the core advantage, the principal limitation and the typical deployment context of robust and adaptive demand forecasting approaches.

Table 4. Comparative Perspective on Robust and Adaptive Demand Forecasting Approaches.

Approach Class	Core Advantage	Principal Limitation	Typical Deployment Context	Key References
Hierarchical/aggregation-aware statistical forecasting	Coherent multi-level planning support	Requires careful aggregation design and reconciliation strategy	Multi-level supply-chain planning	[2,30]
Probabilistic forecasting	Direct support for risk-aware inventory policies	Calibration and evaluation complexity	Safety-stock and service-level optimization	[2,24]
Global ML forecasting	Cross-series learning improves sparse/volatile item performance	Sensitivity to feature engineering and drift	Large SKU-location portfolios	[30,33]
Deep sequence models (TFT/Transformer family)	Captures nonlinear dynamics and long-range dependencies	Data and infrastructure intensity; interpretability constraints	High-frequency, high-dimensional forecasting	[5,6,25]
Competition-informed ensemble pipelines	Strong empirical performance at scale	Transfer requires data/process alignment	Retail demand platforms with large panel data	[32,33]

The comparative structure of Table 4 reveals two patterns that cut across forecasting method families. First, there is a consistent trade-off between predictive power and operational deployability: deep sequence models and global ML approaches achieve superior accuracy on complex demand patterns, but impose higher infrastructure, interpretability, and maintenance requirements than statistical and hierarchical methods. This means that the optimal method choice depends not only on forecast accuracy but also on the organization's data maturity, model governance capacity, and tolerance for retraining overhead. Second, drift sensitivity appears as a shared vulnerability across all data-driven approaches, whether statistical, ML, or deep learning, differing only in the speed and severity of degradation. This observation supports the case for ensemble and hybrid pipelines not primarily for accuracy gains, but as a risk management strategy: methodological diversity reduces the probability that a single mode of failure simultaneously degrades all forecast components.

3.3. Industrial Impact

Industrial demand forecasting has moved decisively toward integrated, adaptive pipelines that combine statistical reliability with ML-based signal extraction and operational decision logic. In retail and consumer-facing sectors, this has improved forecast responsiveness at fine granularity, enabling better synchronization of replenishment cycles, lower stockout exposure, and tighter inventory control. The strongest gains are reported

where pipelines ingest multi-source covariates and continuously re-estimate demand states rather than relying on static periodic re-forecasting [25,32,33].

From a comparative perspective, method families contribute differently to industrial value creation. Statistical and hierarchical methods remain strong on stability, auditability, and coherence across decision levels; deep and global learning models provide superior adaptivity and cross-series transfer in highly volatile panels; probabilistic outputs create direct value by translating predictive uncertainty into explicit service-risk tradeoffs. The evidence reviewed in this section points to a dominant implementation pattern that is hybrid: organizations increasingly combine methods rather than selecting a single paradigm, using each where its failure modes are best controlled [2,25,30,33].

3.3.1. Ready-Made Tools for Adaptive Forecasting

Commercial platforms now expose these capabilities in production-ready form. SAP IBP provides integrated demand planning, demand sensing, multilevel supply planning, and what-if simulation in a unified planning process [20]. Blue Yonder's planning stack emphasizes AI/ML-driven forecasting, LP-enabled planning, and real-time scenario modeling for connected demand-supply decisions [28]. o9 similarly offers a unified planning data model with mixed-solver support, touchless automation, and scenario-centric planning workflows [29].

Amazon Forecast has been widely used as a managed service for time-series and quantile forecasting across inventory, workforce, and demand use cases, including large-scale SKU-level settings [34,35]. Across these tools, practical differentiation is less about headline algorithms and more about integration maturity, hierarchy handling, probabilistic support, and operationalization speed.

3.3.2. Representative Quantitative Case Study

A representative quantitative deployment is the More Retail demand-forecasting and automated-ordering program built with Amazon Forecast and integrated into ERP operations [36]. The case addresses fresh-produce demand at high granularity across a large network (hundreds of stores and thousands of store-SKU combinations), where forecasting error directly translates into both spoilage and stockout risk.

Reported outcomes are substantial: forecasting accuracy improved from 24% to 76%, wastage declined by up to 30% in fresh produce, in-stock rate improved from 80% to 90%, and gross profit increased by 25% [36]. Methodologically, the program combined segmentation (ABC-XYZ), extensive feature experimentation, quantile-based ordering logic, and end-to-end automation. The same ecosystem reports additional quantitative gains in other deployments, including forecast-accuracy improvements and inventory/service benefits across multiple industries [35].

This case is important for the present review because it demonstrates the operational mechanism by which robust forecasting creates value: not by improving a benchmark metric in isolation, but by coupling adaptive forecasts with ordering policy, uncertainty-aware thresholds, and production-grade execution pipelines. It thus illustrates the broader thesis of this section that resilient forecasting is an integrated capability spanning model, data, and decision process design.

It is important to acknowledge the conditions that enabled these results. The More Retail deployment [36] benefited from a large store network generating high-frequency transactional data, a clearly defined and high-cost loss function (fresh-produce waste), organizational commitment to end-to-end automation, and a technically mature integration with ERP systems. These factors, data richness, process redesign, and implementation quality, are not universally present across industries or firm sizes. The reported quantitative

improvements should therefore be understood as achievable under favorable conditions rather than as expected outcomes of adopting similar algorithms in less mature environments. Similarly, the broader Amazon Forecast customer evidence [35] typically reflects early-adopter organizations with above-average data infrastructure and analytics capability, which limits direct extrapolation to the general population of supply chain operators.

4. Harnessing Digitalization, AI, and Deep Learning for Next-Generation Integrated Supply Chain Management

4.1. Digital Transformation and Its Supply-Chain Implications

Digital transformation in supply chains has shifted from isolated digitization initiatives to system-level redesign of planning and execution architectures. Contemporary supply chains increasingly operate as cyber–physical, data-intensive networks in which transactional, sensor, and contextual streams are integrated for near-real-time decision support. This transition has materially changed the operating logic of supply-chain management: planning cycles are shorter, control becomes more continuous, and resilience is treated as a design objective rather than a post-disruption response capability [7,9,19].

A central implication of this shift is that digitalization and intelligence must be co-designed. Data pipelines without adaptive analytics produce limited operational value, while AI without trusted data integration and process embedding fails at scale. Recent work on viability and resilience reinforces that high-performing supply chains require dynamic reconfiguration capabilities, cross-functional data visibility, and decision mechanisms that can operate under severe uncertainty [17,18]. Consequently, integrated digital supply chains are best understood as socio-technical systems in which algorithmic quality, data governance, and organizational readiness are jointly determinant of performance.

4.2. Recent Methodological Developments

4.2.1. End-to-End Visibility, Digital Twins, and Continuous Monitoring

The first major development is the maturation of end-to-end visibility architectures based on platform integration, event-driven data flow, and digital representation of operational states. Digital twins have emerged as a key mechanism in this trajectory, enabling synchronized analysis of network structure, control policies, and disruption scenarios. In practical terms, digital twin environments support stress testing, policy comparison, and contingency design before physical deployment, thereby reducing the lag between sensing and response [7,9,10,19].

Importantly, visibility is no longer confined to descriptive tracking. It increasingly supports predictive and prescriptive layers, including anticipated bottleneck detection, lead-time risk profiling, and dynamic allocation analysis. Recent work further demonstrates that integrating digital twins with IoT sensing and deep reinforcement learning (e.g., DDPG) enables simultaneous optimization of cost, delivery time, and carbon emissions in agile supply chain settings [37]. This is one reason digitalization is strongly linked to resilience outcomes in recent empirical literature: visibility improves adaptation quality when uncertainty is high and response windows are short [38,39].

4.2.2. AI-Driven Planning, Adaptive Optimization, and Autonomous Decision Support

A second development concerns the operationalization of AI across tactical and execution-level decisions. Reinforcement-learning and deep-learning approaches are being used to complement traditional optimization in inventory control, supply allocation, and policy adaptation under stochastic demand and supply conditions [4,11]. Rather than replacing OR foundations, leading architectures integrate AI components where nonlinear

signal extraction and fast policy adjustment are most valuable, while preserving feasibility constraints and controllability through optimization modules.

In parallel, forecasting and planning models increasingly share representations and pipelines, allowing demand intelligence to propagate more directly into replenishment and allocation decisions [5,6,25]. The managerial significance is substantial: integrated AI planning can reduce reaction latency and improve alignment between forecast uncertainty and operational action. However, literature consistently indicates that deployment success depends on explainability, robustness monitoring, and failure-mode governance in high-stakes environments [17,30].

While explainability is frequently cited as a requirement for AI adoption in supply chain decision-making, the specific technical progress of Explainable AI (XAI) within operations research contexts remains limited. Post hoc attribution methods such as SHAP (SHapley Additive exPlanations) have been applied to demand forecasting models to identify which input features, promotions, seasonality, price changes, drive individual predictions, providing forecast consumers with feature-level transparency. However, SHAP and similar perturbation-based methods explain predictions, not decisions: they do not reveal why an optimization model selected a particular replenishment quantity, sourcing allocation, or safety-stock level given the forecast. Explaining constrained decisions requires a different class of techniques, including constraint attribution (identifying which constraints are binding and how relaxing them would change the solution), sensitivity reporting (quantifying how objective value responds to parameter perturbation), and counterfactual explanation (showing which minimal input changes would alter the recommended action). These methods exist in the OR toolkit as duality analysis and parametric programming but have not been systematically integrated with ML-based planning components into unified explanation interfaces. A further gap concerns the interpretability of reinforcement learning policies, where the mapping from state to action is encoded in neural network weights that resist intuitive summarization. Early approaches using decision-tree distillation or attention-based policy architectures show promise but have not yet been validated in production supply chain settings. Until XAI methods are developed that can jointly explain the predictive and prescriptive layers of hybrid AI-OR systems, explainability will remain a barrier to adoption in regulated and high-stakes environments.

4.2.3. Interoperability, Data Trust, and Human-Centered Governance

The third development addresses integration barriers across heterogeneous enterprise landscapes. Legacy ERP/Manufacturing Execution System (MES) cores, cloud-native planning services, and edge-level automation components often coexist with inconsistent semantics, update cadence, and ownership rules. Blockchain-based smart contracts offer a complementary integration mechanism by enhancing traceability, transparency, and accountability across supply chain partners while eliminating single points of failure in transactional record-keeping [40]. As a result, interoperability and trust architecture have become methodological concerns, not merely IT concerns. Current research and practice emphasize the need for standardized interfaces, data lineage controls, and governance frameworks that ensure transparency and accountability of AI-assisted decisions [10,38,41,42].

Human-in-the-loop design is also increasingly treated as a structural requirement for industrial AI. In integrated supply chains, escalation pathways, exception handling, and decision explainability are critical to maintain operator trust, regulatory defensibility, and coordinated execution across organizational boundaries. This aligns with broader resilience findings: digital capability contributes most when combined with organizational flexibility and clear governance of model-driven decisions [18,38].

Table 5 summarizes the principal contribution, the key limitation and the typical deployment of digitalization, AI and deep learning approaches for integrated supply chains.

Table 5. Comparative Analysis of Digitalization, AI, and Deep-Learning Approaches for Integrated Supply Chains.

Approach	Principal Contribution	Key Limitation	Typical Deployment Context	Key References
Digital supply twins	Scenario-based resilience analysis and policy stress testing	Model-data synchronization burden	Network redesign, disruption planning	[7,9,10,19]
Real-time data visibility platforms	Faster anomaly detection and control transparency	Integration complexity across fragmented systems	Multi-echelon monitoring and control towers	[10,38,39]
AI-assisted demand-to-replenishment pipelines	Shorter reaction cycles and improved tactical alignment	Data quality and drift sensitivity	Retail/consumer planning loops	[5,25,36]
RL-based adaptive inventory/policy control	Dynamic response under stochastic uncertainty	Stability, safety, explainability constraints	Multi-echelon inventory and replenishment	[4,11]
Hybrid AI-OR planning architectures	Feasibility plus adaptivity in one stack	Higher implementation complexity	Tactical planning under volatility	[4,17,42]
Cloud-native planning suites	Scalable collaboration and faster deployment	Data governance and migration risk	Cross-regional planning and S&OP	[20,28,29]
Multi-enterprise digital collaboration networks	Better coordination across suppliers and channels	Incentive and interoperability challenges	Extended enterprise ecosystems	[28,39,43]
Data-analytics-enabled resilience capabilities	Stronger adaptability and recovery performance	Capability-development and change effort	Organizations with high disruption exposure	[18,38]
Viability-oriented supply-chain design	Long-horizon survivability under severe disruptions	Measurement and operationalization difficulty	Strategic resilience and transformation planning	[17,18]
Human-in-the-loop AI governance	Improves trust, adoption, and exception quality	May increase decision cycle time in some contexts	Regulated or high-impact operations	[10,39,41]

Across Table 5, two structural patterns deserve emphasis. First, the tension between integration depth and implementation complexity is visible in every row: approaches that deliver the greatest decision value, digital twins, hybrid AI-OR architectures, multi-enterprise collaboration networks, are also the most demanding in terms of data governance, organizational change, and cross-functional coordination. This suggests that the adoption frontier is defined less by algorithmic availability than by institutional readiness. Second, explainability and trust constraints recur across AI-assisted planning, RL-based control, and autonomous decision support, indicating that the human-governance dimension is not a peripheral concern but a binding constraint on the deployment envelope of advanced methods. These patterns collectively support the paper's central thesis that next-generation supply chain management is fundamentally a systems-integration challenge in which technological capability must be matched by organizational and governance maturity.

4.3. Industrial Impact

The combined adoption of digitalization, AI, and deep learning is already producing measurable shifts in supply-chain performance, especially in sectors operating under high volatility and high service-level pressure. Firms report improved responsiveness through tighter sensing-planning-execution coupling, more frequent replanning, and greater resilience to disruption propagation. The practical pattern is clear: value is highest when digital visibility is operationally connected to adaptive decision logic and embedded into execution workflows, rather than deployed as stand-alone analytics [9,11,36].

From a comparative perspective, digital twin and simulation environments provide superior diagnostic depth for structural resilience and contingency design, whereas AI-driven planning components provide higher short-cycle adaptivity in daily operations. Cloud planning suites and platform ecosystems improve scalability and collaboration, but their performance impact depends heavily on data governance maturity and process redesign quality. Therefore, the reviewed evidence does not support dominance of any single technology class across all objectives; rather, the strongest reported industrial outcomes come from layered architectures that combine visibility, optimization, learning, and governance [7,10,20,28,29].

4.3.1. Ready-Made Tools in Current Practice

Commercial tool landscapes now reflect this layered design. SAP IBP supports integrated demand, supply, inventory, and scenario workflows in a cloud planning context [28]. Blue Yonder provides AI/ML-enabled planning with optimization and scenario capabilities in a connected suite [28]. o9 emphasizes a unified planning model, solver flexibility, and automation-oriented decision support [29]. anyLogistix remains a widely used environment for network design, simulation, and risk analysis, particularly in strategic and tactical what-if studies [27]. Amazon Forecast has been used as a managed forecasting component in broader operational architectures and integrated with enterprise ordering systems [34,36].

Across these tools, selection is increasingly driven by interoperability characteristics, governance fit, scenario depth, and operationalization speed, rather than by isolated algorithmic claims. In other words, tool performance is architecture-dependent: capabilities generate value when integrated with enterprise data flows, process ownership, and execution controls.

4.3.2. Representative Quantitative Case Study

A representative implementation is the More Retail and Ganit deployment using Amazon Forecast integrated with Oracle ERP for automated ordering in a large grocery network [36]. The program targeted fresh-produce demand at store-SKU-day level, where both under-forecasting (stockouts) and over-forecasting (wastage) carry material financial penalties. The deployment combined feature-rich ML forecasting, quantile-based ordering logic, and end-to-end automation.

Reported quantitative outcomes include an increase in forecast accuracy from 24% to 76%, wastage reduction by up to 30% in fresh produce, in-stock-rate improvement from 80% to 90%, and gross-profit increase by 25% [36]. The broader Amazon Forecast customer evidence also reports quantified improvements in related contexts, including 9.9% SKU-management improvement valued at over £110 million (The Very Group), 8% forecast-accuracy improvement with projected annual savings (Foxconn, New Taipei City, Taiwan), and measurable inventory/service gains in retail deployments [35].

The case demonstrates the central argument of this section: industrial impact does not arise from model sophistication alone, but from integrated architecture. Forecasting,

ordering, ERP connectivity, and governance were implemented as a coherent decision system, converting digital and AI capability into measurable operational outcomes.

As noted in Section 3.3.2, the quantitative outcomes reported here reflect a deployment with favorable data conditions, strong organizational commitment, and end-to-end process integration; these contextual factors are integral to the result and should be considered when assessing transferability to other settings.

Translating these governance principles into engineering practice requires a set of concrete technical mechanisms that are increasingly documented in applied AI deployment literature but remain underspecified in most supply chain optimization studies. Model monitoring entails continuous tracking of prediction-versus-actual distributions, with automated alerts triggered when statistical divergence (e.g., via population stability index or Kolmogorov–Smirnov tests) exceeds predefined thresholds, signaling potential concept drift. Retraining cadence must be determined not by fixed schedules but by monitored performance degradation, with pre-validated retraining pipelines that can execute without production downtime. Drift detection should operate at both the input level (covariate shift in demand signals, lead-time distributions, or supplier behavior) and the output level (forecast error trend, policy reward degradation in RL systems). Confidence thresholds govern when an AI-generated recommendation is auto-executed versus escalated for human review: a replenishment order within historical bounds and high model confidence may proceed autonomously, whereas an anomalous demand spike triggering unusual allocation should invoke a manual override pathway. Fallback procedures define what happens when the AI component fails entirely, for instance, reverting to a rules-based safety-stock policy or the last validated plan. Exception management protocols must specify escalation rules, resolution authority, and time limits to prevent decision queues from stalling execution. Finally, auditability requires that every AI-assisted decision is logged with its input state, model version, confidence score, and whether human override occurred, creating a traceable decision record for both operational review and regulatory compliance. Without these mechanisms, governance remains aspirational rather than operational, and the reliability gap between pilot demonstrations and production-grade deployment persists.

5. Conclusions and Future Research Directions

5.1. Integrative Conclusions

This review shows that the strongest advances in supply chain optimization come from integration rather than substitution: AI methods improve perception, prediction, and adaptation, while OR models preserve structure, feasibility, and decision quality under constraints. Across the literature, digital twins, learning-based forecasting, and optimization models each contribute different strengths, but the reviewed evidence indicates that their value increases substantially when they are connected in a single decision architecture, particularly in environments with compound disruption risk and multi-objective performance requirements. In practical terms, organizations benefit most when forecasting outputs are translated into scenario-aware plans, those plans are stress-tested in simulation, and final decisions are validated against operational constraints and service-level targets.

A second central conclusion is that resilience and viability cannot be treated as afterthoughts to efficiency. Recent disruptions have shown that models optimized only for nominal conditions often fail when lead times, demand patterns, or logistics capacity shift abruptly. Based on the converging evidence from the studies reviewed, we observe an emerging consensus that robust performance requires multi-objective formulations in which cost, service, carbon impact, and recoverability are jointly managed. This shifts the

goal from static optimality to adaptive optimality, where decision policies remain effective across a wider range of operating states.

To align the strength of these conclusions with the underlying evidence, we distinguish three tiers of confidence. The first tier, consistently supported across the reviewed literature, includes the superiority of integrated method architectures over isolated deployment in volatile multi-echelon settings, the mediating role of decision rules between forecast quality and operational outcomes, and the binding nature of governance and data quality as deployment constraints. The second tier, promising but context-dependent, includes the performance advantage of hybrid AI-OR architectures (which diminishes in stable, low-complexity environments), the adaptive potential of reinforcement learning for inventory control (where production-grade longitudinal evidence remains limited), and the operational value of digital twins (which depends on synchronization fidelity that many organizations have not yet achieved). The third tier, forward-looking rather than empirically established, includes the vision of continuously learning decision ecosystems, multi-objective optimization that jointly manages cost, service, resilience, and environmental impact at scale, and the potential role of foundation models in planning workflows.

5.2. Managerial and Industrial Implications

Drawing on the deployment patterns and case evidence reviewed in Sections 2–4, we suggest that the primary practical implication is architectural: firms benefit from avoiding fragmented analytics pipelines where forecasting, planning, and execution tools operate independently. A layered implementation is more effective, with data governance and feature engineering at the base, forecasting and risk sensing in the middle, and optimization plus execution control at the top. This structure reduces decision latency, improves traceability, and supports faster replanning when disruptions occur.

Implementation success also depends on organizational capability, not only model sophistication. Companies that realize measurable value typically establish shared ownership between operations, data science, and IT; define model-refresh and fallback procedures; and track impact through operational KPIs rather than isolated model metrics. In this context, deployment discipline matters as much as algorithm choice: a moderately complex but well-governed pipeline can outperform an advanced model that is poorly integrated into planning and execution workflows.

5.3. Methodological Gaps and Research Priorities

Despite rapid progress, several methodological gaps remain. Many published studies still evaluate models in narrow experimental settings, which limits generalizability across sectors and disruption regimes. Future research should prioritize cross-domain benchmarking with common datasets, transparent evaluation protocols, and stress-test scenarios that reflect realistic supply uncertainty. This would improve reproducibility and make performance claims more decision-relevant for practitioners.

Another priority is interpretable and risk-aware decision intelligence. As hybrid AI-OR systems become more autonomous, decision-makers need clear explanations for why a recommendation changes, how uncertainty propagates, and which constraints are binding. Research should therefore emphasize explainability mechanisms, uncertainty quantification, and human-in-the-loop control policies. In parallel, more attention is needed for lifecycle issues such as concept drift, model decay, retraining cadence, and governance of digital twin fidelity over time.

5.4. Final Perspective

The trajectory suggested by the reviewed literature points from isolated optimization models toward continuously learning decision ecosystems. The next wave of contribution

will likely come from systems that couple predictive accuracy with operational accountability: they will forecast better, optimize under competing objectives, and adapt safely under disruption while remaining transparent to human decision-makers. In that direction, the most valuable research is not the one with the most complex standalone algorithm, but the one that demonstrably improves end-to-end supply chain decisions under real constraints and real uncertainty.

Funding: This work was supported by a grant of the Ministry of Research, Innovation and Digitization, CNCS/CCCDI—UEFISCDI, project number ERANET-CHISTERA-IV-REMINDER, within PNCDI IV.

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: No new data were created or analyzed in this study. Data sharing is not applicable to this article.

Conflicts of Interest: The authors declare no conflict of interest.

Abbreviations

The following abbreviations are used in this manuscript:

ABC-XYZ	Inventory/demand segmentation classification
AI	Artificial Intelligence
AI-OR	Artificial Intelligence and Operations Research
ARIMA	AutoRegressive Integrated Moving Average
DL	Deep Learning
ERP	Enterprise Resource Planning
IBP	Integrated Business Planning
IoT	Internet of Things
IT	Information Technology
KPI/KPIs	Key Performance Indicator(s)
LP	Linear Programming
LSTM	Long Short-Term Memory
M5	M5 forecasting competition, commonly referring to the fifth Makridakis competition
MES	Manufacturing Execution System
MIP	Mixed-Integer Programming
ML	Machine Learning
OR	Operations Research
RL	Reinforcement Learning
SAP	Systems, Applications, and Products
S&OP	Sales and Operations Planning
SKU	Stock Keeping Unit
TFT	Temporal Fusion Transformer

References

1. Chopra, S.; Meindl, P. *Supply Chain Management: Strategy, Planning, and Operation*, 7th ed.; Pearson: London, UK, 2019.
2. Hyndman, R.J.; Athanasopoulos, G. *Forecasting: Principles and Practice*, 3rd ed.; OTexts: Melbourne, Australia, 2021.
3. Akbari, M.; Do, T.N.A. A systematic review of machine learning in logistics and supply chain management: Current trends and future directions. *Benchmarking Int. J.* **2021**, *28*, 2977–3009. [[CrossRef](#)]
4. Stranieri, F.; Fadda, E.; Stella, F. Combining deep reinforcement learning and multi-stage stochastic programming to address the supply chain inventory management problem. *Int. J. Prod. Econ.* **2024**, *271*, 109217. [[CrossRef](#)]
5. Wen, Q.; Zhou, T.; Zhang, C.; Chen, W.; Ma, Z.; Yan, J.; Sun, L. Transformers in time series: A survey. *arXiv* **2022**, arXiv:2202.07125.

6. Zhou, H.; Zhang, S.; Peng, J.; Zhang, S.; Li, J.; Xiong, H.; Zhang, W. Informer: Beyond efficient transformer for long sequence time-series forecasting. In *Proceedings of the 35th AAAI Conference on Artificial Intelligence*; AAAI: Palo Alto, CA, USA, 2021; pp. 11106–11115.
7. Ivanov, D.; Dolgui, A. A digital supply chain twin for managing the disruption risks and resilience in the era of Industry 4.0. *Prod. Plan. Control* **2021**, *32*, 775–788. [[CrossRef](#)]
8. Daios, A.; Kladovasilakis, N.; Kelemis, A.; Kostavelis, I. AI Applications in Supply Chain Management: A Survey. *Appl. Sci.* **2025**, *15*, 2775. [[CrossRef](#)]
9. Burgos, D.; Ivanov, D. Food retail supply chain resilience and the COVID-19 pandemic: A digital twin-based impact analysis and improvement directions. *Transp. Res. Part E Logist. Transp. Rev.* **2021**, *152*, 102412. [[CrossRef](#)]
10. Bhandal, R.; Meriton, R.; Kavanagh, R.E.; Bowe, A. The application of digital twin technology in operations and supply chain management: A bibliometric review. *Supply Chain Manag. Int. J.* **2022**, *27*, 182–206. [[CrossRef](#)]
11. Rolf, B.; Jackson, I.; Müller, M.; Lang, S.; Huttman, L.; Schambach, H.; Knoch, S. A review on reinforcement learning algorithms and applications in supply chain management. *Int. J. Prod. Res.* **2023**, *61*, 7962–7988. [[CrossRef](#)]
12. Douaioui, K.; Oucheikh, R.; Benmoussa, O.; Benyoussef, A. Machine learning and deep learning models for demand forecasting in supply chain management: A critical review. *Appl. Syst. Innov.* **2024**, *7*, 93. [[CrossRef](#)]
13. Ivanov, D.; Dolgui, A. OR-methods for coping with the ripple effect in supply chains during COVID-19 pandemic: Managerial insights and research implications. *Int. J. Prod. Econ.* **2021**, *232*, 107921. [[CrossRef](#)]
14. Llaguno, A.; Mula, J.; Campuzano-Bolarin, F. State of the art, conceptual framework and simulation analysis of the ripple effect on supply chains. *Int. J. Prod. Res.* **2021**, *60*, 2044–2066. [[CrossRef](#)]
15. Ivanov, D. Predicting the impacts of epidemic outbreaks on global supply chains: A simulation-based analysis on the coronavirus outbreak (COVID-19/SARS-CoV-2) case. *Transp. Res. Part E Logist. Transp. Rev.* **2020**, *136*, 101922. [[CrossRef](#)] [[PubMed](#)]
16. Li, Y.; Chen, K.; Collignon, S.; Ivanov, D. Ripple effect in the supply chain network: Forward and backward disruption propagation, network health and firm vulnerability. *Eur. J. Oper. Res.* **2021**, *291*, 1117–1131. [[CrossRef](#)]
17. Ruel, S.; El Baz, J.; Ivanov, D.; Das, A. Supply chain viability: Conceptualization, measurement, and nomological validation. *Ann. Oper. Res.* **2024**, *335*, 1107–1136. [[CrossRef](#)]
18. El Baz, J.; Ruel, S. Can supply chain risk management practices mitigate the disruption impacts on supply chains' resilience and robustness? Evidence from an empirical survey in a COVID-19 outbreak era. *Int. J. Prod. Econ.* **2021**, *233*, 107972. [[CrossRef](#)]
19. Frazzon, E.M.; Freitag, M.; Ivanov, D. Intelligent methods and systems for decision-making support: Toward digital supply chain twins. *Int. J. Inf. Manag.* **2021**, *57*, 102281. [[CrossRef](#)]
20. SAP. SAP Integrated Business Planning (Product Page). Available online: <https://www.sap.com/products/scm/integrated-business-planning.html> (accessed on 9 April 2026).
21. Ma, C.; Zhang, L.; You, L.; Tian, W. A Review of Supply Chain Resilience: A Network Modeling Perspective. *Appl. Sci.* **2025**, *15*, 265. [[CrossRef](#)]
22. Leoni, L.; Ferraro, S.; Cantini, A.; Rinaldi, R.; De Carlo, F. MaMoReS: A Robust Tool for Evaluating Supply Chain Resilience Through Systematic Literature and Network Analysis. *Appl. Sci.* **2024**, *14*, 9568. [[CrossRef](#)]
23. dos Santos, C.R.; de Oliveira, U.R.; Aprigliano, V. Supplier Risk in Supply Chain Risk Management: An Updated Conceptual Framework. *Appl. Sci.* **2025**, *15*, 7128. [[CrossRef](#)]
24. Makridakis, S.; Spiliotis, E.; Assimakopoulos, V. Statistical and Machine Learning forecasting methods: Concerns and ways forward. *PLoS ONE* **2018**, *13*, e0194889. [[CrossRef](#)]
25. Lim, B.; Arik, S.O.; Loeff, N.; Pfister, T. Temporal Fusion Transformers for interpretable multi-horizon time series forecasting. *Int. J. Forecast.* **2021**, *37*, 1748–1764. [[CrossRef](#)]
26. Piao, M.; Zhang, D.; Lu, H.; Li, R. A Supply Chain Inventory Management Method for Civil Aircraft Manufacturing Based on Multi-Agent Reinforcement Learning. *Appl. Sci.* **2023**, *13*, 7510. [[CrossRef](#)]
27. anyLogistix. Supply Chain Design and Optimization Software (Product Page). Available online: <https://www.anylogistix.com/> (accessed on 9 April 2026).
28. Blue Yonder. Supply Chain Planning (Product Page). Available online: <https://blueyonder.com/solutions/supply-chain-planning> (accessed on 9 April 2026).
29. o9 Solutions. AI-Driven Supply Chain Planning Software for Enterprises (Product Page). Available online: <https://www.o9solutions.com/solutions/supply-chain-planning/> (accessed on 9 April 2026).
30. Babai, M.Z.; Boylan, J.E.; Rostami-Tabar, B. Demand forecasting in supply chains: A review of aggregation and hierarchical approaches. *Int. J. Prod. Res.* **2022**, *60*, 324–348. [[CrossRef](#)]
31. Aldahmani, E.; Alzubi, A.; Iyiola, K. Demand Forecasting in Supply Chain Using Uni-Regression Deep Approximate Forecasting Model. *Appl. Sci.* **2024**, *14*, 8110. [[CrossRef](#)]
32. Theodorou, E.; Wang, S.; Kang, Y.; Spiliotis, E.; Makridakis, S.; Assimakopoulos, V. Exploring the representativeness of the M5 competition data. *Int. J. Forecast.* **2022**, *38*, 1500–1506. [[CrossRef](#)]

33. Bandara, K.; Hewamalage, H.; Godahewa, R.; Gamakumara, P. A fast and scalable ensemble of global models with long memory and data partitioning for the M5 forecasting competition. *Int. J. Forecast.* **2022**, *38*, 1400–1404. [CrossRef]
34. Amazon Web Services. Amazon Forecast Product Page. Available online: <https://aws.amazon.com/forecast/> (accessed on 9 April 2026).
35. Amazon Web Services. Amazon Forecast Customers Page. Available online: <https://aws.amazon.com/forecast/customers/> (accessed on 9 April 2026).
36. Banerjee, S.; Kankaria, G.H.; Shivaprasad, K.T. From Forecasting Demand to Ordering—An Automated Machine Learning Approach with Amazon Forecast to Decrease Stockouts, Excess Inventory, and Costs. AWS Machine Learning Blog, 2021. Available online: <https://aws.amazon.com/blogs/machine-learning/from-forecasting-demand-to-ordering-an-automated-machine-learning-approach-with-amazon-forecast-to-decrease-stock-outs-excess-inventory-and-costs/> (accessed on 9 April 2026).
37. Nozari, H.; Szmelter-Jarosz, A.; Weiland, D. A Fuzzy Multi-Objective Sustainable and Agile Supply Chain Model Based on Digital Twin and Internet of Things with Adaptive Learning Under Environmental Uncertainty. *Appl. Sci.* **2025**, *15*, 10399. [CrossRef]
38. Dubey, R.; Gunasekaran, A.; Childe, S.J.; Wamba, S.F.; Roubaud, D.; Foropon, C. Empirical investigation of data analytics capability and organizational flexibility as complements to supply chain resilience. *Int. J. Prod. Res.* **2021**, *59*, 110–128. [CrossRef]
39. Zouari, D.; Ruel, S.; Viale, L. Does digitalising the supply chain contribute to its resilience? *Int. J. Phys. Distrib. Logist. Manag.* **2021**, *51*, 149–180. [CrossRef]
40. Yigit, E.; Dag, T. Improving Supply Chain Management Processes Using Smart Contracts in the Ethereum Network Written in Solidity. *Appl. Sci.* **2024**, *14*, 4738. [CrossRef]
41. Choi, T.-M. Risk analysis in logistics systems: A research agenda during and after the COVID-19 pandemic. *Transp. Res. Part E* **2021**, *145*, 102190. [CrossRef]
42. Ivanov, D.; Tang, C.S.; Dolgui, A.; Battini, D.; Das, A. Researchers’ perspectives on Industry 4.0: Multi-disciplinary analysis and opportunities for operations management. *Int. J. Prod. Res.* **2020**, *59*, 2055–2078. [CrossRef]
43. Wamba, S.F.; Queiroz, M.M. Industry 4.0 and the supply chain digitalisation: A blockchain diffusion perspective. *Prod. Plan. Control* **2022**, *33*, 193–210. [CrossRef]

Disclaimer/Publisher’s Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.