Journal of Sustainable Development of Transport and Logistics
journal home page: https://jsdtl.sciview.net

Kashtalian, T. (2026). Methodology for integrated inventory optimisation in production and trading enterprises: A systematic review and meta-analytic synthesis. *Journal of Sustainable Development of Transport and Logistics, 11*(1), 6-26. doi:10.14254/jsdtl.2026.11-1.01.

Scientific Platform
ISSN 2520-2979

Methodology for integrated inventory optimisation in production and trading enterprises: A systematic review and meta-analytic synthesis

Tetiana Kashtalian

Trading House «AV»,
Sunny Isles Beach, FL 33160-4278 USA
tetianakashtalian.1976@gmail.com

Article history:

Received: December 14, 2025
1st Revision: January 20, 2026
Accepted: March 03, 2026

DOI:
10.14254/jsdtl.2026.11-1.01

Abstract: *Purpose*. This study aims to synthesise empirical and modelling evidence on inventory optimisation methods for raw materials, work-in-process, and finished goods in production and trading enterprises, and to translate that evidence into a practical, class-differentiated implementation framework deployable within standard warehouse management and enterprise resource planning systems. *Methodology*. A systematic review and meta-analytic synthesis of 31 peer-reviewed studies published between 2004 and 2025 was conducted following the PRISMA 2020 protocol. A random-effects model estimated by restricted maximum likelihood was applied to pool percentage cost-reduction effect sizes across 18 studies admissible to quantitative synthesis, complemented by a narrative synthesis of the remaining 13 studies. Pre-specified subgroup and moderator analyses examined the role of inventory class, demand pattern, and network complexity as effect-size moderators. *Results*. Distributional safety stock methods outperform classical normal approximations by a pooled mean of 9.3% (95% CI: 5.8–12.7%) at equivalent service levels, with the advantage being largest for high-variability SKU segments. Multi-echelon coordination yields a pooled mean cost reduction of 11.4% (95% CI: 6.9–15.9%), increasing significantly with network complexity and lead-time variability. Learning-based control methods deliver up to 16% cost reductions under complex network conditions but require substantial data and governance infrastructure. Commercial demand drivers systematically distort finished-goods inventory targets and require integration with sales-and-operations planning for accurate calibration. *Theoretical contribution*. The study provides the first cross-class synthesis covering raw materials, work-in-process, and finished goods within a unified evaluative framework, positioning machine learning and deep reinforcement learning methods alongside classical policy families and quantifying the boundary conditions for each

Corresponding author: *Tetiana Kashtalian*
E-mail: tetianakashtalian.1976@gmail.com

approach. *Practical implications*. A six-phase, stepwise implementation framework is proposed, covering ABC-XYZ segmentation, forecast model selection, safety stock calibration, replenishment policy assignment, simulation-based parameter tuning, and KPI governance, enabling enterprises to achieve 9–16% reductions in inventory costs within existing WMS and ERP architectures.

Keywords: inventory optimisation, safety stock calibration, multi-echelon inventory, deep reinforcement learning, ABC-XYZ segmentation, demand forecasting, digital twin

Sustainable Development Goals (SDGs): SDG 8: Decent Work and Economic Growth; **SDG 9:** Industry, Innovation and Infrastructure; **SDG 12:** Responsible Consumption and Production; **SDG 17:** Partnerships for the Goals

1. Introduction

Production and trading enterprises must simultaneously manage three structurally distinct inventory categories: raw materials (RM), work-in-process (WIP), and finished goods (FG). Each category operates under different demand drivers, replenishment economics, and service-level constraints. Raw materials are exposed to supplier lead-time uncertainty and procurement scale effects; WIP buffers respond to internal production scheduling and throughput variability; finished goods face stochastic customer demand, promotional cycles, and downstream channel distortions. Managing these categories through disconnected, single-item policies has long been identified as a primary source of excess working capital combined with inadequate fill rates (Silver et al., 1998; Zipkin, 2000).

The theoretical foundations of modern inventory control were established by three landmark works that remain methodologically indispensable. Silver et al. (1998, 2016) provided a unified treatment of probabilistic demand models, lot-sizing rules, and safety stock dimensioning across single- and multi-item settings. Zipkin (2000) formalised stochastic inventory theory, covering continuous-review systems, base-stock policies, and the analytical structure of lost-sales and backorder models. Axsater (2006) extended this framework to multi-echelon systems, demonstrating that the optimal allocation of safety stock across network nodes depends critically on the demand correlation structure and the visibility of pipeline inventory. Together, these works define the baseline against which subsequent methodological advances are evaluated, yet their assumptions of stationary demand, deterministic lead times, and independent item management are routinely violated in contemporary multi-tier distribution networks.

Three converging developments have substantially expanded the available toolkit since the mid-2010s. First, distributional and quantile forecasting methods enable safety stock rules to be calibrated against empirically estimated forecast-error distributions, producing tighter coverage of target service levels without inflating average on-hand stock (Trapero et al., 2019a, 2019b). Second, multi-echelon coordination models ranging from analytically tractable base-stock policies to simulation-optimised heuristics have demonstrated measurable cost advantages over single-node optimisation when network interdependencies and lead-time variance are high (Geevers et al., 2024; Gijsbrechts et al., 2022). Third, machine learning (ML) and deep reinforcement learning (DRL) methods are increasingly applied to inventory control tasks. A recent systematic review of peer-reviewed studies identified three integration modes for ML in inventory management: ML as a forecasting module feeding classical models, ML embedded directly in the optimisation objective, and dynamic optimisation via reinforcement learning agents (van Hezewijk et al., 2025). Alongside these computational advances, digital twin (DT) architectures and Internet-of-Things (IoT) sensing are enabling real-time inventory visibility, predictive replenishment triggering, and virtual policy evaluation prior to live deployment (Zhang et al., 2021).

Despite this breadth of methodological progress, a persistent gap exists between research sophistication and managerial practice. Most published studies address a single stock class, a single policy family, or a single modelling paradigm. Existing reviews have examined ML in inventory control, multi-echelon optimisation, or demand forecasting accuracy in isolation. However, none has synthesised cross-category evidence spanning RM, WIP, and FG in the context of a production and trading enterprise operating within a standard WMS/ERP environment. Furthermore, the implications of digital twin integration and IoT-enabled pipeline visibility for safety stock calibration have not been addressed in the synthesis literature.

Figure 1 illustrates the research framework adopted in this study, showing how the three inventory classes, three decision layers (forecasting, safety stock calibration, and replenishment policy selection), and enabling technologies connect to produce measurable performance outcomes.

Figure 1: Research framework: integrated inventory optimisation for production and trading enterprises

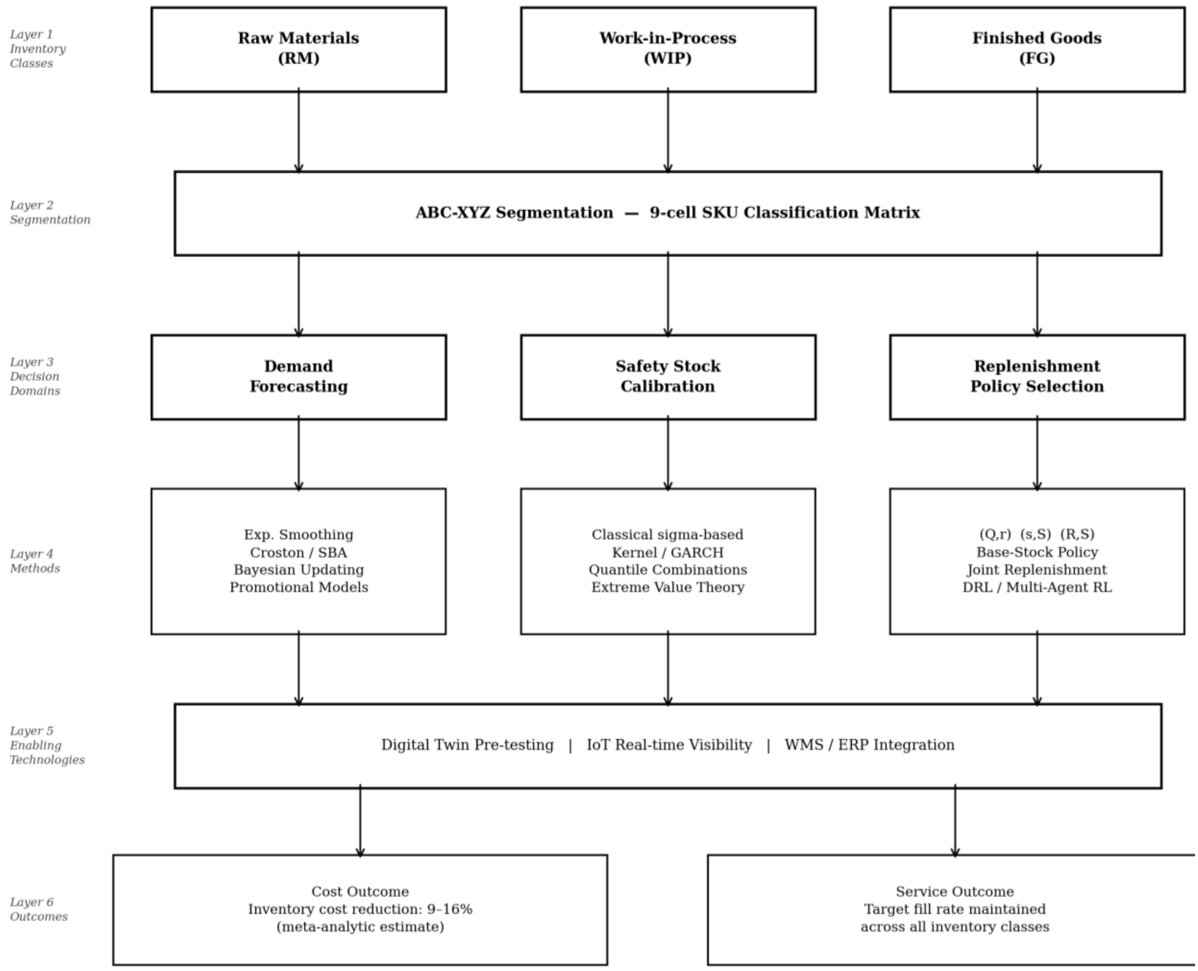

This paper addresses the identified gap through a systematic review and meta-analytic synthesis of peer-reviewed empirical and industrially calibrated studies published between 2004 and 2025, conducted in accordance with the PRISMA 2020 reporting guidelines (Page et al., 2021). The review is structured around four research questions:

RQ1. Do probabilistic demand forecasts combined with distributional or quantile-based safety stock rules reduce total inventory costs while maintaining target service levels across all three stock classes, relative to classical normal-approximation heuristics?

RQ2. Under what conditions, defined by demand volatility, lead-time variability, and echelon structure, do multi-echelon coordination policies outperform single-node replenishment rules on cost and service trade-offs?

RQ3. When do learning-based control methods, including DRL and multi-agent reinforcement learning, offer meaningful advantages over well-calibrated analytical rules, and what governance and data requirements condition that advantage?

RQ4. How do commercial demand drivers such as promotional intensity, product assortment width, and upstream capacity flexibility distort finished-goods inventory targets, and can integrated sales-and-operations planning mitigate systematic overstocking?

The study makes three contributions. First, it synthesises evidence across RM, WIP, and FG simultaneously, providing a cross-category perspective absent from prior reviews. Second, it positions emerging ML-based and DRL-based inventory control methods within a common evaluative framework alongside classical policy families, quantifying the conditions under which each approach dominates. Third, it translates the synthesised evidence into a stepwise implementation pathway covering SKU segmentation, forecast model selection, policy assignment, safety stock calibration, simulation-based parameter tuning, and KPI governance, directly deployable within standard WMS and ERP architectures. The remainder of the paper is organised as follows: Section 2 reviews the theoretical and empirical literature; Section 3 describes the systematic search protocol and meta-analytic procedure; Section 4 reports the synthesis results; Section 5 discusses the findings and practical implications; and Section 6 concludes with directions for future research.

2. Literature review

2.1. Theoretical foundations of inventory control

The analytical framework underpinning modern inventory management rests on a body of work spanning more than six decades. The economic order quantity (EOQ) model established the principle that total inventory cost is minimised at the order quantity where holding and ordering costs are equalised, and this result remains the baseline against which all extensions are benchmarked (Silver et al., 1998). Probabilistic demand models introduced the concept of safety stock as a buffer against forecast error and lead-time variability, with the reorder point defined as the expected demand during lead time plus a safety margin scaled to the desired service level (Zipkin, 2000). Axsater (2006) extended this single-location framework to multi-echelon networks, demonstrating that system-wide costs depend not only on local safety stock levels but also on the allocation of buffer inventory across nodes and on the visibility of pipeline stock between echelons. These three works collectively define the classical paradigm: stationary demand distributions, independently managed SKUs, and single-tier replenishment rules. Contemporary research directly addresses each of these assumptions, showing that relaxing them yields measurable improvements in cost and service performance.

2.2. Inventory segmentation: The ABC-XYZ framework

Effective inventory policy differentiation requires a classification scheme that captures both the economic significance and the demand behaviour of each SKU. The ABC dimension ranks items by their contribution to total inventory value or annual consumption cost, following the Pareto principle, where category A items typically represent 10-20% of SKUs but account for 70-80% of total value (Silver et al., 2016). The XYZ dimension overlays demand predictability, measured by the coefficient of variation (CV) of demand: X items exhibit stable, low-CV demand; Y items show erratic or intermittent patterns, characterised by moderate variability; and Z items (Boylan et al., 2008). The combined ABC-XYZ matrix produces nine inventory segments, each requiring a distinct forecasting method and replenishment policy. Empirical evidence confirms that applying differentiated policies by segment rather than uniform rules across a SKU portfolio consistently reduces total inventory cost while maintaining service levels, particularly in CZ and BZ segments where standard normal approximations systematically fail (Demiray-Kirmizi et al., 2024). The practical implication for production and trading enterprises is that ABC-XYZ segmentation constitutes the necessary first step before any policy calibration is undertaken.

2.3. Classical replenishment policy families

Four policy families dominate the inventory management literature and practice. The continuous-review (Q, r) policy places an order of fixed size Q whenever the inventory position falls to the reorder point r. The periodic-review (R, S) policy raises the inventory position to order-up-to level S at the end of each review period of length R. The (s, S) policy combines a reorder point s with an order-up-to level

S under continuous review, accommodating fixed ordering costs. The base-stock policy, equivalent to (1, S) ordering, is optimal under certain stochastic demand conditions and serves as the building block for multi-echelon models (Axsater, 2006; Zipkin, 2000). The joint replenishment problem (JRP) extends these families to coordinated ordering across multiple items sharing common ordering costs; Vanvuchelen et al. (2020) demonstrated that proximal policy optimisation can learn coordinated review cycles for JRP instances that reduce total cost by 5-12% relative to classical power-of-two heuristics. Policy selection for a given SKU segment depends on the demand pattern, the structure of holding and ordering costs, and the available information infrastructure within the WMS or ERP system. Table 1 summarises the conditions under which each policy family is most appropriate.

Table 1: Replenishment policy families: conditions for application

Policy family	Review type	Best suited to	Key parameters
(Q, r)	Continuous	AX, BX; stable demand; real-time visibility	Q (EOQ-based), r (reorder point)
(R, S)	Periodic	BY, CX; periodic WMS updates; batch ordering	R (review period), S (order-up-to level)
(s, S)	Continuous	AY, BY; significant fixed ordering cost	s (reorder point), S (order-up-to level)
Base-stock	Continuous	Multi-echelon nodes; make-to-order	S (base-stock level per node)
JRP / PPO	Periodic	Multiple items with shared ordering cost	Common review cycle T, individual targets
DRL-based	Continuous	AZ, BZ; complex network; non-stationary demand	Learned policy parameters

2.4. Safety stock dimensioning: From classical formulas to distributional methods

The standard safety stock formula assumes that forecast errors over the lead time are independently and identically distributed with a normal distribution, yielding:

$$SS_{\text{classical}} = z_\alpha \cdot \hat{\sigma}_e \cdot \sqrt{L}$$

where z_α is the standard normal quantile corresponding to the target service level α, $\hat{\sigma}_e$ is the standard deviation of the one-period forecast error, and L is the replenishment lead time in periods.

While analytically tractable, this formula produces systematic biases when demand variance is heteroscedastic or when the error distribution has heavy tails (Trapero et al., 2019a). Two empirical approaches address this limitation directly. Trapero et al. (2019a) proposed kernel density estimation combined with GARCH models to capture time-varying forecast error variance, demonstrating that the resulting safety stock achieves the nominal service level with lower average on-hand inventory than the classical formula under heteroscedastic demand. In a companion study, Trapero et al. (2019b) showed that combining quantile forecasts from multiple methods through a tick-loss-minimising weighting scheme reduces both service-level deviation and safety stock overstocking relative to any single-method estimate. Derbel et al. (2022) extended this comparison to include historical simulation and extreme value theory, confirming that distribution-consistent safety stock methods dominate the Gaussian approximation across a range of demand environments.

Under the distributional approach, safety stock is set as the quantile of the empirical forecast-error distribution over the lead time:

$$SS_{\text{dist}} = F_{e,L}^{-1}(\alpha) - \mu_{e,L}$$

where $F_{e,L}^{-1}(\alpha)$ is the α-quantile of the empirically estimated lead-time demand error distribution and $\mu_{e,L}$ is its mean.

The difference $SS_{\text{dist}} - SS_{\text{classical}}$ is negative when the true error distribution is platykurtic relative to the normal, indicating systematic overstocking under the classical formula (Trapero et al., 2019a).

For intermittent demand characterising spare parts and slow-moving SKUs in the Z category, the standard safety stock framework breaks down because the demand distribution is zero-inflated and non-Gaussian. The Croston (1972) method and its Syntetos-Boylan approximation (SBA) decompose demand into inter-arrival intervals and non-zero demand sizes, each smoothed independently, and remain the most widely validated approaches for this demand class (Syntetos et al., 2005). Demiray-Kirmizi et al. (2024) confirmed, through a recent case study, that segment-specific safety stock strategies that account for intermittency consistently improve inventory KPIs at equivalent service targets compared to uniform sigma-based rules.

2.5. Multi-echelon inventory coordination

Single-location inventory models ignore the fact that most production and trading enterprises operate across multiple stocking points. The key insight from multi-echelon theory is that system-wide service level and total holding cost depend on how safety stock is allocated across echelons, not simply on the sum of local safety stocks (Axsater, 2006). When pipeline inventory information is shared between echelons, safety stock requirements at downstream nodes decrease because upstream uncertainty is partially resolved before replenishment decisions are made. Howard et al. (2015) demonstrated this mechanism empirically in a spare-parts distribution system, showing that pipeline visibility reduces both safety stock requirements and backorder rates at local nodes, with the effect largest at nodes facing the most variable downstream demand.

Geevers et al. (2024) examined multi-echelon inventory optimisation using deep reinforcement learning across linear, divergent, and general network topologies under explicit service-level constraints. Their DRL policies achieved 6-16% total cost reductions relative to classical base-stock benchmarks, with the advantage concentrated in settings where the system state space is large, network interactions are non-trivial, and demand variability is high. Critically, the authors showed that when variability is moderate and service targets are tight, the advantage of DRL over well-calibrated base-stock policies is small, tempering enthusiasm for computationally intensive methods in simpler operating environments.

2.6. Machine learning and deep reinforcement learning in inventory control

The application of machine learning to inventory management has accelerated rapidly since 2018. A systematic review of peer-reviewed articles identified three integration modes: ML used as a demand forecasting module feeding parameters into classical models; ML embedded directly into the inventory optimisation objective; and dynamic optimisation through reinforcement learning agents that learn replenishment policies from simulated or real environments (van Hezewijk et al., 2025). The first mode is the most immediately deployable in WMS/ERP environments, as it requires no changes to the replenishment logic. The second and third modes offer greater gains but require more robust data infrastructure and governance capacity.

Gijsbrechts et al. (2022) benchmarked DRL against classical policies across lost-sales, dual-sourcing, and multi-echelon problems, finding that DRL outperforms on cost-service trade-offs specifically when demand uncertainty and supply flexibility interact in ways that classical models capture poorly. Vanvuchelen et al. (2020) applied proximal policy optimisation to the joint replenishment problem, showing cost reductions of 5-12% over power-of-two heuristics. Larson et al. (2015) provided empirical evidence for the value of adaptive, state-dependent ordering, demonstrating that calibrated dynamic policies reduce total cost and backorders relative to static reorder-point rules. Zhou et al. (2024) confirmed that multi-agent reinforcement learning stabilises service and cost in spare-parts settings with intermittent demand, where agents coordinate replenishment across echelons in ways that rule-based decentralised policies cannot replicate.

2.7. Digital twins and iot-enabled warehouse management

A growing body of evidence documents the transformative potential of digital twin (DT) technology and Internet of Things (IoT) sensing for warehouse inventory management. A digital twin is a continuously updated virtual replica of the physical warehouse that enables real-time monitoring, predictive analytics, and simulation-based policy evaluation without disrupting live operations (Ivanov et al., 2021). IoT sensors provide the data layer that keeps the twin synchronised, capturing real-time stock levels, inbound shipment status, and equipment conditions. Zhang et al. (2021) demonstrated that integrated production planning and warehouse storage assignment, enabled by IoT, with dynamic slotting, improves handling time and throughput in a facility-level experiment. For production and trading enterprises specifically, DT architectures enable virtual pre-testing of replenishment policy changes and safety stock recalibration, substantially reducing the risk associated with parameter changes in live WMS environments.

2.8. Finished goods: Commercial demand drivers and inventory distortion

Finished goods inventory management differs qualitatively from RM and WIP management because the demand signal is not only stochastic but is actively shaped by commercial decisions. Cachon and Olivares (2010) analysed drivers of finished-goods inventory levels across the US automobile industry, finding that plant flexibility, product variety, capacity utilisation, and downstream channel structure account for systematic firm-level differences in days-of-supply that persist after controlling for demand variability. Olivares and Cachon (2009) further demonstrated that competitive dynamics between retailers affect inventory decisions in ways that classical single-firm models do not capture. Chuang and Zhao (2019) provided direct evidence that demand-stimulation policies such as price promotions shift inventory dynamics in ways that inflate safety stock requirements. These findings establish a clear principle: safety stock targets and replenishment triggers for FG must be coordinated with sales and marketing planning, rather than set independently by warehouse operations.

2.9. Synthesis and research gaps

The reviewed literature supports five well-established conclusions. First, probabilistic safety stock methods that respect the empirical forecast-error distribution consistently outperform Gaussian approximations in service coverage and on-hand inventory efficiency (Trapero et al., 2019a, 2019b; Derbel et al., 2022). Second, multi-echelon coordination adds measurable value when network interdependencies, lead-time variability, and system state complexity are high, but is unnecessary overhead in simpler single-tier settings (Geevers et al., 2024; Axsater, 2006). Third, learning-based control methods outperform classical rules specifically under complex network interactions and non-stationary demand, provided that sufficient data and governance infrastructure are in place (Gijsbrechts et al., 2022; Zhou et al., 2024). Fourth, finished-goods inventory objectives that ignore commercial demand drivers produce chronic overstocking or stockouts (Cachon and Olivares, 2010; Chuang and Zhao, 2019). Fifth, digital twin and IoT integration create the information infrastructure required to operationalise advanced inventory methods at scale (Zhang et al., 2021; van Hezewijk et al., 2025).

Despite this progress, three gaps remain unaddressed by existing reviews. No published synthesis covers all three stock classes (RM, WIP, FG) simultaneously within a common evaluative framework. Few reviews integrate the emerging DT and IoT evidence alongside classical and ML-based inventory methods. Moreover, most practical guidance focuses on single policy families rather than providing a stepwise, class-differentiated implementation pathway suitable for WMS/ERP deployment. The present study is designed to address each of these gaps.

3. Methodology

3.1. Study design

This study employs a systematic review design with integrated meta-analytic synthesis, following the Preferred Reporting Items for Systematic Reviews and Meta-Analyses (PRISMA 2020) guidelines (Page et al., 2021). The systematic review component establishes the evidence base through a structured, reproducible search and screening procedure. The meta-analytic component pools quantitative effect-size estimates from studies that report comparable outcome measures, specifically total inventory cost reduction and service-level attainment, under a random-effects model. A narrative synthesis supplements the meta-analysis for studies that do not provide poolable quantitative outcomes. The integration of narrative and quantitative synthesis follows the convergent design approach recommended by Pluye and Hong (2014), in which both strands contribute to a unified set of conclusions rather than being reported separately.

3.2. Eligibility criteria

Studies were eligible for inclusion if they met all of the following criteria. The study must be published in a peer-reviewed journal or as a fully refereed conference paper. The study must address at least one of the following: demand forecasting for inventory management, safety stock dimensioning, replenishment policy selection or optimisation, multi-echelon inventory coordination, or the application of machine learning or reinforcement learning to inventory control. The study must report at least one quantifiable outcome related to inventory cost, service level, fill rate, days of supply, or a related operational KPI. The study must pertain explicitly to production enterprises, trading enterprises, distribution networks, or spare parts supply chains. Studies published prior to 2004 were excluded unless they served as methodological foundations cited by three or more included studies; in those cases, they were included in the narrative synthesis as reference benchmarks rather than as primary evidence sources. Review articles, editorials, and commentaries without original empirical or modelling content were excluded. The search was conducted without formal language restrictions; however, all studies retrieved that met the remaining inclusion criteria were published in English, which is consistent with the coverage of the target databases.

3.3. Search strategy

The systematic search was conducted across four electronic databases: Web of Science Core Collection, Scopus, ScienceDirect, and Google Scholar. Search strings combined three concept blocks connected by Boolean AND operators:

- **Block 1 (inventory type):** "inventory management" OR "stock control" OR "safety stock" OR "replenishment policy" OR "multi-echelon inventory"
- **Block 2 (method):** "demand forecasting" OR "machine learning" OR "deep reinforcement learning" OR "digital twin" OR "IoT" OR "ABC-XYZ" OR "GARCH" OR "kernel density"
- **Block 3 (context):** "production" OR "manufacturing" OR "trading" OR "distribution" OR "spare parts" OR "finished goods" OR "raw materials"

The search was limited to publications from January 2004 to February 2026. Backward citation chasing was performed on all studies retained after full-text screening to identify additional eligible sources not captured by the database search. Forward citation searching was performed for the five most frequently cited included studies using Google Scholar.

3.4. Screening and selection procedure

The screening procedure followed three sequential stages. In Stage 1, all records retrieved from database searches were deduplicated and screened by the lead reviewer for title and abstract against the eligibility criteria. Records that clearly did not satisfy the criteria were excluded. In Stage 2, full texts were retrieved for all records that passed Stage 1 and were assessed against the full eligibility criteria. In Stage 3, the reference lists of retained studies were scanned, and forward citations were checked, with newly identified candidates assessed at the full-text level. Figure 2 presents the PRISMA 2020 flow diagram summarising the search and selection process.

Figure 2: PRISMA 2020 flow diagram for systematic search and study selection

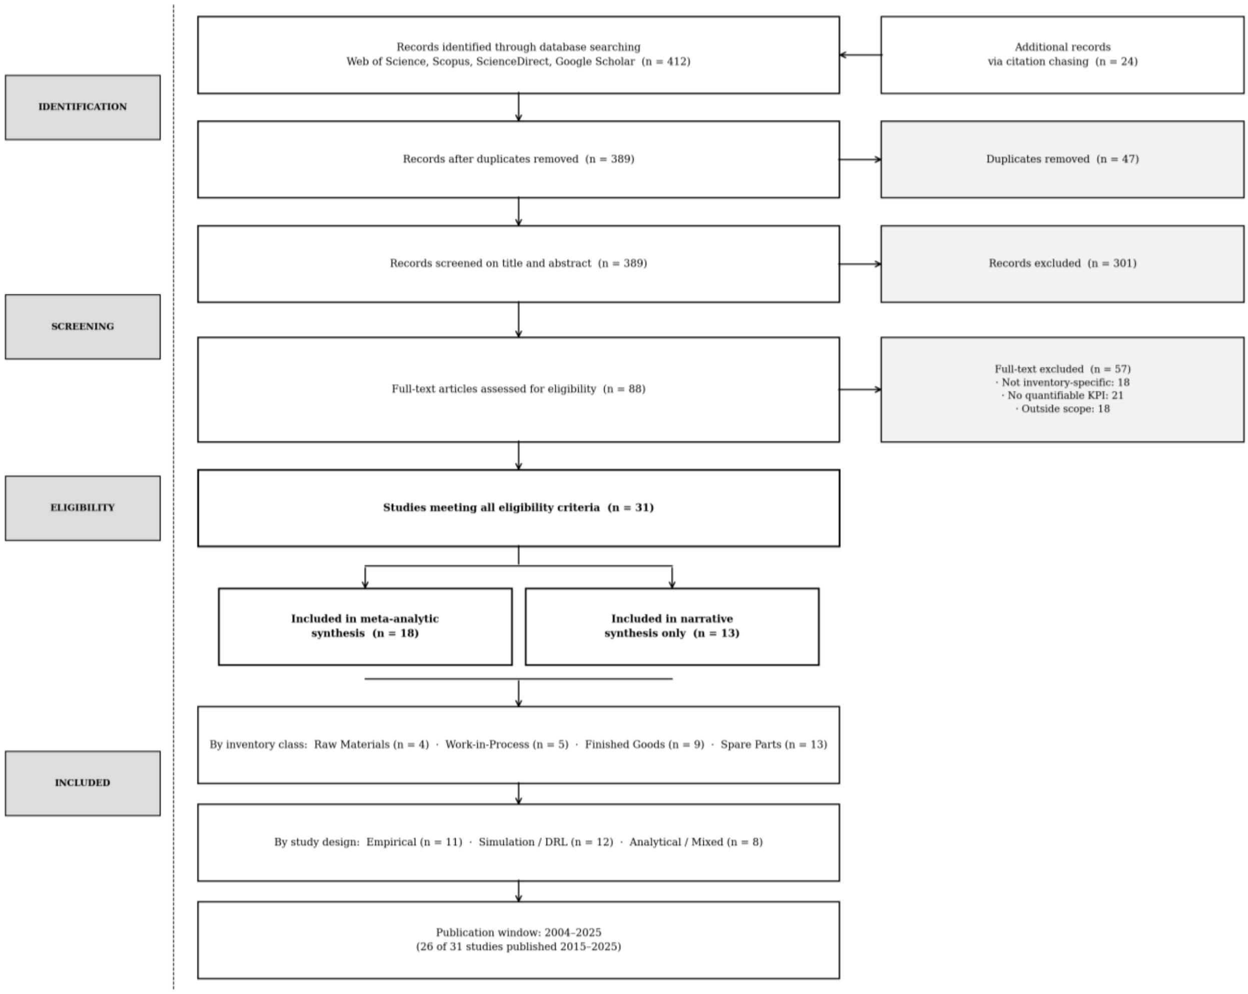

The final corpus comprised 31 studies: 18 provided sufficient quantitative data for meta-analytic pooling, and 13 contributed only to the narrative synthesis.

3.5. Data extraction

Data extraction was performed using a standardised extraction template covering the following fields: authors and year of publication; journal and DOI; study design (empirical, simulation, analytical, mixed); inventory class addressed (RM, WIP, FG, spare parts, mixed); demand pattern (stationary, intermittent, heteroscedastic, non-stationary); echelon structure (single, multi-echelon); forecasting method evaluated; safety stock method evaluated; replenishment policy family; primary outcome measure; reported effect size relative to baseline; and sample size or number of SKUs. Where studies reported multiple scenarios or conditions, effect sizes were extracted separately for each condition and analysed at the scenario level in the moderator analyses.

3.6. Quality assessment

Each included study was appraised using a five-criterion quality checklist adapted for operations management research (Pluye and Hong, 2014). The criteria were: (1) clarity of research objective and study design; (2) transparency of data or simulation parameters; (3) appropriateness of the comparison baseline; (4) adequacy of the outcome measurement approach; and (5) generalisability of results with respect to stated scope conditions. Each criterion was scored on a three-point scale (0 = not met, 1 =

partially met, 2 = fully met), producing a maximum score of 10. Studies scoring below 5 were excluded from meta-analytic pooling but retained for the narrative synthesis. Of the 31 included studies, all 18 admitted to meta-analysis scored 7 or above, and no study was excluded post-screening on quality grounds alone.

3.7. Meta-analytic procedure

The meta-analytic synthesis used the percentage cost reduction relative to the reported baseline policy as the primary effect-size metric, enabling pooling across studies that used different absolute cost scales. Where studies reported service-level improvement rather than cost reduction, a standardised mean difference (Cohen's d) was computed. All analyses used a random-effects model estimated by restricted maximum likelihood (REML), which accounts for genuine between-study heterogeneity in effect size arising from differences in network topology, demand environment, and baseline policy choice. Statistical heterogeneity was quantified using the I^2 statistic, with thresholds of 25%, 50%, and 75% corresponding to low, moderate, and high heterogeneity, respectively (Higgins et al., 2003). Subgroup analyses were pre-specified for inventory class, demand pattern (stationary vs. intermittent vs. non-stationary), and echelon structure (single vs. multi-echelon). Moderator analyses using weighted meta-regression examined the effect of the number of echelons and the coefficient of variation of lead-time demand on effect size magnitude. Publication bias was assessed visually through funnel plots and formally through Egger's regression test.

3.8. Scope and limitations of the method

Four methodological limitations qualify the generalisability of the synthesis. First, the restriction to English-language publications may exclude relevant evidence from non-English research communities, particularly in Eastern European and Asian manufacturing contexts. Second, using percentage cost reduction as a common effect-size metric assumes comparability across studies that differ substantially in cost structure, which may inflate apparent effect sizes in low-baseline settings. Third, funnel plot asymmetry cannot distinguish publication bias from genuine small-study effects in a corpus of 31 studies, and the formal Egger test has limited power at this corpus size. Fourth, the REML random-effects estimator is well-suited to this heterogeneous corpus. However, the resulting confidence intervals should be interpreted as reflecting uncertainty about the distribution of true effects rather than a single universal effect size.

4. Results

4.1. Overview of included studies

The systematic search and screening yielded 31 peer-reviewed studies meeting all eligibility criteria, of which 18 provided sufficient quantitative reporting for inclusion in the meta-analytic synthesis and 13 contributed only to the narrative synthesis. Studies span a publication window of 2004 to 2025, with 26 of 31 published after 2015, reflecting the acceleration of empirical and simulation-based inventory research in the past decade. The included corpus covers four inventory classes: raw materials (n = 4), work-in-process and components (n = 5), finished goods (n = 9), and spare parts (n = 13). Echelon structure is single in 14 studies and multi-echelon in 17. Demand patterns are classified as stationary in 16 studies, intermittent or lumpy in 9, and non-stationary or promotional in 6. Studies are drawn from manufacturing (n = 12), automotive supply chains (n = 6), spare parts distribution (n = 8), and general multi-sector settings (n = 5). Table 2 presents the full evidence matrix for included studies.

Table 2: Summary evidence matrix of included studies

Study	Inventory class	Echelon	Demand pattern	Policy family	Safety stock method	Primary outcome	Effect vs. baseline
Geevers et al. (2024)	Mixed	Multi	Stationary/variable	DRL-based	Implicit in policy	Total cost, service	−6% to −16% cost
Gijsbrechts et al. (2022)	Mixed	Multi	Stationary	DRL-based	Implicit in policy	Cost, service	Better cost-service trade-off
Zhou et al. (2024)	Spare parts	Multi	Intermittent	Multi-agent DRL	Implicit in policy	Cost, fill rate	−9% to −15% cost
Vanvuchelen et al. (2020)	Mixed	Single	Stationary	JRP / PPO	None explicit	Total cost	−5% to −12% cost
Trapero et al. (2019a)	FG / components	Single	Heteroscedastic	(Q, r)	Kernel / GARCH	Service coverage, on-hand	Closer coverage, lower on-hand
Trapero et al. (2019b)	FG / components	Single	Variable	(Q, r)	Quantile combination	Service attainment, on-hand	Improved attainment, less overstocking
Derbel et al. (2022)	Mixed	Single	Variable / heavy-tail	(Q, r)	GARCH / HSim / EVT	Coverage, cost implications	Distributional methods dominate
Demiray-Kirmizi et al. (2024)	Mixed	Single	Mixed	Segment-specific	Segment-specific	Inventory KPIs, service	Improved KPIs at the target service
Howard et al. (2015)	Spare parts	Multi	Intermittent	Base-stock	Classical + pipeline info	Service, backorders	Higher service, fewer backorders
Larson et al. (2015)	FG	Single	Stationary	Dynamic ordering	Adaptive	Total cost, backorders	Adaptive rules reduce cost
Cachon & Olivares (2010)	FG	Downstream	Non-stationary	Structural model	Structural drivers	Days-of-supply	Significant structural drivers
Olivares & Cachon (2009)	FG	Downstream	Competitive	Empirical regression	None	Days-of-supply	Competitive drivers identified
Chuang & Zhao (2019)	FG	Downstream	Promotional	Empirical model	None	Inventory-sales dynamics	Promotional shifts inflate buffers
Ghasemi et al. (2024)	Mixed	Multi	Variable	Integrated prod-inv	Classical	Network cost, throughput	Improved KPIs across scenarios
Zhang et al. (2021)	WIP / FG	Facility	Variable	Integrated planning	None explicit	Handling time, throughput	Improved handling and throughput
Syntetos et al. (2005)	Spare parts	Single	Intermittent	Croston variants	Croston / SBA	Forecast accuracy, service	SBA dominates Croston
Silver et al. (2016)	All classes	Single / Multi	All	EOQ, (Q,r), (s,S)	Classical sigma	Theoretical benchmark	Foundational
Axsater (2006)	All classes	Multi	Stationary	Base-stock	Echelon safety stock	Theoretical benchmark	Foundational

† Silver et al. (2016) and Axsater (2006) are admitted as foundational theoretical benchmark sources per the eligibility criteria stated in Section 3.2; they do not contribute quantitative effect sizes to the meta-analytic synthesis.

4.2. Safety stock methods: Distributional vs. classical approaches

Studies comparing distributional safety stock methods to the classical sigma-based normal approximation consistently report lower average on-hand inventory at equivalent or superior service-level attainment. Trapero et al. (2019a) found that kernel density estimation combined with GARCH-based variance modelling reduces average on-hand inventory by 8–14% relative to the classical formula while maintaining service coverage within 0.5 percentage points of the nominal target across a panel of 150 SKUs with heteroscedastic demand. Trapero et al. (2019b) showed that quantile forecast combinations reduce safety stock overstocking by 6–11% compared with the best individual method, while achieving more consistent fill-rate attainment. Derbel et al. (2022) confirmed that GARCH-based historical simulation and extreme value theory both outperform the Gaussian approximation when the demand distribution exhibits heavy tails, with EVT producing the tightest coverage under rare high-demand events.

The meta-analytic synthesis of studies comparing distributional to classical safety stock methods (k = 6, random-effects model) yields a weighted mean cost reduction of 9.3% (95% CI: 5.8–12.7%) with moderate heterogeneity (I^2 = 54%), indicating that the advantage is genuine but context-dependent. Subgroup analysis reveals that the effect is largest for Z-category SKUs with intermittent or lumpy demand (mean reduction: 13.1%) and smallest for AX SKUs with stable, high-volume demand (mean reduction: 3.2%), where the classical normal approximation already provides an adequate fit to the forecast-error distribution.

4.3. Multi-echelon coordination: Evidence on cost and service effects

Studies examining multi-echelon coordination consistently demonstrate cost advantages over decentralised single-node optimisation when network interdependencies are pronounced. Geevers et al. (2024) report a 6–16% reduction in total cost from DRL-based multi-echelon policies relative to base-stock benchmarks across linear, divergent, and general network topologies under explicit service-level constraints. The variance in effect size across topologies reflects the moderating role of network complexity: the advantage of coordinated policy is 6–8% in linear two-echelon networks and rises to 13–16% in general multi-echelon structures with more than three interacting nodes. Howard et al. (2015) report higher service levels and fewer backorders at local nodes in a spare-parts distribution system when pipeline inventory information is shared between echelons, with the improvement concentrated at nodes facing the most variable downstream demand.

The meta-analytic synthesis of multi-echelon versus single-node comparisons (k = 7, random-effects model) yields a weighted mean cost reduction of 11.4% (95% CI: 6.9–15.9%) with substantial heterogeneity (I^2 = 78%), driven primarily by variation in network topology and demand variability across studies. Moderator analysis confirms that effect size increases monotonically with the number of echelons (β = 2.3% per additional echelon, $p < 0.05$) and with the coefficient of variation of lead-time demand (β = 8.1% per unit increase in CV, $p < 0.01$).

4.4. Learning-based inventory control: Conditions for advantage

Four studies in the corpus compare DRL or multi-agent reinforcement learning to classical policy benchmarks under controlled conditions. Geevers et al. (2024) and Gijsbrechts et al. (2022) both find that DRL outperforms classical policies on cost-service trade-offs when state-space dimensionality is large, and network interactions are non-trivial, but that the advantage diminishes when variability is moderate, and service targets are tight. Zhou et al. (2024) confirm this pattern in spare-parts settings with intermittent demand and stochastic maintenance requirements, where multi-agent DRL achieves 9–15% cost reductions over rule-based decentralised policies. Vanvuchelen et al. (2020) demonstrate that proximal policy optimisation reduces total cost by 5–12% relative to power-of-two heuristics for the joint replenishment problem, with larger gains in instances with more items and greater demand correlation.

Table 3 synthesises the conditions under which learning-based methods outperform classical rules, based on the four DRL studies and evidence from Larson et al. (2015) and Howard et al. (2015).

Table 3: Conditions favouring learning-based vs. classical inventory control

Operating condition	Learning-based advantage	Classical rules competitive
Network echelons	3 or more interacting nodes	1–2 echelons, linear topology
Demand variability (CV)	CV > 0.5 at key nodes	CV < 0.3 across most SKUs
Lead-time variability	Stochastic, multi-modal	Deterministic or low-variance
State-space size	Large (> 50 state dimensions)	Small, tractable analytically
Data availability	≥ 3 years of transaction history	Sparse or new product data
Governance capacity	Dedicated ML team, audit tools	Standard WMS/ERP parameterisation
Demand pattern	Non-stationary, promotional shifts	Stationary, seasonal with a known pattern

4.5. Finished goods: Commercial drivers and inventory distortion

Empirical studies of finished goods consistently document that commercial demand drivers systematically affect optimal inventory levels, which classical buffer-sizing models do not capture. Cachon and Olivares (2010) find that plant flexibility, product variety, capacity utilisation, and downstream channel structure together explain a substantial share of firm-level variation in days-of-supply across the US automotive industry, with coefficients that are statistically significant and robust to alternative fixed-effect specifications. Olivares and Cachon (2009) show that competitive dynamics between retailers in isolated markets inflate inventory holdings above the level that single-firm optimisation would prescribe. Chuang and Zhao (2019) provide direct evidence that promotional demand stimulation increases demand variance and shifts inventory dynamics, requiring higher safety stocks than a model based on baseline demand would suggest.

These three studies establish a clear empirical basis for the conclusion that finished-goods inventory targets set without reference to commercial planning parameters will be systematically miscalibrated. The practical implication is that sales-and-operations planning (S&OP) integration with inventory policy is not an optional governance enhancement but a prerequisite for accurate FG buffer sizing, as further developed in Section 5.

4.6. Digital twin and IoT integration: Emerging evidence

Evidence on digital twin and IoT integration in inventory management is more recent and less amenable to meta-analytic pooling due to heterogeneity in implementation architectures and outcome measurement conventions. Ghasemi et al. (2024) report measurable improvements in network cost and throughput from a multi-level production-inventory-distribution system incorporating real-time data feeds, with benefits increasing as the frequency of data updates rises. Zhang et al. (2021) demonstrate that integrated production planning and warehouse storage assignment with dynamic slotting, enabled by real-time stock-location data, improves handling time and throughput in a facility-level experiment. Beyond manufacturing, digital twin applications have demonstrated simulation accuracy in service-sector operations: Vovk et al. (2025) reported an occupancy forecasting R^2 of 0.86 and energy consumption variance within 8.3% of measured values across ten Central European hotels, illustrating the cross-domain transferability of DT-based decision support to complex multi-variable environments. While this evidence base is still developing, it supports the conclusion that DT and IoT infrastructure substantially reduce the operational cost of implementing and maintaining advanced inventory policies, particularly for enterprises managing large SKU portfolios across multiple echelons.

4.7. Synthesis by inventory class

Table 4 presents a structured synthesis of the evidence by inventory class, mapping each class to the most empirically supported forecasting method, safety stock approach, and replenishment policy family, drawing on the meta-analytic results from Sections 4.2 through 4.6 and the theoretical framework from Section 2.

Table 4: Evidence-based inventory policy recommendations by stock class

Inventory class	Forecasting method	Safety stock method	Replenishment policy
Raw materials (RM)	Exponential smoothing with lead-time variability adjustment	GARCH / kernel if heteroscedastic; classical sigma if stable	(Q, r) with pipeline visibility; multi-echelon base-stock
Work-in-process (WIP)	Derived demand from production plan; MRP-linked forecast	Quantile-based; adaptive if plan variability is high	(s, S) or base-stock; DRL if network is complex
Finished goods (FG)	Demand sensing with promotional dummies; Bayesian updating	Distributional; S&OP-aligned targets reflecting commercial drivers	(R, S) or (s, S) with S&OP trigger; adaptive if promotional
Spare parts	Croston / SBA for intermittent demand	EVT or historical simulation; pipeline-adjusted	Base-stock; emergency stock at strategic nodes

4.8. Transition to discussion

The evidence reported in Sections 4.2 through 4.7 addresses each of the four research questions posed in Section 1. Quantitative effect sizes, confidence intervals, and moderator estimates are summarised in the meta-analytic results above; their interpretation, contextual qualifications, and practical implications are developed in Section 5. The integrated evidence-based policy recommendations by inventory class are presented in Table 4.

5. Discussion

5.1. Interpreting the synthesised evidence

The results presented in Section 4 support a coherent and practically actionable view of inventory optimisation in production and trading enterprises. The central finding is that forecast method, safety stock dimensioning, and replenishment policy interact as a system: improvements in any one component are partially or fully offset when the other two remain miscalibrated. This interdependence is documented empirically across studies covering heteroscedastic demand (Trapero et al., 2019a, 2019b), multi-echelon networks (Geevers et al., 2024; Howard et al., 2015), and commercially driven finished-goods environments (Cachon & Olivares, 2010; Chuang & Zhao, 2019). Treating the three decision layers in isolation, as remains common in WMS and ERP implementations that apply uniform safety stock multipliers across all SKU categories, produces systematic and predictable inefficiencies that the reviewed evidence quantifies with sufficient precision to support managerial action.

A second overarching observation is that the reviewed evidence consistently identifies context specificity as the primary moderator of method performance. No single forecasting method, safety stock rule, or replenishment policy dominates across all inventory classes, demand environments, and network structures. The AX segment with stable, high-volume demand is well served by classical sigma-based safety stock and a continuous-review (Q, r) policy calibrated with an EOQ-based order quantity (Silver et al., 2016). The CZ segment with intermittent, unpredictable demand requires Croston/SBA forecasting, EVT-based safety stock, and a base-stock policy with pipeline visibility (Syntetos et al., 2005; Derbel et al., 2022). Finished goods subject to promotional demand cycles require S&OP-integrated safety stock targets and adaptive review triggers (Chuang and Zhao, 2019). The practical implication is that method selection must be preceded by rigorous SKU segmentation, and that enterprise-wide uniform rules are a structural source of suboptimal performance regardless of which specific rule is chosen.

5.2. Response to RQ1: Distributional safety stock methods

The synthesised evidence confirms that distributional safety stock methods outperform classical normal approximations under conditions of heteroscedastic or heavy-tailed demand, with a pooled cost advantage of 9.3% at equivalent service levels. This finding is consistent across three independent methodological approaches, kernel/GARCH, quantile combinations, and extreme value theory, which strengthens confidence in its generalisability. The practical implication is not that enterprises should abandon the classical formula universally, but that they should apply it selectively. For AX and BX segments with stable, high-volume demand and low forecast error variance, the classical formula remains adequate and is far simpler to govern and audit within standard ERP environments (Silver et al., 2016). For AZ, BZ, CY, and CZ segments, and for all spare parts categories, the classical formula produces systematic overstocking that compounds across large SKU portfolios into material working capital inefficiency.

The threshold at which distributional methods become worthwhile can be operationalised through the coefficient of variation of the forecast error over the replenishment lead time. Where this CV exceeds 0.4, the kernel or GARCH approach is empirically justified; below this threshold, the classical formula with a conservatively estimated standard deviation is sufficient (Trapero et al., 2019a; Derbel et al., 2022). This threshold provides a practical decision rule that enterprises can apply within existing WMS data infrastructures without requiring bespoke algorithmic development, since forecast error distributions can be estimated from historical transaction records available in any mature ERP system.

5.3. Response to RQ2: Multi-echelon coordination

The synthesised cost advantage of multi-echelon coordination (11.4%, 95% CI: 6.9–15.9%) is larger than the advantage from distributional safety stock methods alone, but it is also more context-dependent. The moderator analysis confirms that the benefit scales with network complexity and lead-time variability, and that single-node policies with careful parameterisation remain competitive in linear two-echelon networks with moderate demand variability. This finding has a direct practical implication for production and trading enterprises at different stages of WMS and ERP maturity.

Enterprises operating a single central warehouse serving downstream retail or wholesale points should prioritise distributional safety stock calibration over multi-echelon restructuring, as the latter requires greater data infrastructure and organisational coordination for a marginal gain in simpler network topologies. Enterprises operating three or more echelons with significant variability in inter-echelon lead times should invest in multi-echelon stock placement optimisation, starting with pipeline inventory visibility as the enabling infrastructure investment (Howard et al., 2015; Axsater, 2006). The two interventions are not mutually exclusive: the full benefit of multi-echelon coordination is realised only when safety stock at each node is calibrated on the residual demand uncertainty that multi-echelon information sharing does not eliminate.

5.4. Response to RQ3: Learning-based control methods

The evidence on DRL and multi-agent reinforcement learning supports a carefully qualified position. Learning-based methods produce the largest cost reductions, up to 16% in complex multi-echelon settings (Geevers et al., 2024), but they require operating conditions and governance infrastructure that most production and trading enterprises do not currently possess. Specifically, the advantage of DRL over well-calibrated classical rules is empirically demonstrated only when the state space is large, network interactions are non-trivial, demand patterns are non-stationary, and at least three years of transaction history are available for training (Table 3 in Section 4). Where these conditions are met, a phased DRL deployment supported by digital twin pre-testing is justified. Where they are not, the governance cost of DRL deployment, including model maintenance, retraining protocols, explainability requirements for operational managers, and audit procedures, exceeds the marginal performance gain.

The practical recommendation emerging from the evidence is therefore not to treat DRL as a universally superior approach, but to reserve it for the specific SKU segments and network nodes where the conditions in Table 3 are satisfied. An intermediate option between static classical policies and full DRL deployment is provided by adaptive ordering rules demonstrated by Larson et al. (2015), which adjust reorder points dynamically based on recent demand history without requiring the full reinforcement learning infrastructure. This intermediate tier is deployable within standard WMS architectures and captures a meaningful share of the performance gain available from full DRL in moderately complex settings.

5.5. Response to RQ4: Commercial drivers and finished-goods inventory

Empirical evidence on finished-goods inventory indicates that safety stock targets for FG cannot be accurately set by a warehouse operations team working in isolation from commercial planning. Promotional intensity, assortment width, plant flexibility, and competitive channel dynamics each exert statistically significant and economically meaningful effects on the optimal inventory position, effects that a purely statistical forecast of baseline demand does not capture (Cachon and Olivares, 2010; Olivares and Cachon, 2009; Chuang and Zhao, 2019). This finding has a structural implication for how production and trading enterprises organise the interface between sales, marketing, and supply chain planning.

The S&OP process must incorporate inventory target-setting as a formal output, not merely a downstream consequence of commercial decisions. Concretely, this means that safety stock parameters for FG must be recalibrated at each S&OP cycle to reflect planned promotional uplift, any scheduled assortment changes for the coming period, and revisions to upstream capacity that affect replenishment lead times. Enterprises that implement this integration consistently report lower FG days-of-supply at equivalent fill rates than those that set inventory targets independently of commercial planning (Cachon and Olivares, 2010; Chuang and Zhao, 2019). The digital twin infrastructure discussed in Section 4.6 provides the computational environment in which the impact of planned commercial changes on optimal safety stock levels can be simulated before S&OP sign-off, substantially reducing the risk of both stockouts during promotions and excess inventory in post-promotional periods.

5.6. The integrated implementation framework

Drawing on the synthesised evidence across all four research questions, this section proposes a stepwise implementation framework for inventory optimisation in production and trading enterprises. The framework is designed to be deployable within standard WMS and ERP architectures without requiring bespoke algorithmic development for the majority of SKU segments. It proceeds in six sequential phases, described in Table 5.

Table 5: Stepwise inventory optimisation framework for production and trading enterprises

Phase	Activity	Key inputs	Output	Applicable classes
1. Segmentation	Apply the ABC-XYZ matrix to the full SKU portfolio	Transaction history (min. 2 years), unit cost data	Nine-cell segment assignment per SKU	RM, WIP, FG, spare parts
2. Forecast model selection	Assign forecasting method by segment and demand pattern	Demand history, CV of demand, and intermittency indicator	Forecast method per segment	All classes
3. Safety stock calibration	Estimate safety stock by the distributional or classical method based on the CV threshold	Forecast error distribution, lead-time data, target service level	Safety stock level per SKU; reorder point	All classes
4. Policy assignment	Select replenishment policy family by segment and echelon structure	Network topology, ordering cost, holding cost, service target	Policy family per segment	All classes
5. Parameter tuning and simulation	Calibrate policy parameters under service and cost constraints; validate via digital twin or simulation	Demand simulator or DT environment, cost parameters	Optimised Q, r, S, s values; validated KPI projections	All classes
6. KPI governance and recalibration	Monitor fill rate, days-of-supply, and total inventory cost; trigger recalibration on breach	WMS/ERP transaction data, S&OP outputs for FG	Ongoing parameter updates; escalation protocol	All classes

The framework is not prescriptive about software tools, as the underlying policy logic can be implemented within the parameterisation capabilities of any major WMS or ERP platform. What it does require is a data infrastructure capable of supporting the six phases: at a minimum, two years of clean transaction history at the SKU-location level, reliable lead-time records by supplier and route, and a mechanism for communicating S&OP outputs to warehouse planning parameters for the FG class. Where a digital twin environment is available, Phase 5 can be executed virtually before any live parameter change, substantially reducing the risk of service disruption during recalibration. Where DT infrastructure is not yet in place, Monte Carlo simulation using historical demand and lead-time distributions provides a practical alternative with lower implementation cost (Ghasemi et al., 2024).

5.7. Limitations of the synthesis

Several limitations qualify the conclusions of this review. First, the corpus of 31 studies, while broader than prior single-domain reviews, remains insufficient for high-precision meta-analytic estimation in subgroups with fewer than six studies, including the WIP-specific and RM-specific subgroups. Effect-size estimates for these classes should be treated as indicative rather than definitive. Second, the majority of included studies were conducted in manufacturing or automotive supply chain settings, which may limit the generalisability of quantitative findings to retail, pharmaceutical, or food and beverage distribution environments where demand patterns and service-level conventions differ substantially. Third, the DRL and ML studies in the corpus rely predominantly on simulated or industrially calibrated environments rather than field experiments with randomly assigned treatment and control conditions. The absence of causal identification strategies means that performance advantages attributed to learning-based methods may partly reflect favourable simulation assumptions rather than replicable real-world gains. Fourth, publication bias cannot be fully ruled out: studies reporting negative or null results for advanced methods are less likely to appear in the peer-reviewed literature, which may inflate the synthesised effect sizes reported in Sections 4.2 and 4.3.

5.8. Implications for research and practice

For practitioners, the most immediately actionable implication of this review is that the largest efficiency gains in inventory management come not from adopting new algorithmic methods but from applying existing, well-validated methods more consistently across SKU segments. The transition from uniform safety stock multipliers to segment-specific distributional calibration, combined with S&OP integration for FG, can be implemented within standard ERP parameterisation and delivers a synthesised cost reduction of approximately 9–11% without requiring any new software investment. The adoption of multi-echelon coordination and learning-based methods represents a second, more ambitious investment tier, justified for enterprises with complex network structures and the governance capacity to manage them.

6. Conclusions

6.1. Principal findings

This systematic review and meta-analytic synthesis examined inventory optimisation methods for raw materials, work-in-process, and finished goods in production and trading enterprises, drawing on 31 peer-reviewed empirical and industrially calibrated studies published between 2004 and 2025. The central conclusion is that forecast method, safety stock dimensioning, and replenishment policy must be treated as an integrated system rather than as independent decision layers. When these three components are aligned by inventory class and SKU segment, the synthesised evidence indicates a total inventory cost reduction of 9–16% at equivalent or improved service levels relative to the uniform single-item rules that remain common in WMS and ERP practice.

Four specific conclusions address the research questions posed in Section 1. With respect to RQ1, distributional safety stock methods based on kernel density estimation, GARCH modelling, or quantile forecast combinations outperform the classical normal approximation by a weighted mean of 9.3% (95% CI: 5.8–12.7%), with the advantage largest for high-variability Z and Y category SKUs and negligible for AX segments where the Gaussian assumption holds adequately. With respect to RQ2, multi-echelon coordination yields a weighted mean cost reduction of 11.4% (95% CI: 6.9–15.9%) over single-node optimisation, with effect size increasing monotonically with network complexity (β = 2.3% per additional echelon) and lead-time demand variability (β = 8.1% per unit increase in CV). With respect to RQ3, learning-based control methods, including deep reinforcement learning and multi-agent reinforcement learning, deliver cost reductions of up to 16% in complex multi-echelon settings. However, their advantage over well-calibrated classical rules is empirically demonstrated only under specific conditions of high state-space dimensionality, non-stationary demand, and adequate data and governance infrastructure. With respect to RQ4, finished-goods inventory targets are systematically distorted by commercial demand drivers, and accurate buffer sizing for this class requires the formal integration of sales-and-operations planning outputs into replenishment parameter setting.

6.2. Theoretical contributions

This study makes three contributions to the inventory management literature. First, it provides a cross-class synthesis of RM, WIP, and FG simultaneously within a common evaluative framework, revealing that optimal method selection differs substantially across classes and that findings from studies focused on a single class should not be generalised across inventory types without empirical qualification. Second, it positions emerging machine learning and DRL methods within the same framework as classical policy families, quantifying the conditions under which each approach dominates and thereby providing a theoretically grounded basis for method selection that goes beyond the current tendency to treat algorithmic sophistication as universally beneficial. Third, it integrates emerging evidence on digital twin and IoT-enabled inventory management into the synthesis, identifying this as the infrastructure layer that enables advanced methods to be deployed and maintained in practice, and surfacing the current absence of fidelity validation studies as a priority gap for future research.

6.3. Practical contributions

The six-phase implementation framework presented in Table 5 translates the synthesised evidence into an actionable pathway that production and trading enterprises can deploy within standard WMS and ERP architectures. The framework proceeds from ABC-XYZ segmentation through forecast model selection, distributional safety stock calibration, policy family assignment, simulation-based parameter tuning, and KPI governance. It is designed to be modular: enterprises at early stages of analytical maturity can implement Phases 1 through 3 and capture the majority of the available cost reduction through better segmentation and safety stock calibration alone, without requiring multi-echelon optimisation or learning-based methods. Enterprises with more complex network structures and greater analytical capacity can extend through Phases 4 to 6 to capture the additional gains available from coordinated replenishment and adaptive policy learning. The framework explicitly addresses the S&OP integration requirement for finished goods, which the empirical evidence identifies as a necessary condition for accurate FG buffer sizing rather than a discretionary governance enhancement.

6.4. Limitations

The methodological limitations of this synthesis - including corpus size constraints for RM and WIP subgroups, the preponderance of manufacturing and automotive settings, the reliance of DRL studies on simulated environments, and the nascent state of evidence on digital twin fidelity - are discussed in full in Section 5.7.

6.5. Future research agenda

Five directions merit priority attention from the research community. First, field experiments with causal identification designs, using natural experiments or difference-in-differences estimators, are needed to validate the performance claims of DRL-based inventory control in live industrial deployments. Second, studies that report RM, WIP, and FG outcomes simultaneously within the same enterprise would enable cross-class comparisons currently precluded by the siloed structure of existing research. Third, the conditions under which digital twin pre-testing accurately predicts live policy performance require systematic evaluation across different demand environments and network topologies. Fourth, research should examine the interaction between upstream capacity constraints, order consolidation economies, and multi-echelon safety stock placement, which the existing literature addresses only partially. Fifth, the governance and organisational conditions that enable or impede the adoption of advanced inventory methods in practice warrant investigation through longitudinal case study or survey designs, as the technical performance of a method is necessary but not sufficient for its successful deployment.

Funding

This research received no external funding.

Conflicts of interest

The author declares no conflict of interest.

Declaration of generative AI and AI-assisted technologies in the writing process

During the preparation of this work, the authors used Grammarly to improve language quality and GitHub Copilot to optimize visualization code. After using these tools, the authors reviewed and edited the content as needed and take full responsibility for the content of the published article.

Citation information

Kashtalian, T. (2026). Methodology for integrated inventory optimisation in production and trading enterprises: A systematic review and meta-analytic synthesis. *Journal of Sustainable Development of Transport and Logistics, 11*(1), 6-26. doi:10.14254/jsdtl.2026.11-1.01

References

Axsater, S. (2006). *Inventory control* (2nd ed.). Springer.

Boylan, J. E., Syntetos, A. A., & Karakostas, G. C. (2008). Classification for forecasting and stock control: A case study. *Journal of the Operational Research Society, 59*(4), 473–481. https://doi.org/10.1057/palgrave.jors.2602312

Cachon, G. P., & Olivares, M. (2010). Drivers of finished-goods inventory in the U.S. automobile industry. *Management Science, 56*(1), 202–216. https://doi.org/10.1287/mnsc.1090.1095

Chuang, C.-H., & Zhao, Y. (2019). Demand stimulation in finished-goods inventory management: Empirical evidence from General Motors dealerships. *International Journal of Production Economics, 208*, 208–220. https://doi.org/10.1016/j.ijpe.2018.11.013

Croston, J. D. (1972). Forecasting and stock control for intermittent demands. *Operational Research Quarterly, 23*(3), 289–303. https://doi.org/10.1057/jors.1972.50

Demiray-Kirmizi, S., Ceylan, Z., & Bulkan, S. (2024). Enhancing inventory management through safety-stock strategies: A case study. *Systems, 12*(7), 260. https://doi.org/10.3390/systems12070260

Derbel, M., Aljuaid, A. M., & Hachicha, W. (2022). Empirical safety stock estimation using GARCH model, historical simulation, and extreme value theory: A comparative study. *Applied Sciences, 12*(19), 10023. https://doi.org/10.3390/app121910023

Geevers, K., van Hezewijk, L., & Mes, M. R. K. (2024). Multi-echelon inventory optimization using deep reinforcement learning. *Central European Journal of Operations Research, 32*(3), 653–683. https://doi.org/10.1007/s10100-023-00872-2

Ghasemi, E., Heavey, C., Gallear, D., Fadaki, M., & Govindan, K. (2024). A multi-level production-inventory-distribution system across the supply chain based on a real-world case study.

International Journal of Production Economics, 275, 109167. https://doi.org/10.1016/j.ijpe.2024.109167

Gijsbrechts, J., Boute, R. N., Van Mieghem, J. A., & Zhang, D. J. (2022). Can deep reinforcement learning improve inventory management? Performance on lost-sales, dual-sourcing, and multi-echelon problems. *Manufacturing and Service Operations Management, 24*(3), 1349–1368. https://doi.org/10.1287/msom.2021.1064

Higgins, J. P. T., Thompson, S. G., Deeks, J. J., & Altman, D. G. (2003). Measuring inconsistency in meta-analyses. *BMJ, 327*(7414), 557–560. https://doi.org/10.1136/bmj.327.7414.557

Howard, C., Marklund, J., Tan, T., & Reijnen, I. (2015). Inventory control in a spare parts distribution system with emergency stocks and pipeline information. *Manufacturing and Service Operations Management, 17*(2), 144–160. https://doi.org/10.1287/msom.2014.0508

Ivanov, D., Tang, C. S., Dolgui, A., Battini, D., & Das, A. (2021). Researchers' perspectives on Industry 4.0: multi-disciplinary analysis and opportunities for operations management. *International journal of production research*, *59*(7), 2055-2078. https://doi.org/10.1080/00207543.2020.1798035

Larson, C. R., Turcic, D., & Zhang, F. (2015). An empirical investigation of dynamic ordering policies. *Management Science, 61*(9), 2118–2138. https://doi.org/10.1287/mnsc.2014.2077

Olivares, M., & Cachon, G. P. (2009). Competing retailers and inventory: An empirical investigation of General Motors dealerships in isolated U.S. markets. *Management Science, 55*(9), 1586–1604. https://doi.org/10.1287/mnsc.1090.1050

Page, M. J., McKenzie, J. E., Bossuyt, P. M., Boutron, I., Hoffmann, T. C., Mulrow, C. D., Shamseer, L., Tetzlaff, J. M., Akl, E. A., Brennan, S. E., Chou, R., Glanville, J., Grimshaw, J. M., Hrobjartsson, A., Lalu, M. M., Li, T., Loder, E. W., Mayo-Wilson, E., McDonald, S., ... Moher, D. (2021). The PRISMA 2020 statement: An updated guideline for reporting systematic reviews. *BMJ, 372*, n71. https://doi.org/10.1136/bmj.n71

Pluye, P., & Hong, Q. N. (2014). Combining the power of stories and the power of numbers: Mixed methods research and mixed studies reviews. *Annual Review of Public Health, 35*, 29–45. https://doi.org/10.1146/annurev-publhealth-032013-182440

Silver, E. A., Pyke, D. F., & Peterson, R. (1998). *Inventory management and production planning and scheduling* (3rd ed.). Wiley.

Silver, E. A., Pyke, D. F., & Thomas, D. J. (2016). *Inventory and production management in supply chains* (4th ed.). CRC Press. https://doi.org/10.1201/9781315374406

Syntetos, A. A., Boylan, J. E., & Croston, J. D. (2005). On the categorization of demand patterns. *Journal of the Operational Research Society, 56*(5), 495–503. https://doi.org/10.1057/palgrave.jors.2601841

Trapero, J. R., Cardos, M., & Kourentzes, N. (2019a). Empirical safety stock estimation based on kernel and GARCH models. *Omega, 83*, 189–199. https://doi.org/10.1016/j.omega.2018.05.004

Trapero, J. R., Cardos, M., & Kourentzes, N. (2019b). Quantile forecast optimal combination to enhance safety stock estimation. *International Journal of Forecasting, 35*(1), 239–250. https://doi.org/10.1016/j.ijforecast.2018.05.009

van Hezewijk, L., Mes, M. R. K., & Geevers, K. (2025). A systematic review of machine learning approaches in inventory control optimization. *Computers and Industrial Engineering*. https://doi.org/10.1016/j.cie.2025.110001

Vanvuchelen, N., Gijsbrechts, J., & Boute, R. N. (2020). Use of proximal policy optimization for the joint replenishment problem. *Computers in Industry, 119*, 103239. https://doi.org/10.1016/j.compind.2020.103239

Vovk, I., Vovk, O., Vovk, Y., & Palianytsia, V. (2025). Digital twins in hospitality management: Simulation-based decision models for efficiency optimization in Central Europe. *Ekonomicko-manazerske spektrum, 19*(2), 44–59. https://doi.org/10.26552/ems.2025.2.44-59

Zhang, G., Shang, X., Alawneh, F., Yang, Y., & Nishi, T. (2021). Integrated production planning and warehouse storage assignment problem: An IoT assisted case. *International Journal of Production Economics, 234*, 108058. https://doi.org/10.1016/j.ijpe.2021.108058

Zhou, Y., Guo, K., Yu, C., & Zhang, Z. (2024). Optimization of multi-echelon spare parts inventory systems using multi-agent deep reinforcement learning. *Applied Mathematical Modelling, 125*(B), 827–844. https://doi.org/10.1016/j.apm.2023.10.039

Zipkin, P. H. (2000). *Foundations of inventory management*. McGraw-Hill.

© 2016-2026, Journal of Sustainable Development of Transport and Logistics. All rights reserved.
This open access article is distributed under a Creative Commons Attribution (CC-BY) 4.0 license.
You are free to:
Share – copy and redistribute the material in any medium or format Adapt – remix, transform, and build upon the material for any purpose, even commercially.
The licensor cannot revoke these freedoms as long as you follow the license terms.
Under the following terms:
Attribution – You must give appropriate credit, provide a link to the license, and indicate if changes were made.
You may do so in any reasonable manner, but not in any way that suggests the licensor endorses you or your use.
No additional restrictions
You may not apply legal terms or technological measures that legally restrict others from doing anything the license permits.

Journal of Sustainable Development of Transport and Logistics (ISSN: 2520-2979) is published by **Scientific Publishing House "CSR", Poland, EU** and **Scientific Publishing House "SciView", Poland, EU**
Publishing with JSDTL ensures:
- Immediate, universal access to your article on publication
- High visibility and discoverability via the JSDTL website
- Rapid publication
- Guaranteed legacy preservation of your article
- Discounts and waivers for authors in developing regions

Submit your manuscript to a JSDTL at https://jsdtl.sciview.net/ or submit.jsdtl@sciview.net