



Contents lists available at IJIECM
 International Journal of Industrial Engineering and Construction
 Management
 Journal Homepage: <http://www.ijiecm.com/>
 Volume 5, No. 1, 2025

IJIECM
 INTERNATIONAL JOURNAL OF
 INDUSTRIAL ENGINEERING
 AND CONSTRUCTION MANAGEMEN

Artificial Intelligence and Data Analytics for Intelligent Operational Decision-Making

Ali Mohammady¹, Ali Daghighi², Seyedkian Rezvanjou³

¹Chapman University, Argyros College of Business & Economics, Orange, California, USA, 92866

²Faculty of Engineering and Natural Sciences, Biruni University, Istanbul, Turkey

³Department of Engineering, California State University East Bay, Hayward, California, USA, 94542

ABSTRACT

The increasing availability of operational data and the growing uncertainty of modern business environments are transforming how organizations plan and make decisions. Traditional deterministic or experience-based approaches often struggle to capture nonlinear patterns or respond effectively to volatile demand, disruptions, and resource constraints. This study presents an integrated decision intelligence framework that combines artificial intelligence (AI), data analytics, and optimization under uncertainty to support intelligent operational decision-making. The framework consists of three complementary components: (i) predictive analytics, which applies machine learning models to forecast key variables such as demand, lead times, and disruption risks using historical and contextual data; (ii) prescriptive optimization, which incorporates these forecasts into stochastic or robust optimization models to generate cost-efficient and service-driven decisions under uncertainty; and (iii) data-driven process improvement, which uses continuous monitoring and feedback mechanisms to refine predictive models and operational policies over time. A representative operational scenario with uncertain, time-varying demand is used to evaluate the framework against conventional deterministic and heuristic methods. Results demonstrate that the integrated approach improves service levels, reduces total operational cost, and enhances robustness to variability. The proposed framework offers a generalizable foundation for embedding AI-driven analytics into operations research models, supporting adaptive, transparent, and evidence-based decision-making in data-rich environments.

ARTICLE INFO

Received: 2025/10/12

Revised: 2025/11/08

Accepted: 2025/11/19

Keywords:

Artificial Intelligence; Data Analytics; Decision Intelligence; Predictive Analytics; Prescriptive Analytics; Optimization under Uncertainty; Machine Learning; Operations Research; Data-Driven Process Improvement; Intelligent Operations Management; Digital Transformation

1. Introduction

Modern industrial and service organizations operate in environments that are increasingly volatile, complex, and data-rich. Demand patterns fluctuate across products and regions, supply disruptions occur with little warning, and capacity constraints must be managed

under stringent cost and service-level targets. At the same time, advances in information systems and sensing technologies have enabled the continuous collection of large volumes of operational data from enterprise resource planning (ERP) systems, manufacturing execution systems, logistics platforms, and customer-facing channels.

These developments have created both an opportunity and a challenge: while abundant data can, in principle, support more informed and proactive decision-making, many organizations still rely on deterministic planning tools, spreadsheet-based analyses, or experience-driven heuristics that are poorly suited to highly uncertain and dynamic environments.

Artificial Intelligence (AI) and data analytics offer powerful mechanisms for extracting value from operational data. Predictive models based on machine learning can uncover nonlinear relationships between explanatory factors and key performance indicators such as demand, lead times, failure probabilities, or service times. These models enable organizations to anticipate future conditions more accurately than traditional time-series or judgmental forecasting alone. In parallel, the field of operations research has developed a rich body of methods for optimization under uncertainty, including stochastic programming, robust optimization, and reinforcement learning. These methods support the design of decision policies that explicitly account for risk and variability in the underlying system. However, in practice, predictive analytics and optimization models are often developed and deployed in isolation. Forecasting systems generate outputs that are not tightly coupled to downstream decision models, while optimization tools frequently rely on simplified or static assumptions that ignore the full information content of available data.

This misalignment leads to several limitations. First, prediction models are typically evaluated with statistical metrics (e.g., mean absolute percentage error) that may not reflect their impact on operational performance, such as cost or service level. Second, optimization models may treat uncertain parameters as fixed point estimates or use ad-hoc safety factors, which can result in solutions that are fragile when confronted with real-world variability. Third, operational decision processes rarely incorporate systematic feedback mechanisms: once a plan is generated and executed, discrepancies between predicted and realized outcomes are not formally used to update either the predictive models or the decision rules. As a consequence, organizations struggle to close the loop between data, models, and decisions, and the potential of AI and analytics to improve operations remains only partially realized.

In response to these challenges, the concept of *decision intelligence* has emerged as a unifying paradigm that integrates data engineering, analytics, and decision science into coherent architectures. Decision intelligence emphasizes the design of end-to-end pipelines in which data collection, model development, optimization, and human oversight are treated as interconnected components of a single decision-making system rather than isolated technical artifacts. Despite growing interest in this concept, there is still a lack of concrete

methodological frameworks that show how to embed AI-driven analytics into operational decision models in a way that is both technically rigorous and practically deployable.

Innovation and research objective. This paper introduces an integrated decision intelligence framework that combines artificial intelligence, data analytics, and optimization under uncertainty to support intelligent operational decision-making. The innovation of the proposed framework lies in three key aspects. First, it adopts a *predict-optimize-improve* architecture that explicitly links predictive analytics to prescriptive optimization and embeds them within a data-driven process improvement loop. Rather than treating prediction and optimization as sequential but independent tasks, the framework uses predictive models to generate distributional scenarios and risk measures that directly inform decision variables and constraints. Second, the framework is designed to be *adaptive*: model parameters and decision policies are updated over time using streaming data and performance feedback, allowing the system to learn from realized outcomes and structural changes in the environment. Third, the framework is *transparent and modular*, providing clearly defined interfaces between data, models, and decisions. This design facilitates implementation within existing information systems and supports managerial interpretation of the resulting policies.

To demonstrate the applicability and benefits of the proposed approach, we instantiate the framework in a representative operational context characterized by multiple products, uncertain and time-varying demand, and finite production or procurement capacity. In this setting, machine learning models are used to generate probabilistic demand forecasts that capture both point predictions and uncertainty measures. These forecasts are then embedded in a stochastic or robust optimization model that determines production or replenishment quantities with the objective of minimizing total cost while satisfying service-level requirements. A feedback mechanism compares planned and realized performance, updating both the forecasting models and selected decision parameters. This case study allows us to quantify the performance gains of the integrated framework relative to conventional deterministic planning and simple heuristic rules.

Contributions. The main contributions of this paper are fourfold:

1. We propose a generalizable *AI-enabled decision intelligence framework* for operational decision-making under uncertainty that tightly couples predictive analytics, prescriptive optimization, and data-driven process improvement.
2. We develop a concrete *predict-optimize-improve*

pipeline in which machine learning-based predictive models produce probabilistic inputs that are directly incorporated into stochastic and robust optimization models, ensuring consistency between forecasting assumptions and decision structures.

3. We introduce a *feedback-based learning mechanism* that uses realized operational outcomes to iteratively update both the predictive models and key decision parameters, thereby enhancing adaptability and robustness over time.
4. Through a detailed numerical case study, we provide an *empirical evaluation* that compares the proposed framework with traditional deterministic and heuristic approaches, demonstrating improvements in total cost, service level attainment, and resilience to variability, and offering managerial insights for practical implementation.

Organization of the paper. The remainder of this paper is structured as follows. Section 2. reviews existing literature on predictive analytics in operations, optimization under uncertainty, and integrated analytics–optimization frameworks. Section 3. presents the proposed decision intelligence framework, detailing the predictive modeling, optimization formulation, and feedback mechanisms. Section ?? describes the case study design, data generation and preprocessing, experimental scenarios, and performance metrics, and reports the numerical results. Section 6. discusses the implications of the findings for researchers and practitioners, including design guidelines and implementation considerations. Finally, Section ?? concludes the paper and outlines promising directions for future research.

2. Related Work

The integration of artificial intelligence, data analytics, and optimization into operational decision-making has attracted significant research attention in recent years. The rapid growth of computational power, cloud-native platforms, and industry-scale data availability has enabled new methodologies that go far beyond traditional operations research paradigms. This section reviews the most relevant contributions from 2020 to 2025 across three main research streams: (i) predictive analytics in operations; (ii) optimization under uncertainty with AI-generated insights; and (iii) integrated decision intelligence and closed-loop analytics–optimization frameworks.

2.1. Predictive Analytics for Operational Decision-Making

Recent advancements in machine learning have improved the accuracy and robustness of demand forecasting, lead-time prediction, failure prediction, and process

monitoring. Studies such as [1–4] demonstrate how deep learning architectures—LSTM variants, temporal convolutional networks, attention-based models, and probabilistic neural networks—can capture nonlinearities and complex temporal patterns typical in supply chain and production systems. More recent work emphasizes uncertainty quantification and probabilistic forecasting using Bayesian neural networks, deep ensembles, and other uncertainty-aware deep learning methods [5, 6]. These models generate full predictive distributions rather than point forecasts, allowing more informed decision-making under uncertainty.

Machine learning has also been used to predict disruption risks (e.g., supplier failure, transportation delays) using multimodal data—including text, sensor streams, and satellite imagery—as demonstrated in [7, 8]. In parallel, domain-specific studies in industrial and infrastructure systems illustrate AI-enabled forecasting and performance improvement in supply chains and marketing, for example through intelligent demand prediction, resource allocation, and campaign optimization [25, 27, 28]. Such approaches underscore the increasing role of contextual and external data sources in operational forecasting beyond traditional ERP records.

2.2. Optimization Under Uncertainty with AI-driven Inputs

Classical stochastic programming and robust optimization remain foundational tools in operations research, but recent work has focused on integrating data-driven uncertainty models learned by AI systems. The “predict–then–optimize” paradigm, surveyed in [9, 10], has evolved to include end-to-end differentiable optimization frameworks that directly tie predictive model training to downstream decision performance [11, 12]. This shift enables models to internalize their impact on cost, risk, and service-level objectives, overcoming limitations of purely error-based training metrics.

Additionally, several studies propose hybrid ML–optimization pipelines for inventory control, scheduling, energy systems, and transportation planning [13–15]. These methods use AI-based scenario generation, distributional forecasting, or residual modeling to enhance the performance of stochastic optimization models under real-world uncertainty. Reinforcement learning (RL) has gained traction as a means of learning adaptive control policies in dynamic environments, as evidenced by the work in [19, 20], while data-driven multi-criteria decision-making techniques have been applied to infrastructure planning problems such as electric vehicle deployment and renewable energy integration [30]. However, RL and complex MCDM models often struggle with interpretability and stability, motivating research into hybrid RL–OR architectures and explainable optimization.

Methodological contributions from industrial engineering further explore how optimization and analytics can be tightly coupled for process design and quality control, providing templates for integrating analytical models into operational workflows [29].

2.3. Integrated Decision Intelligence and Closed-Loop Frameworks

A growing research frontier focuses on closing the loop between prediction, optimization, and continuous improvement—aligning with the emerging concept of *decision intelligence*. Several studies highlight the need for integrated systems that can process streaming data, update predictive models, and refine operational decisions dynamically [21–23]. Digital twins have been especially influential, serving as real-time simulation engines that evaluate alternative policies under evolving conditions and enable experimentation with “what-if” scenarios before implementation.

Despite progress, many existing systems lack modularity, transparent data-to-decision pipelines, and adaptive self-correction. The literature repeatedly notes that prediction and optimization components are often implemented separately, creating inconsistencies between forecast assumptions and decision models. As summarized in [24], the absence of unified architectures remains a barrier to full adoption of data-driven decision-making in industry.

2.4. Research Gap

Although significant advances have been made, several important limitations persist:

- Predictive models are often evaluated independently of their impact on operational performance.
- Optimization models frequently rely on simplified uncertainty assumptions that do not leverage the richness of AI-generated predictive distributions.
- Few studies offer truly modular, end-to-end pipelines that integrate prediction, optimization, monitoring, and iterative improvement.
- Existing frameworks rarely incorporate continuous feedback loops capable of adjusting forecasts and decisions as new data arrives.

This gap motivates the integrated *predict–optimize–improve* framework proposed in this paper, which unifies AI-driven predictive analytics, stochastic and robust optimization, and data-driven performance refinement into a coherent decision intelligence pipeline.

3. Methodology

This section presents the proposed *predict–optimize–improve* decision intelligence framework. The methodology is designed to integrate machine learning-based predictive analytics with stochastic and robust optimization, and a continuous feedback mechanism that refines both forecasting accuracy and operational performance. The framework consists of three interconnected stages: (i) data engineering and predictive modeling, (ii) prescriptive optimization under uncertainty using AI-derived distributions, and (iii) a closed-loop performance improvement layer.

3.1. Stage 1: Data Engineering and Predictive Modeling

The pipeline begins with the construction of a unified operational data environment. Historical demand, lead times, supplier disruptions, production capacity, and contextual features (e.g., weather, promotions, macroeconomic indicators) are consolidated into a structured data warehouse.

We employ a set of modern machine learning architectures suited for temporal and contextual forecasting, including:

- Long Short-Term Memory (LSTM) and temporal convolutional networks (TCN),
- Transformer-based sequence models for long-range dependencies,
- Deep ensembles and Bayesian neural networks for uncertainty quantification,
- Diffusion-based generative models for scenario generation (2023–2025 advancements).

The predictive module outputs full predictive *distributions* rather than point forecasts. These distributions provide quantiles, variance measures, and sample paths that directly inform downstream decision models. Model selection follows a hyperparameter search and cross-validation strategy optimized for decision relevance rather than solely statistical accuracy.

3.2. Stage 2: Optimization Under AI-Derived Uncertainty

The second component embeds predictive distributions into a prescriptive optimization model. Let $\tilde{D}_t$ denote the AI-generated demand distribution at time t , and let x_t represent the operational decision (e.g., replenishment, production quantity). We formulate the decision problem as:

$$\min_x \mathbb{E}_{\tilde{D}}[C(x, \tilde{D})] + \lambda \Phi(x, \tilde{D}),$$

where $C(\cdot)$ represents operational cost (production, inventory, shortage), and $\Phi(\cdot)$ captures risk measures such as CVaR or service-level penalties.

Two optimization strategies are supported:

Stochastic Optimization. AI-generated samples from $\tilde{D}_t$ are used to generate scenario sets for a scenario-based stochastic program. Decision variables include production/replenishment levels, safety-stock settings, and capacity allocations.

Robust Optimization. Predictive intervals form uncertainty sets $\mathcal{U}$, resulting in:

$$\min_x \max_{d \in \mathcal{U}} C(x, d).$$

This formulation yields decisions that remain feasible and cost-efficient under worst-case realizations of uncertainty.

Both approaches ensure consistency between forecasting assumptions and operational decisions, which is a major improvement over traditional deterministic planning.

3.3. Stage 3: Closed-Loop Feedback and Continuous Improvement

The final stage implements an adaptive feedback system that monitors discrepancies between predicted and realized outcomes. Let $\epsilon_t = D_t^{\text{actual}} - D_t^{\text{predicted}}$ denote forecast error. The system performs three updates:

1. **Predictive Model Update:** Model parameters are fine-tuned periodically using recent data, with greater weighting on periods where ϵ_t was large.
2. **Decision Policy Update:** Sensitivity analysis is used to adjust optimization parameters (e.g., risk aversion λ , service-level constraints).
3. **Process Improvement:** Operational KPIs (e.g., stockouts, utilization, backorder rates) feed into root-cause analysis to refine upstream data collection or feature engineering.

To support real-time adaptability, the framework enables incremental learning and rolling-horizon optimization. As new data arrive, the forecasting and optimization modules update dynamically, ensuring decisions remain aligned with current operating conditions.

3.4. System Architecture Summary

Figure 1 summarizes the full architecture of the pipeline. The system is modular, interpretable, and deployable in both on-premise and cloud-native environments. Each stage exposes configurable hyperparameters, enabling organizations to balance accuracy, computational cost, and responsiveness.

3.5. Reproducibility and Implementation Details

All experiments are executed using Python-based machine learning libraries (TensorFlow / PyTorch), combined with standard optimization solvers (Gurobi, CPLEX, or Pyomo). The predictive models are trained using an 80/10/10 historical data split, with rolling-window validation. Optimization runs are performed on a rolling horizon of 4–8 periods depending on the scenario.

We release code templates, data generators, and solver configurations as supplementary materials to ensure full reproducibility.

4. Case Study and Problem Formulation

To demonstrate the applicability of the proposed decision intelligence framework, we consider a multi-product production–inventory system operating under uncertain, time-varying demand. The case study is generic enough to represent a wide range of manufacturing or supply chain settings (e.g., a plant supplying multiple retailers or an assembly facility producing several product families), yet structured to highlight the benefits of integrating AI-driven predictive analytics with optimization under uncertainty.

4.1. Operational Setting

We consider a finite planning horizon $T = \{1, \dots, H\}$ divided into discrete periods (e.g., weeks), and a set of products $I = \{1, \dots, N\}$. In each period $t \in T$, stochastic customer demand $\tilde{D}_{it}$ is realized for each product $i \in I$. The firm decides on production or replenishment quantities x_{it} subject to capacity and resource constraints. Unsatisfied demand is either backordered or lost, and inventory is carried over between periods.

The decision-maker observes historical data for each product, including past demand, lead times, and relevant contextual features. As described in Section 3., these data feed the predictive modeling stage, which produces probabilistic forecasts for future demand. These forecasts are then embedded in a stochastic or robust optimization model used to determine the production/replenishment plan.

4.2. Notation and Decision Variables

We summarize the main notation used in the formulation.

- I : set of products, indexed by i .
- T : set of periods, indexed by t .




Figure 1: Overview of the integrated *predict-optimize-improve* decision intelligence framework linking operational data, AI-based prediction, scenario generation, optimization, and closed-loop feedback.

- $\tilde{D}_{it}$: random demand for product i in period t .
- D_{it}^s : demand realization for product i in period t under scenario s .
- S : set of demand scenarios generated from the AI-based predictive distributions.
- p^s : probability of scenario $s \in S$, with $\sum_{s \in S} p^s = 1$.

Decision variables:

- x_{it} : production or order quantity of product i in period t .
- I_{it}^s : inventory level of product i at the end of period t in scenario s .
- B_{it}^s : backorder level of product i at the end of period t in scenario s .

Parameters:

- c_i : unit production (or procurement) cost for product i .
- h_i : unit inventory holding cost per period for product i .
- b_i : unit backorder (or penalty) cost per period for product i .
- C_t : total production capacity available in period t .
- a_i : capacity consumption per unit of product i (e.g., processing time).
- I_{i0} : initial inventory of product i at the start of the horizon.

4.3. Scenario-Based Stochastic Programming Model

The AI-based predictive models described in Section 3. generate empirical demand distributions for each product and period. From these distributions, we construct a finite set of demand scenarios S using sampling or generative techniques. For each scenario $s \in S$, we obtain a trajectory $\{D_{it}^s\}_{i,t}$ that reflects both point forecasts and uncertainty patterns.

The stochastic program seeks a production/replenishment plan $\{x_{it}\}$ that minimizes the expected total cost across scenarios while satisfying material balance and capacity constraints.

4.3.1 Objective Function

The expected total cost over the planning horizon is given by:

$$\min_{x, I^s, B^s} Z = \sum_{t \in T} \sum_{i \in I} c_i x_{it} + \sum_{s \in S} p^s \left(\sum_{t \in T} \sum_{i \in I} (h_i I_{it}^s + b_i B_{it}^s) \right). \quad (1)$$

The first term captures production/procurement cost (deterministic across scenarios), while the second term represents the expected inventory and backorder costs under stochastic demand.

4.3.2 Inventory Balance Constraints

For each product i , period t , and scenario s , the inventory balance is:




Figure 2: Layered system architecture of the AI-driven decision-intelligence framework, illustrating how data ingestion, predictive analytics, scenario generation, optimization, and feedback loops interact in a unified operational planning environment.

$$I_{it}^s - B_{it}^s = I_{i,t-1}^s - B_{i,t-1}^s + x_{it} - D_{it}^s, \quad \forall i \in I, t \in T, s \in S, \quad (2)$$

with initial conditions:

$$I_{i0}^s = I_{i0}, \quad \forall i \in I, s \in S, \quad (3)$$

$$B_{i0}^s = 0, \quad \forall i \in I, s \in S. \quad (4)$$

Nonnegativity constraints enforce:

$$I_{it}^s \geq 0, \quad B_{it}^s \geq 0, \quad \forall i \in I, t \in T, s \in S. \quad (5)$$

4.3.3 Capacity Constraints

Production in each period is constrained by available capacity:

$$\sum_{i \in I} a_i x_{it} \leq C_t, \quad \forall t \in T. \quad (6)$$

We also impose nonnegativity on production quantities:

$$x_{it} \geq 0, \quad \forall i \in I, t \in T. \quad (7)$$

4.4. Robust Optimization Variant

While the stochastic program relies on explicit demand scenarios, some decision-makers prefer solutions that remain feasible and cost-effective across a range of plausible demand realizations. Let $\mathcal{U}_{it}$ denote an uncertainty set for demand of product i in period t , derived from the predictive distributions (e.g., confidence intervals or quantile bands).

A simple robust counterpart can be written as:




Figure 3: Circular representation of the robust optimization decision pipeline, showing how AI-based predictive models, uncertainty sets, robust optimization, robust plans, and stress-testing feedback form a closed decision loop.

$$\min_x \max_{\{d_{it} \in \mathcal{U}_{it}\}} \left\{ \sum_{t \in T} \sum_{i \in I} c_i x_{it} + \sum_{t \in T} \sum_{i \in I} (h_i I_{it}(d) + b_i B_{it}(d)) \right\}, \quad (8)$$

subject to inventory balance and capacity constraints defined as in (2) and (6), but now for all $d_{it} \in \mathcal{U}_{it}$. In practice, tractable robust formulations (e.g., budgeted or polyhedral uncertainty sets) are used, ensuring linear or mixed-integer programming structure.

4.5. Baseline Policies for Comparison

To assess the value added by the proposed AI-driven decision intelligence framework, we compare its performance against several benchmark policies commonly used in practice:

- **Deterministic planning with point forecasts:** demand is set to the mean forecast, and a deterministic linear program is solved ignoring uncertainty.
- **Safety stock heuristic:** production is based on mean forecasts plus fixed safety factors determined by historical variability, without explicit optimization under uncertainty.
- **Myopic policy:** each period's production is determined solely by current-period demand forecasts and inventory levels, without considering future periods.

These baselines represent typical approaches used in many organizations and provide a meaningful reference

for evaluating the performance benefits of the integrated *predict-optimize-improve* framework.

5. Results

This section presents the experimental evaluation of the proposed *predict-optimize-improve* framework. We assess its performance relative to commonly used baseline policies across stochastic, volatile operational conditions. All experiments use the multi-product production-inventory case study described in Section 4.

5.1. Experimental Setup

5.1.1 Computational Environment

All experiments were executed on a workstation equipped with an Intel Xeon 3.2 GHz CPU, 64 GB RAM, and Python-based machine learning libraries (TensorFlow/PyTorch) for predictive modeling. Optimization problems were solved using Gurobi 11.0 with parallelism disabled to ensure reproducibility.

5.1.2 Data and Scenario Construction

We simulate a rolling-horizon planning environment with $N = 5$ products, $H = 20$ planning periods, and time-varying demand distributions. Historical demand patterns are generated using a mixture of:

- seasonal components,
- random shocks,




Figure 4: Horizontal workflow for robust optimization using AI-based predictive distributions and uncertainty modeling.

- level shifts representing disruptions.

The predictive models (LSTM, Transformer, and Bayesian deep models) produce full predictive distributions for each product-period pair. From these distributions, we generate a scenario set S using:

1. **Deep generative sampling** (100 scenarios),
2. **Stratified quantile sampling** (20 scenarios),
3. **Hybrid ensemble sampling** (50 scenarios).

Unless stated otherwise, the stochastic program uses $|S| = 50$ scenarios with calibrated probabilities.

5.1.3 Baseline Methods

We compare the proposed method against three classical operational decision policies:

- **Deterministic planning:** point forecasts + linear programming.
- **Safety stock heuristic:** mean demand + factor-based buffers.
- **Myopic planning:** period-by-period replenishment.

These baselines represent the planning approaches still commonly used in practice.

5.1.4 Evaluation Metrics

Performance is evaluated along four dimensions:

1. **Total operational cost:** production + holding + backorder.
2. **Service level:** proportion of demand satisfied on time.
3. **Inventory dynamics:** average inventory, stockouts, and variability.
4. **Robustness:** performance degradation under unanticipated demand shocks.




Figure 5: Normalized total cost across decision policies (lower is better).

5.2. Main Quantitative Results

5.2.1 Total Cost Comparison

Table 1 reports total expected cost across the four methods. The integrated stochastic optimization approach achieves the lowest cost due to its ability to proactively hedge against uncertainty.

The proposed framework reduces total cost by:

- 32% vs. myopic,
- 22% vs. safety stock,
- 17% vs. deterministic optimization.

5.2.2 Service Level and Responsiveness

Table 2 shows service performance. The integrated method maintains high service levels even under demand volatility.

Higher service levels are achieved without overstocking due to more accurate uncertainty modeling and dynamic adjustment.

5.2.3 Inventory and Backorder Dynamics

Compared to baselines, the proposed method:


- reduces average inventory by 14–20%,
- reduces backorder frequency by 30–50%,

Table 1: Total expected cost (normalized units) over 30 Monte Carlo replications.

Method	$E[\text{Cost}]$	Std. Dev.
Myopic policy	1.00	0.06
Safety stock heuristic	0.87	0.05
Deterministic planning	0.82	0.04
Proposed predict–optimize–improve	0.68	0.03

Table 2: Service-level performance under uncertainty.

Method	Service Level (%)
Myopic	87.4
Safety stock	92.1
Deterministic	93.5
Proposed	97.8

**Figure 6:** Service-level performance under demand uncertainty.

- stabilizes inventory variance through scenario-driven hedging.

These improvements demonstrate the effectiveness of integrating predictive distributions with optimization.

Detailed Inventory and Backorder Statistics


Table 3 reports normalized inventory and backorder statistics averaged over all products, periods, and Monte Carlo replications. Average inventory is normalized so that the myopic policy equals 1.00.

These results show that the proposed framework maintains substantially lower average inventory while also cutting the fraction of periods with stockouts by roughly 35–50% relative to the baselines.

5.3. Robustness Analysis

To evaluate robustness, we introduce unanticipated shocks:

- 30% demand spike for selected products,
- temporary capacity reduction,

**Figure 7:** Robustness to unanticipated demand and capacity shocks.

- demand distributional shift (mean and variance drift).

Table 4 shows cost degradation relative to nominal conditions.

The proposed framework shows the lowest degradation, illustrating superior adaptability.

5.4. Ablation Study: Contribution of Each Component

We evaluate each component of the framework independently:

- using **only prediction + deterministic optimization**,
- using **only stochastic optimization with naive distributions**,
- full **predict–optimize–improve** pipeline.

Each stage contributes measurable value; the full integration yields the largest gains.

5.5. Managerial Insights

Several practical insights emerge:

1. **Uncertainty-aware planning outperforms point forecasts.** Even modest improvements in uncertainty modeling significantly reduce costs.
2. **Scenario diversity matters.** Generating a rich scenario set (via deep generative models) improves

Table 3: Average inventory and backorder frequency (normalized units).

Method	Avg. Inventory	Backorder Frequency (% of periods)
Myopic policy	1.00	16.8
Safety stock heuristic	0.98	13.2
Deterministic planning	0.95	11.9
Proposed predict-optimize-improve	0.82	7.8

Table 4: Performance degradation under shocks (% increase in cost).

Method	Cost Increase (%)
Myopic	38.6
Safety stock	27.4
Deterministic	22.1
Proposed	11.8

Table 5: Impact of pipeline components.

Configuration	Total Cost
Prediction + deterministic optimization	0.79
Naive stochastic optimization	0.75
Full framework	0.68

robustness.



3. **Closed-loop feedback is essential.** Incorporating recent discrepancies accelerates adaptation to volatility.
4. **Balanced inventory is a byproduct of better forecasting, not higher buffers.** Organizations can increase service level *and* reduce inventory simultaneously.

5.6. Summary

The experimental results demonstrate that the proposed integrated framework significantly improves cost efficiency, service levels, and robustness relative to traditional methods. The results validate the value of unifying predictive analytics, stochastic optimization, and continuous improvement into a coherent decision intelligence platform.

5.7. Visualization of Comparative Performance

To make the quantitative differences more interpretable for practitioners, we provide simple bar charts summarizing total cost and service-level performance across the four methods.

**Figure 8:** Comparison of normalized total cost across planning policies.**Figure 9:** Service-level comparison under demand uncertainty.

6. Discussion

The results presented in Section ?? demonstrate that integrating predictive analytics, stochastic or robust optimization, and closed-loop feedback offers a substantial performance advantage over traditional operational planning methods. This section discusses the broader implications of these findings, the conditions under which the proposed framework is most effective, and the managerial and theoretical contributions to intelligent operations management.

6.1. Interpreting the Performance Gains

The superior cost efficiency and service-level improvements achieved by the proposed framework highlight the central role of uncertainty modeling in modern operational systems. Traditional deterministic approaches rely on point forecasts that systematically

underestimate risk, leading to over-reactive or myopic decisions. In contrast, the predictive models employed here generate full predictive distributions that capture nonlinear patterns, structural breaks, and temporal dependencies. Embedding these distributions into a stochastic optimization model enables the decision-maker to proactively hedge against variability, especially when demand fluctuations or disruption risks are substantial.

The robustness analysis further illustrates that decision policies informed by probabilistic forecasts degrade more gracefully under unanticipated shocks, such as abrupt demand spikes or capacity reductions. By incorporating scenario-diversity—derived from deep generative models and stratified sampling techniques—the system maintains stability even when data-generating processes shift over time. This resilience underscores the value of uncertainty-aware optimization as a foundation for operational decision-making in volatile environments.

6.2. When the Framework Provides the Largest Benefits

The integrated decision intelligence framework is most beneficial in settings characterized by one or more of the following conditions:

- **High demand variability:** Organizations with significant temporal fluctuations, seasonality, or promotion-driven demand exhibit large performance improvements due to better hedging.
- **Frequent disruptions or supply uncertainty:** The framework's robustness becomes critical when lead times shift or upstream suppliers become unreliable.
- **Multi-product, capacity-constrained systems:** As resource allocation becomes more complex, scenario-based planning significantly outperforms heuristic rules.
- **Data-rich environments:** Firms with access to ERP, MES, IoT, or customer transaction data benefit from advanced forecasting models, which in turn support more accurate optimization.

Conversely, when uncertainty is minimal or demand patterns are nearly stationary, the performance gap between advanced methods and deterministic planning narrows. However, even in such cases, the closed-loop learning component ensures adaptability as conditions evolve.

6.3. Managerial and Organizational Implications

The findings of this study underscore several managerial insights. First, **investing in predictive analytics yields operational value only when coupled with**

decision-focused optimization. Isolated forecasting improvements do not automatically translate into better outcomes unless the predictions are operationalized through cost-sensitive models.

Second, the results highlight the importance of **scenario diversity.** Relying on narrow distributions or simplistic uncertainty assumptions may produce biased or brittle decisions. Managers should therefore prioritize scenario generation approaches—such as deep generative models or ensemble sampling—that capture a wide spectrum of plausible futures.

Third, the closed-loop feedback mechanism demonstrates that **decision intelligence is inherently iterative.** Organizations that continuously monitor forecast errors, operational KPIs, and constraint violations can gradually align their data-generating, modeling, and optimization processes to create a self-improving operational system.

Finally, the study provides a practical blueprint for shifting from experience-driven decisions to **data-driven, evidence-based planning.** As firms navigate increasingly volatile supply chains, this transition becomes essential to maintaining competitiveness, cost efficiency, and service-level commitments.

6.4. Theoretical Contributions

Beyond its practical significance, the study makes several theoretical contributions. It unifies three traditionally independent streams of research—predictive modeling, stochastic optimization, and continuous improvement—into a coherent methodological pipeline. This integration provides a generalizable framework for embedding learning mechanisms into classical operations research models, bridging the gap between machine learning and optimization under uncertainty.

Additionally, the use of deep generative models for scenario construction advances the literature on data-driven stochastic programming, offering an alternative to traditional bootstrapping or distribution-fitting approaches. The experimental evidence supports the claim that richer scenario sets lead to more resilient decision policies.

Taken together, these contributions position decision intelligence as a paradigm shift in operations management, moving beyond reactive, static planning toward adaptive and self-correcting operational systems.

7. Conclusion

This study presented an integrated decision-intelligence framework that unifies predictive analytics, stochastic optimization, and continuous process improvement into a coherent pipeline for intelligent operational decision-making. Motivated by the increasing volatility and data intensity of modern supply chains, the

framework addresses fundamental limitations of traditional deterministic or heuristic planning approaches, which struggle to represent nonlinear uncertainty, react to disruptions, or adapt to shifting data-generating processes.

Across a multi-product production-inventory environment characterized by time-varying demand, supply disruptions, and resource constraints, the proposed predict-optimize-improve architecture consistently outperformed classical baselines. Experimental results demonstrated significant reductions in total operational cost, higher service levels, improved responsiveness to demand shocks, and greater robustness under distributional shifts. Ablation studies confirmed that each pipeline component—probabilistic forecasting, scenario-based optimization, and closed-loop learning—adds measurable value, with the full integration providing the strongest performance.

Beyond empirical gains, the work advances the theory of intelligent operations management by bridging machine learning, deep generative uncertainty modeling, and operations research decision models. The framework provides a replicable method for embedding learning mechanisms into optimization workflows, enabling systems that adapt, recalibrate, and self-correct as conditions evolve.

In practice, the results highlight a clear managerial message: predictive analytics produce operational value only when tightly coupled with decision-focused optimization and iterative feedback mechanisms. Organizations seeking resilience and competitiveness in data-rich, uncertain environments can benefit from transitioning toward evidence-based, uncertainty-aware, and iteratively improving decision processes.

Future research may extend this work in several directions, including multi-echelon supply chain settings, reinforcement learning-based adaptive policies, tighter integrations between digital-twin environments and stochastic optimization, and applications to highly perishable or service-oriented systems. As industrial systems continue their digital transformation, the proposed decision-intelligence architecture represents a robust and scalable foundation for next-generation operational planning.

8. Limitations and Future Research

While the proposed predict-optimize-improve framework demonstrates strong performance and provides a unified architecture for intelligent operational decision-making, several limitations warrant discussion and present valuable opportunities for future research.

8.1. Modeling and Data Limitations

First, the predictive models rely on the availability of sufficiently rich historical data, including past demand, contextual variables, and operational events. Organizations with limited or highly fragmented data may experience reduced predictive accuracy, which can subsequently affect the quality of scenario generation. Additionally, although probabilistic forecasting models capture heteroscedasticity and temporal dynamics, they may still struggle under extreme structural breaks or rare-event disruptions that fall outside the training distribution.

Second, the scenario-based stochastic optimization model depends on the quality and representativeness of generated scenarios. While deep generative and ensemble-based sampling methods improve diversity, guaranteeing full coverage of all plausible uncertainty realizations remains challenging, especially in high-dimensional, multivariate settings.

8.2. Computational and Scalability Considerations

The integration of deep learning and stochastic optimization introduces nontrivial computational requirements. Solving large-scale multi-stage stochastic programs may require substantial computational resources or specialized solvers, particularly when the number of scenarios grows. Although progressive hedging and decomposition methods offer partial relief, scaling the framework to very large production networks, multi-echelon supply chains, or real-time decision environments may require further algorithmic innovations.

8.3. Practical Implementation Challenges

From an organizational perspective, deploying an end-to-end decision-intelligence system requires alignment across data engineering, analytics teams, and operational managers. Challenges include model governance, interpretability, integration with ERP/MES systems, and maintaining reliable feedback loops. Moreover, human-AI collaboration in operational decision-making remains an open topic: determining the appropriate balance between automated policies and managerial oversight is both context-specific and understudied.

8.4. Future Research Directions

Building on the findings of this work, several avenues for future investigation are promising:

- **Multi-echelon and networked systems:** Extending the framework to multi-echelon supply chains, logistics networks, and distributed production systems would enhance its applicability.

- **Reinforcement learning integration:** Combining stochastic optimization with reinforcement learning could produce adaptive policies capable of reacting to continuous feedback without relying on discrete scenario sets.
- **Digital twins for real-time decision intelligence:** Integrating the framework into a high-fidelity digital twin environment would enable real-time simulation, anomaly detection, and robust policy stress-testing.
- **Explainable AI for decision-focused models:** Developing transparent, interpretable versions of probabilistic forecasting models and optimization outputs would improve managerial trust and adoption.
- **Robustness to rare events and black swan disruptions:** Enhancing generative scenario models to better cover tail-risk phenomena—such as pandemics, geopolitical shocks, or supplier failures—remains a critical research direction.
- **Sustainability and environmental metrics:** Incorporating carbon footprint, energy efficiency, or circular economy metrics into optimization objectives would expand the framework's relevance to sustainability-driven operations.

Overall, while the proposed framework provides a solid foundation for embedding AI-driven analytics into operations research models, continued methodological innovation and empirical validation will further strengthen the practical and theoretical contributions of decision intelligence for next-generation operational planning.

References

- [1] Salinas, D., Flunkert, V., Gasthaus, J., & Januschowski, T. (2020). DeepAR: Probabilistic forecasting with autoregressive recurrent networks. *International Journal of Forecasting*, 36(3), 1181–1191.
- [2] Hewamalage, H., Bergmeir, C., & Bandara, K. (2021). Recurrent neural networks for time series forecasting: Current status and future directions. *International Journal of Forecasting*, 37(1), 388–427.
- [3] Lim, B., & Zohren, S. (2021). Time series forecasting with deep learning: A survey. *Philosophical Transactions of the Royal Society A*, 379(2194), 20200209.
- [4] Taghizadeh, A., Mirzaei, A., & Reisi, M. (2023). Demand forecasting in enterprise resource planning system: A machine learning approach. *Journal of Applied Research in Industrial Engineering*, 10(1), 54–72.
- [5] Gal, Y., & Ghahramani, Z. (2016). Dropout as a Bayesian approximation: Representing model uncertainty in deep learning. In *Proceedings of the 33rd International Conference on Machine Learning* (pp. 1050–1059).
- [6] Lakshminarayanan, B., Pritzel, A., & Blundell, C. (2017). Simple and scalable predictive uncertainty estimation using deep ensembles. In *Advances in Neural Information Processing Systems*, 30, 6402–6413.
- [7] Bassiouni, M. M., Chakraborty, R. K., Hussain, O. K., & Rahman, H. F. (2022). Advanced deep learning approaches to predict supply chain risks under COVID-19 restrictions. *Expert Systems with Applications*, 211, 118604.
- [8] Kalisetty, S., Pandugula, C., & Mallesham, G. (2023). Leveraging artificial intelligence to enhance supply chain resilience: A study of predictive analytics and risk mitigation strategies. *Journal of Artificial Intelligence and Big Data*, 3(4), 1–23.
- [9] Elmachoub, A. N., & Grigas, P. (2022). Smart “predict, then optimize”. *Management Science*, 68(1), 9–26.
- [10] Donti, P. L., Amos, B., & Kolter, J. Z. (2017). Task-based end-to-end model learning in stochastic optimization. In *Advances in Neural Information Processing Systems*, 30, 5484–5494.
- [11] Mandi, J., Bucarey, V., Guns, T., & Hutter, F. (2020). Smart predict-and-optimize for hard combinatorial optimization problems. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 34(02), 1603–1610.
- [12] Wilder, B., Dilkina, B., & Tambe, M. (2019). Melding the data-decisions pipeline: Decision-focused learning for combinatorial optimization. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 33(01), 1658–1665.
- [13] Bertsimas, D., & Thiele, A. (2006). A robust optimization approach to inventory theory. *Operations Research*, 54(1), 150–168.
- [14] Bertsimas, D., Gupta, V., & Kallus, N. (2018). Robust sample average approximation. *Mathematical Programming*, 171(1–2), 217–282.
- [15] Birge, J. R., & Louveaux, F. (2011). *Introduction to Stochastic Programming* (2nd ed.). Springer.
- [16] Bertsimas, D., & Kallus, N. (2020). From predictive to prescriptive analytics. *Management Science*, 66(3), 1025–1044.

- [17] Ben-Tal, A., El Ghaoui, L., & Nemirovski, A. (2009). *Robust Optimization*. Princeton University Press.
- [18] Shapiro, A., Dentcheva, D., & Ruszczyński, A. (2014). *Lectures on Stochastic Programming: Modeling and Theory* (2nd ed.). SIAM.
- [19] Powell, W. B. (2019). *Unified Framework for Sequential Decisions: From Markov Decision Processes to Reinforcement Learning*. Princeton University Press.
- [20] Sutton, R. S., & Barto, A. G. (2018). *Reinforcement Learning: An Introduction* (2nd ed.). MIT Press.
- [21] Tao, F., Zhang, H., Liu, A., & Nee, A. Y. C. (2019). Digital twin in industry: State-of-the-art. *IEEE Transactions on Industrial Informatics*, 15(4), 2405–2415.
- [22] Ivanov, D., & Dolgui, A. (2020). Viability of intertwined supply networks: Extending the supply chain resilience angles towards survivability. *International Journal of Production Research*, 58(10), 2904–2925.
- [23] Mikalef, P., Krogstie, J., Pappas, I. O., & Pavlou, P. (2021). Exploring the relationship between big data analytics capability and competitive performance: The mediating roles of dynamic and operational capabilities. *Information & Management*, 58(3), 103434.
- [24] Provost, F., & Fawcett, T. (2013). Data science and its relationship to big data and data-driven decision making. *Big Data*, 1(1), 51–59.
- [25] Shoushtari, F., Ghafourian, E., & Talebi, M. (2021). Improving performance of supply chain by applying artificial intelligence. *International Journal of Industrial Engineering and Operational Research*, 3(1), 14–23.
- [26] Shoushtari, F., Daghighi, A., & Ghafourian, E. (2024). Application of artificial intelligence in project management. *International Journal of Industrial Engineering and Operational Research*, 6(2), 49–63.
- [27] Shoushtari, F., Bashir, E., Hassankhani, S., & Rezvanjou, S. (2023). Optimization in marketing enhancing efficiency and effectiveness. *International Journal of Industrial Engineering and Operational Research*, 5(2), 12–23.
- [28] Lotfi, R., Shoushtari, F., Ali, S. S., Davoodi, S. M. R., Afshar, M., & Sharifi Nevisi, M. M. (2024). A viable and bi-level supply chain network design by applying risk, robustness and considering environmental requirements. *Central European Journal of Operations Research*, 32, 1–29.
- [29] Zadeh, M. S. N., Shoushtari, F., & Talebi, M. (2024). Optimization of analytical methods in industrial engineering: Enhancing decision-making in process design and quality control. *International Journal of Industrial Engineering and Construction Management*, 2(1), 27–40.
- [30] Talebi, M., Rezvanjou, S., Ghafourian, E., & Shoushtari, F. (2025). Machine learning-driven multi-criteria decision-making models for optimizing sustainable infrastructure deployment: Electric vehicles and renewable energy systems. *SSRN Working Paper*, 5121862.
- [31] Makridakis, S., Spiliotis, E., & Assimakopoulos, V. (2018). Statistical and machine learning forecasting methods: Concerns and ways forward. *PLoS One*, 13(3), e0194889.
- [32] Makridakis, S., Spiliotis, E., & Assimakopoulos, V. (2020). The M5 competition: Results, findings and conclusions. *International Journal of Forecasting*, 36(1), 54–74.
- [33] Benidis, K., Rangapuram, S., Flunkert, V., et al. (2022). Deep learning for time series forecasting: Tutorial and survey. *ACM Computing Surveys*, 55(6), 1–36.
- [34] Goodfellow, I., Bengio, Y., & Courville, A. (2016). *Deep Learning*. MIT Press.
- [35] Kingma, D. P., & Ba, J. (2015). Adam: A method for stochastic optimization. In *International Conference on Learning Representations*.
- [36] Bengio, Y., Courville, A., & Vincent, P. (2013). Representation learning: A review and new perspectives. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 35(8), 1798–1828.
- [37] Silver, D., Schrittwieser, J., Simonyan, K., et al. (2017). Mastering the game of Go without human knowledge. *Nature*, 550(7676), 354–359.
- [38] Ivanov, D. (2020). Predicting the impacts of epidemic outbreaks on global supply chains: A simulation-based analysis on the coronavirus outbreak (COVID-19/SARS-CoV-2) case. *Transportation Research Part E*, 136, 101922.
- [39] Makridakis, S., Spiliotis, E., Assimakopoulos, V., & Petropoulos, F. (2022). The M competitions: What have we learned so far? *International Journal of Forecasting*, 38(3), 787–806.