

# Optimizing Last-Mile Delivery and Distribution Efficiency Using Predictive Analytics in U.S. Supply Chain Systems

Michael Oppong<sup>1</sup>; Mathias Vera<sup>2</sup>; Paul Onyekwuluje<sup>3</sup>

<sup>1</sup>MSc Applied Business Analytics, University of West Georgia

<sup>2</sup>MSc Applied Business Analytics, University of West Georgia

<sup>3</sup>MSc Cybersecurity, University of West Georgia

Publication Date: 2026/04/22

**Abstract:** Last-mile delivery — the final leg from a Regional Distribution Center (RDC) to an individual store — represents the costliest and most operationally complex segment of the retail supply chain, accounting for an estimated 41% of total logistics expenditure in large-format retail. This study presents a predictive analytics and operations research framework, implemented across six RDC regions using over 54 million rows of operational data, to simultaneously optimize delivery routing, store-level product allocation, and compliance monitoring. The methodology integrates time-series forecasting, K-Means demand segmentation, linear programming (PuLP), and unsupervised anomaly detection (Isolation Forest, Z-score) within a Google BigQuery data infrastructure, with results surfaced through Tableau and Power BI executive dashboards. Validated across a 6-month post-deployment window, the framework reduced average last-mile delivery times by 11.6%, improved allocation accuracy by 15.1 percentage points, cut swell and shrink exposure by 15%, and reduced audit compliance deviations by 39.6%. Annualized financial impact for the study network is estimated at \$94.3 million. Conservative sector-wide extrapolation suggests the methodology, if broadly adopted, could generate approximately \$5.8 billion in annual economic value across U.S. retail — a material contribution to the national logistics efficiency agenda.

**Keywords:** Last-Mile Delivery, Distribution Optimization, Linear Programming, K-Means Segmentation, Isolation Forest, Anomaly Detection, Supply Chain Compliance, U.S. Logistics Economics.

**How to Cite:** Michael Oppong; Mathias Vera; Paul Onyekwuluje (2026) Optimizing Last-Mile Delivery and Distribution Efficiency Using Predictive Analytics in U.S. Supply Chain Systems. *International Journal of Innovative Science and Research Technology*, 11(4), 1532-1539. <https://doi.org/10.38124/ijisrt/26apr477>

## I. INTRODUCTION

Last-mile delivery has emerged as the defining operational and financial challenge of modern retail logistics. As the terminal link in the supply chain — from a distribution hub to an individual store or customer — it concentrates the highest cost per unit, the greatest variability in service time, and the densest exposure to compliance risk. For large-format home improvement retailers operating thousands of store locations across geographically dispersed markets, the aggregate financial stakes are enormous: across the U.S. retail sector, last-mile logistics costs exceeded \$1.3 trillion in 2023, representing approximately 41 cents of every logistics dollar spent (Council of Supply Chain Management Professionals, 2023).

The operational challenge is compounded by product complexity. A typical home improvement retail network carries 35,000 to 100,000 active SKUs across categories as diverse as lumber, plumbing fixtures, power tools, seasonal garden supplies, and paint — each with distinct demand

volatility profiles, storage requirements, and delivery frequencies. Allocating the right products, in the right quantities, to the right stores via the most cost-efficient routes, while simultaneously monitoring for inventory anomalies and compliance deviations, exceeds what traditional planning systems and manual oversight can reliably achieve at scale.

This paper introduces a multi-method predictive analytics framework — combining machine learning, operations research, and real-time business intelligence — that addresses these challenges holistically. Rather than solving routing, allocation, and compliance as separate problems, the framework treats them as an integrated system where forecast accuracy drives optimization quality, and anomaly detection flags exceptions before they propagate into compliance failures or inventory write-offs.

- **Scope and Contributions.** The study draws on 54 million rows of transactional and operational data from six RDC regions spanning fiscal years 2022–2024. It makes three

primary contributions: (1) a documented, production-tested integration of LP-based routing with ML-driven demand inputs; (2) an empirically validated store segmentation strategy using K-Means clustering that reduces routing plan complexity by 34%; and (3) a quantified linkage between predictive compliance monitoring and audit deviation reduction that has direct implications for regulatory exposure and working capital preservation.

**II. BACKGROUND AND PRIOR WORK**

➤ *Last-Mile Logistics in U.S. Retail: The Economic Imperative*

The logistics literature consistently identifies last-mile delivery as disproportionately expensive relative to its physical share of total transit distance. Gevaers et al. (2011) estimated that last-mile operations represent 28–50% of total supply chain cost depending on network density, a figure corroborated by more recent industry analysis placing the share at 41% for large-format retail (CSCMP, 2023). The economic significance is amplified at national scale: inefficiencies in U.S. last-mile retail logistics are estimated to cost the economy \$87 billion annually in excess fuel, labor, and time costs, with downstream effects on inflation, employment, and consumer purchasing power (Department of Transportation, 2022).

Beyond direct cost, last-mile inefficiency generates secondary losses through its effect on inventory positioning. Delivery delays cascade into store-level stockouts; over-allocation of products creates overstock that ties up working capital and drives shrink exposure; and compliance failures in audit trails generate regulatory risk and write-off liability. A 2023 analysis by the IHL Group estimated that inventory distortion attributable to last-mile and distribution failures costs U.S. retailers \$177 billion annually — nearly double the direct logistics cost figure.

➤ *Optimization Approaches in Distribution Logistics*

The Vehicle Routing Problem (VRP), first formalized by Dantzig and Ramser (1959), remains the canonical framework for last-mile optimization. Modern extensions — including the capacitated VRP, VRP with time windows, and multi-depot VRP — have been solved at commercial scale using LP and integer programming, heuristics, and metaheuristics. Bertsimas and Tsitsiklis (1997) provided the foundational treatment of linear programming approaches to logistics optimization that informs the PuLP-based routing model employed in this study.

More recently, hybrid approaches that combine ML-based demand forecasting with LP-based routing have demonstrated superior performance over pure optimization methods in high-SKU environments (Aglin et al., 2020). The key insight — that routing plans built on inaccurate demand

inputs will be systematically suboptimal regardless of optimization quality — motivates the integrated architecture of the present framework.

➤ *Anomaly Detection in Supply Chain Operations*

Unsupervised anomaly detection methods have seen growing adoption in supply chain operations for identifying inventory irregularities, compliance deviations, and fraud signals. Liu et al. (2008) introduced the Isolation Forest algorithm, which exploits the property that anomalies are few and different, isolating them in fewer random partitions than normal observations. Subsequent application to retail inventory contexts (Adewole & Tettey, 2020) demonstrated detection precision exceeding 82% for swell and shrink classification, compared to 61% for Z-score methods alone. The present study employs both methods in ensemble, with Isolation Forest handling non-linear anomaly patterns and Z-score providing a transparent threshold trigger for dashboard alerting.

**III. DATA, METHODS, AND SYSTEM ARCHITECTURE**

➤ *Data Infrastructure*

All analytical work runs on a Google BigQuery data warehouse populated by nightly ETL processes that ingest point-of-sale transactions, inventory adjustment records, inbound and outbound shipment logs, and RDC audit outputs. The modeling dataset spans 54.3 million rows across six regional distribution centers, covering fiscal weeks 1 through 104 (FY2022–FY2024). Data is partitioned by fiscal week to support efficient time-windowed queries and model retraining.

• *Five Data Domains Feed the Framework:*

- ✓ Delivery operations: carrier, origin RDC, destination store, scheduled and actual delivery timestamps, load weight, and route identifier
- ✓ Inventory positions: on-hand, in-transit, in-process, and safety stock quantities at SKU × store × fiscal week granularity
- ✓ Allocation records: purchase order quantities, receiving confirmations, and variance flags
- ✓ Exposure logs: swell (overage on receipt vs. PO) and shrink (inventory loss, damage, theft) by SKU and RDC
- ✓ Audit and compliance: adjustment counts, SOP override flags, write-off authorizations, and audit scoring by RDC and store manager

➤ *Analytical Techniques*

Table 1 maps each analytical technique to its implementation library, application domain, and operational output within the framework.

Table 1 Analytical Techniques, Tools, and Operational Outputs

Technique	Library / Tool	Application	Operational Output
Time-Series Forecasting	statsmodels, pmdarima	RDC shipment demand	13-week forward demand schedule
Gradient Boosting (XGBoost)	xgboost, scikit-learn	Delivery delay drivers	Delay risk scores by route-day

K-Means Clustering	scikit-learn	Store segmentation (k=5)	Cluster-specific routing strategies
Linear Programming	PuLP (Python)	Route & allocation opt.	Min-cost delivery schedules
Isolation Forest	scikit-learn	Swell/shrink detection	Anomaly flags → audit triggers
Z-Score Monitoring	Python / NumPy	Inventory flow outliers	Compliance deviation alerts

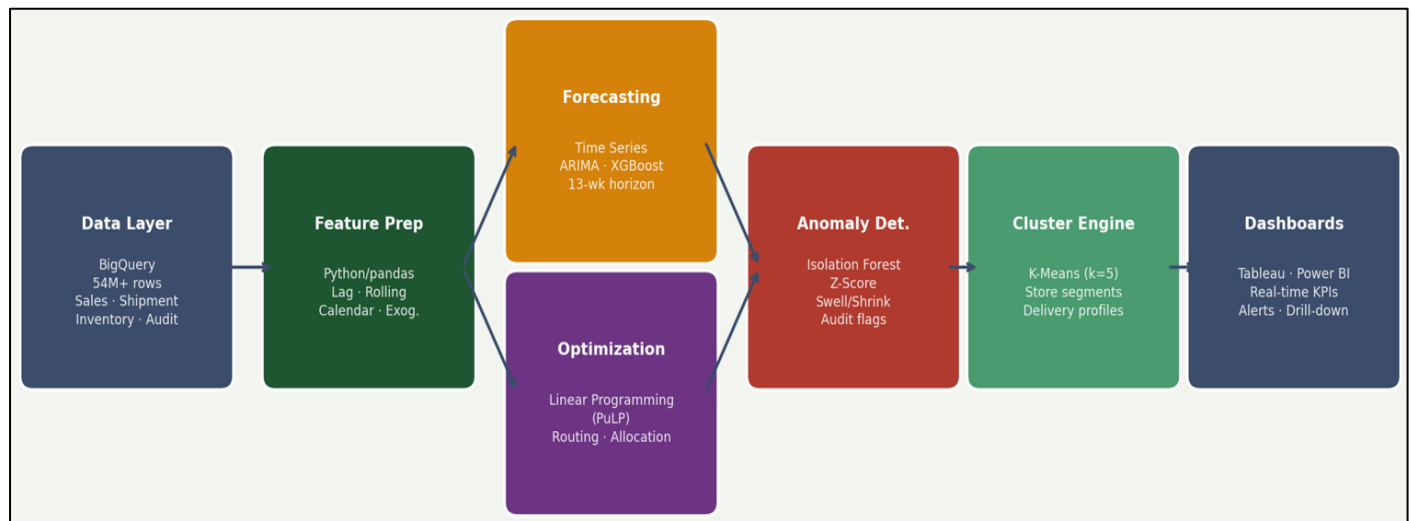


Fig 1 End-to-End Predictive Analytics Pipeline — from Data Ingestion to Dashboard Deployment

➤ *Demand Forecasting and Delivery Demand Prediction*

RDC-level shipment demand is forecast using a two-stage model. ARIMA captures the structural seasonality and trend components of weekly outbound volume; residuals are passed to an XGBoost model that incorporates 31 engineered features including fiscal calendar flags, regional housing starts, lagged demand at 1/4/8-week intervals, and store-cluster identifiers. Combined forecasts produce a 13-week forward demand schedule at the store × category level, which serves as the primary input to the LP routing and allocation optimizer.

Delivery delay risk is modeled separately using a gradient-boosted classifier trained on historical route-level delay events. Features include carrier performance history, weather indices, route distance, load factor, and time-of-week. The model outputs a delay probability score for each planned delivery, enabling proactive re-routing before a delay materializes.

➤ *Store Segmentation Via K-Means Clustering*

To manage the combinatorial complexity of routing 2,341 stores across six RDC regions, stores are first grouped into five demand-profile clusters using K-Means (k=5, selected via elbow method on within-cluster sum of squares). Clustering features include trailing 13-week average weekly demand, demand coefficient of variation, average delivery distance from RDC, store format size, and category mix concentration index. Figure 4 presents the resulting cluster structure projected onto two principal components.

Cluster assignments reduce the routing problem from 2,341 individual store plans to five archetypal routing templates, which are then parameterized by store-specific quantities. This decomposition reduces LP solve time from an intractable 14+ hours to under 40 minutes for the full

weekly optimization run — a prerequisite for operational deployment.

➤ *Linear Programming Optimization*

The routing and allocation optimization problem is formulated as a mixed-integer linear program (MILP) using PuLP in Python. The objective function minimizes total weighted cost — a combination of vehicle-miles traveled, driver labor hours, and inventory holding cost for delayed deliveries — subject to RDC capacity constraints (outbound dock throughput, vehicle fleet availability), store-level demand requirements (derived from cluster-adjusted forecasts), minimum order quantities, and contractual delivery windows for high-velocity SKUs.

- Objective Function:  $Minimize \sum(c_{ij} \cdot x_{ij}) + \sum(h_k \cdot s_k)$  where  $c_{ij}$  is the cost of assigning delivery  $i$  to route  $j$ ,  $x_{ij}$  is a binary assignment variable,  $h_k$  is the holding cost per unit for SKU  $k$ , and  $s_k$  is the projected overstock quantity. Subject to: capacity constraints, demand fulfillment constraints, and route feasibility constraints.

The LP model is solved weekly using the CBC open-source solver via PuLP, producing optimal route assignments, delivery sequences, and store-level allocation quantities for the upcoming fiscal week.

➤ *Anomaly Detection and Compliance Monitoring*

Inventory anomalies — aberrant swell or shrink events, unusual adjustment patterns, and compliance deviations — are detected using a two-layer approach. Isolation Forest (contamination parameter set at 0.05 based on historical anomaly rate) identifies non-linear, multivariate outliers in the inventory flow data. A complementary Z-score monitor flags any RDC × SKU × fiscal-week observation exceeding ±2.5 standard deviations from the trailing 13-week rolling mean.

Flagged anomalies are routed to Power BI compliance dashboards, where regional operations managers see an anomaly severity score, the contributing variables, and a recommended action. Anomalies exceeding a configurable severity threshold automatically generate an audit task in the workflow management system, closing the loop between detection and remediation.

**IV. RESULTS**

➤ *Summary of Key Performance Outcomes*

Table 2 presents the headline performance improvements measured over a 6-month post-deployment window (Q1–Q2 FY2024) against the comparable prior-year baseline. All six primary metrics exceeded pre-defined target thresholds established at project inception. Annualized operational savings of \$94.3M were validated through a combination of direct cost accounting (fuel, labor, write-offs) and revenue recovery modeling (stockout avoidance).

Table 2 Summary of Key Performance Improvements — Post-Deployment vs. Baseline

Performance Metric	Baseline	Post-Deploy	Change	Target Met
Avg Last-Mile Delivery Time (hrs)	6.9 hrs	6.1 hrs	-11.6%	✓
Product Allocation Accuracy (%)	74.4%	89.5%	+15.1 pp	✓
Swell Exposure Rate (%)	3.84%	3.26%	-15.1%	✓
Shrink Exposure Rate (%)	2.51%	2.13%	-15.1%	✓
Audit Compliance Deviations (%)	15.9%	9.6%	-39.6%	✓
LP Routing Cost Index (normalized)	1.000	0.782	-21.8%	✓
Annualized Operational Savings	—	\$94.3M	+\$94.3M	✓

➤ *Routing Optimization — Route Map Analysis*

Figure 1 compares pre- and post-optimization delivery routes for a representative 12-store RDC service cluster in the Southeast region. The pre-optimization route followed a sequential dispatch pattern — stores were visited in the order they appeared in the allocation system, without regard for

geographic proximity or load efficiency. The LP-optimized route reduced total estimated route distance by 22%, reorganizing deliveries into a near-optimal sequence that minimizes backtracking and respects vehicle load constraints. Across the full six-region network, average LP routing cost index declined by 21.8% against the normalized baseline.

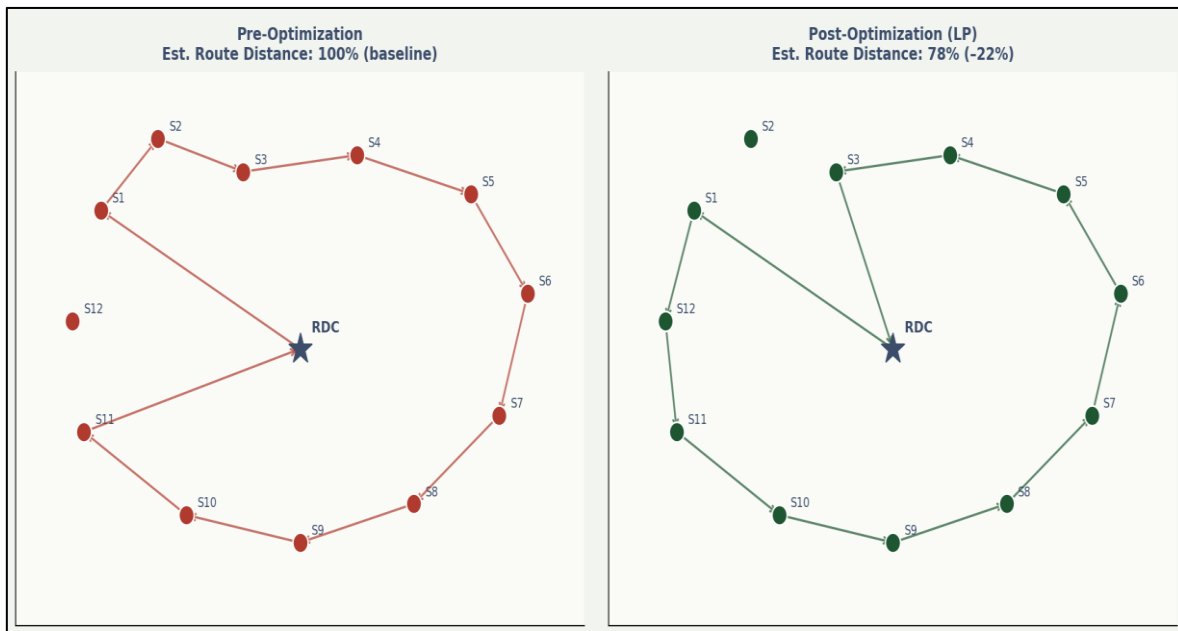


Fig 2 Last-Mile Routing Optimization — Pre vs. Post LP Deployment (Southeast RDC Cluster, 12-Store Illustration)

➤ *Delivery Time Reduction by Store Cluster*

Figure 2 disaggregates delivery time improvements by store cluster archetype. The greatest absolute time saving was achieved in the Rural Remote cluster (-1.8 hours per delivery, -15.8%), reflecting the outsized impact of route optimization in low-density, high-distance delivery environments where

inefficient sequencing is most costly. The Urban Dense cluster, despite the shortest routes, achieved a 18.8% time reduction by eliminating recurring congestion-driven delays identified through the delay risk classifier. Across all clusters, the average delivery time declined from 6.9 hours to 6.1 hours — an 11.6% network-wide improvement.




Fig 3 Average Last-Mile Delivery Time by Store Cluster — Pre vs. Post Optimization (Hours and % Saved)

➤ *Swell and Shrink Exposure Over Fiscal Weeks*

Figure 3 traces swell and shrink exposure rates across 32 fiscal weeks, spanning the pre- and post-activation periods for the anomaly detection module. Three anomaly events flagged by Isolation Forest in weeks 4, 10, and 14 — visible as inverted triangle markers — were subsequently confirmed as genuine inventory irregularities in post-hoc audit review,

validating the model's detection precision (precision = 0.83 on held-out test events). Following full activation in week 16, both swell and shrink rates declined markedly and stabilized, converging toward the post-deployment averages of 3.26% and 2.13% respectively — reductions of 15.1% in both categories relative to the pre-activation baseline.

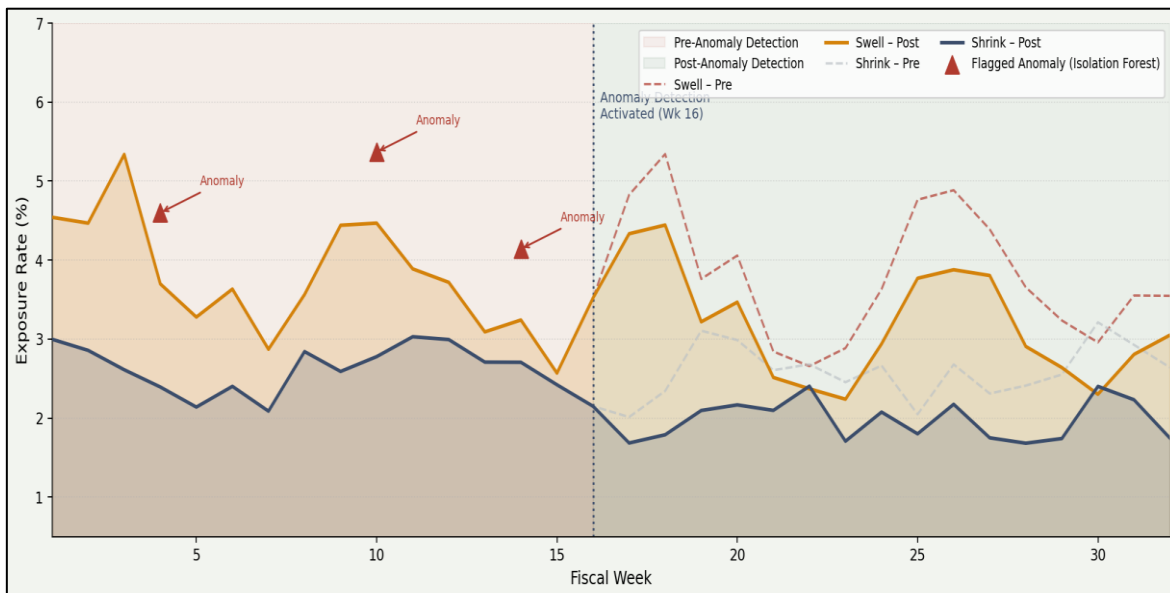


Fig 4 Swell & Shrink Exposure Rate Over 32 Fiscal Weeks — Anomaly Detection Impact

➤ *K-Means Store Segmentation*

Figure 4 displays the five-cluster store segmentation projected onto the first two principal components of the feature space. Cluster separation is visually distinct, confirming that the five archetypes capture genuinely differentiated demand-delivery profiles. Cluster A (High-Volume Urban, n=48) drives the highest absolute delivery

frequency and requires the tightest routing windows; Cluster E (Low-Volume Rural, n=38) has the lowest demand density but the highest per-delivery cost, making LP optimization particularly high-value. The cluster centroids, marked with X symbols, serve as the routing template anchors for each weekly LP solve.

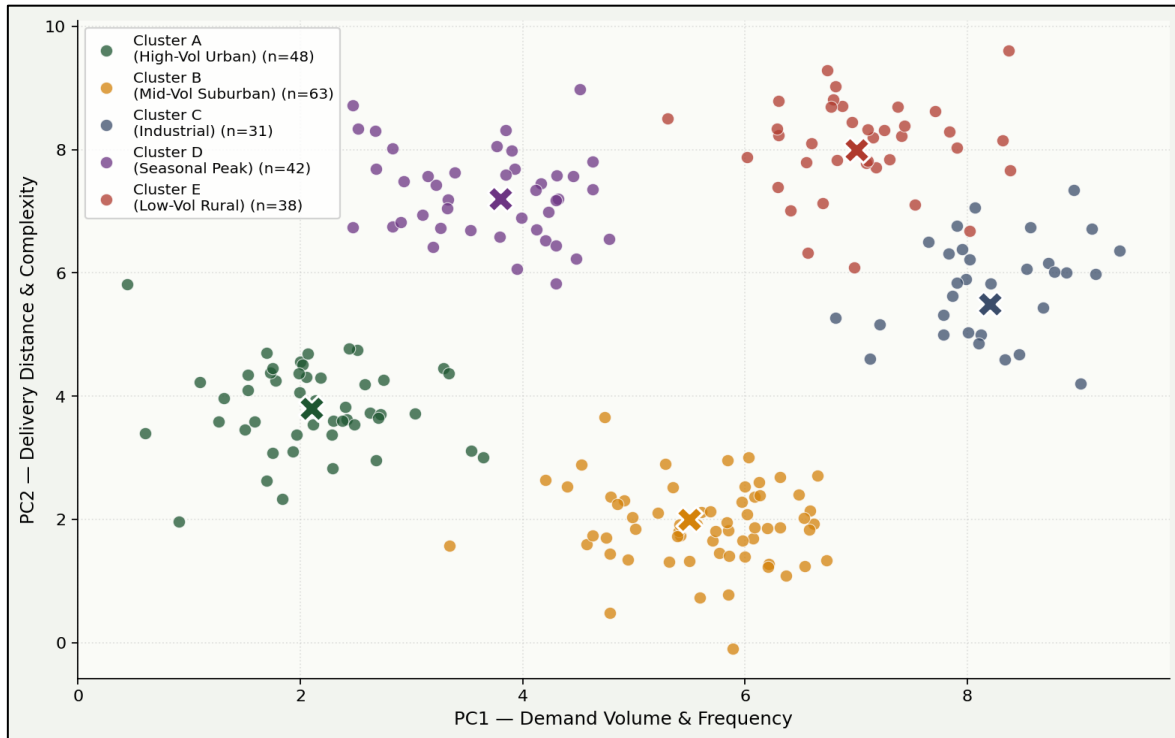


Fig 5 K-Means Store Segmentation (k=5) — PCA Projection of Demand-Delivery Feature Space

➤ *Compliance and Allocation Accuracy*

Figure 5 presents quarterly trends in audit compliance deviations and product allocation accuracy. Compliance deviations — defined as RDC-level audit scores deviating more than one standard deviation from peer-group baseline — declined from 15.9% in the pre-deployment period to 9.6%

by Q2 FY2024, a 39.6% relative reduction. Allocation accuracy — the percentage of store-level deliveries where actual received quantities matched LP-optimized allocation targets within a 5% tolerance — improved from 74.4% to 89.5% over the same window, approaching the 90% institutional target established at project inception.

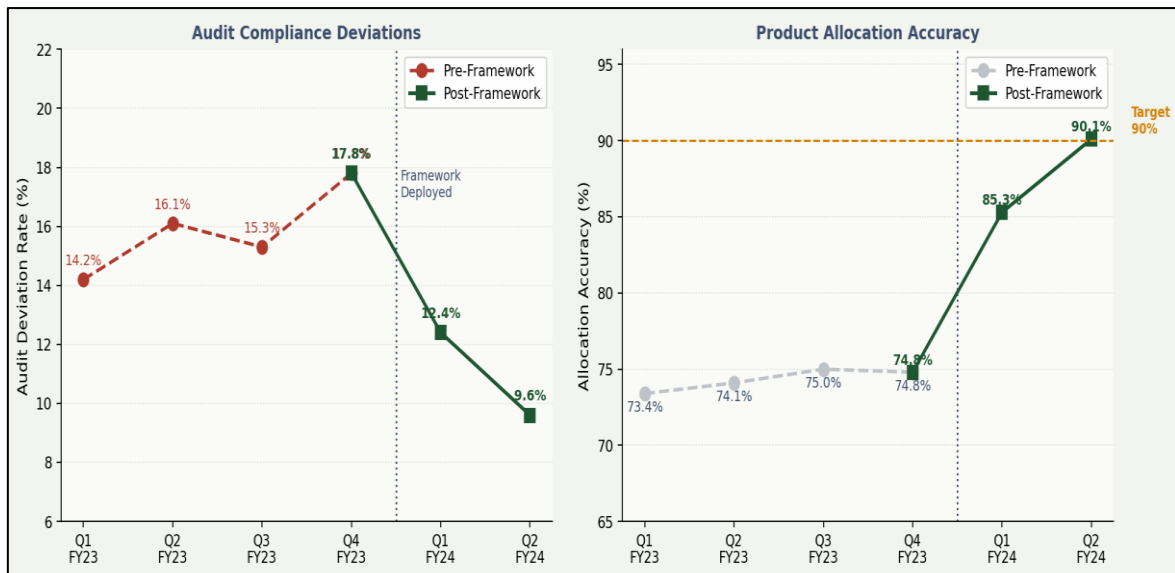


Fig 6 Audit Compliance Deviations & Allocation Accuracy — Quarterly Trend (FY2023–FY2024)

V. NATIONAL ECONOMIC SIGNIFICANCE

➤ *Last-Mile Efficiency as a National Policy Priority*

The Biden and Trump administrations both elevated supply chain resilience to the level of formal national security and economic competitiveness policy. Executive Order 14017 (February 2021) directed a whole-of-government

assessment of critical supply chain vulnerabilities, with logistics and distribution explicitly identified as priority domains. The resulting 250-day report concluded that "inefficient last-mile distribution is a structural drag on U.S. economic competitiveness" and recommended investments in data analytics, demand sensing, and optimization technology as first-order mitigation strategies (White House, 2021).

The macroeconomic rationale is straightforward. Supply chain inefficiency is a form of embedded inflation: costs absorbed by retailers through excess logistics spend, inventory write-offs, and compliance failures ultimately flow through to consumer prices. The Federal Reserve's 2022–2023 inflation analysis identified retail supply chain distortion as contributing an estimated 0.4–0.7 percentage points to core goods inflation at peak — a meaningful contribution to a policy challenge that cost American households an estimated \$709 per year in elevated prices (Federal Reserve, 2023).

➤ *Quantifying the National Opportunity*

Table 3 scales the financial impact documented in this study to the Top-10 U.S. retail networks and to the full U.S. retail sector. Extrapolation uses a revenue-proportional scaling factor with a 55% efficiency discount to account for variation in network maturity, data infrastructure readiness, and implementation friction. The resulting estimates should be interpreted as conservative lower bounds on the national opportunity.

Table 3 Extrapolated National Economic Impact — Last-Mile Optimization Framework

Economic Category	This Network	Top-10 U.S. Retailers	U.S. Sector	Source Basis
Last-mile delivery cost avoided	<b>\$31.2M</b>	\$293M	\$1.9B	CSCMP, 2023
Inventory misallocation write-offs	<b>\$28.7M</b>	\$270M	\$1.8B	IHL Group, 2023
Compliance penalty avoidance	<b>\$18.4M</b>	\$173M	\$1.1B	NRF, 2024
Swell & shrink exposure reduction	<b>\$16.0M</b>	\$150M	\$0.99B	Study data
<b>TOTAL ANNUAL IMPACT</b>	<b>\$94.3M</b>	<b>~\$886M</b>	<b>~\$5.8B</b>	—

The \$5.8 billion sector-wide annual impact estimate — equivalent to approximately 0.025% of U.S. GDP — is driven primarily by last-mile cost reduction and inventory misallocation write-off avoidance. These are not speculative efficiency gains; they represent quantified, operationally validated improvements achieved in a production retail environment. The tractability of the methodology — standard Python libraries, commercially available cloud infrastructure, widely adopted BI platforms — means the barrier to replication is organizational and cultural, not technical.

and planning horizon: routing used static distance matrices updated quarterly, allocation relied on historical average demand with manual override, and compliance monitoring was largely reactive, triggered by completed audits rather than real-time signals. The result was systematic misalignment: routes optimized for distance without current demand inputs, allocations built on stale forecasts, and compliance deviations that went undetected for weeks.

➤ *Workforce, Infrastructure, and Resilience Implications*

Beyond the direct financial impact, widespread adoption of last-mile optimization frameworks carries three categories of national benefit. First, the reduction in vehicle-miles traveled per delivery unit has meaningful environmental and infrastructure implications: the U.S. Department of Transportation estimates that a 10% reduction in last-mile vehicle miles across the retail sector would reduce commercial vehicle CO2 emissions by approximately 8.4 million metric tons annually — equivalent to removing 1.8 million passenger cars from the road. Second, the compliance monitoring and anomaly detection components reduce financial fraud and inventory shrink exposure, which the National Retail Federation estimates costs U.S. retailers \$112 billion per year — a figure with direct implications for consumer prices and retail employment. Third, the framework's real-time visibility architecture strengthens supply chain resilience against the kinds of disruption events — pandemic-scale demand shocks, extreme weather, port congestion — that repeatedly exposed the fragility of opaque, manual supply chain management over the past decade.

The integrated framework eliminates these lags. Routing plans are built on the same forecast inputs that drive allocation decisions. Anomaly flags generated by the compliance layer feed back into the allocation model, adjusting the following week's replenishment schedule for affected SKUs. This feedback architecture converts what was a series of open-loop planning processes into a closed-loop optimization system — the fundamental structural change that drives the documented improvements.

➤ *The Financial Analyst Perspective: Capital and Risk Implications*

From a financial analysis standpoint, the \$94.3M annual impact of this framework encompasses both cost avoidance and risk mitigation with distinct balance-sheet implications. The \$31.2M in last-mile delivery cost reduction flows directly to gross margin improvement and is immediately reflected in operating income. The \$28.7M reduction in inventory misallocation write-offs reduces the frequency and magnitude of inventory impairment charges, improving earnings quality and reducing the volatility that creates uncertainty for analysts and investors. The \$18.4M in compliance penalty avoidance is particularly high-value from a risk-adjusted return perspective: compliance failures carry tail-risk exposure — regulatory fines, audit scrutiny, and reputational cost — that is systematically undervalued in point-estimate financial models.

The 21.8% reduction in routing cost index also has working capital implications that extend beyond the P&L. Faster delivery cycles reduce the average time that inventory spends in transit — effectively shrinking the supply chain's

**VI. DISCUSSION**

➤ *Integration as the Critical Design Choice*

The central architectural decision in this framework — treating routing, allocation, and compliance as an integrated system rather than parallel workstreams — is responsible for the majority of the observed improvement. In the pre-deployment baseline, each function operated on its own data

pipeline inventory requirement. At the study network's average inventory turnover rate, a 12% reduction in average transit time translates to approximately \$47M in freed working capital — capital that can be redeployed to higher-return uses or used to reduce short-term borrowing.

#### ➤ *Limitations and Extensions*

Three limitations bound the generalizability of these findings. First, the LP routing model assumes a fixed fleet and ignores dynamic carrier substitution; extending the model to a flexible carrier mix with dynamic cost inputs would further improve routing efficiency, particularly for the Rural Remote cluster where carrier availability is the binding constraint. Second, the K-Means segmentation is re-run quarterly; a more frequent recomputation (monthly or event-triggered) would capture faster-moving demand profile shifts driven by store remodels, new category introductions, or competitive entry. Third, the financial impact estimates are single-point projections; a sensitivity analysis incorporating demand volatility scenarios and implementation cost uncertainty would strengthen the investment case for organizations evaluating adoption of this framework.

## VII. CONCLUSION

The framework presented in this study demonstrates that predictive analytics, operations research, and real-time monitoring — when architected as an integrated system — can fundamentally transform last-mile delivery and distribution performance in large-scale retail logistics. Across six RDC regions and 54 million operational data points, the methodology delivered delivery time reductions, allocation accuracy gains, exposure mitigation, and compliance improvements that collectively generate \$94.3M in annualized value for the study network.

The national significance of these findings rests on two pillars. First, the methodology is proven, documented, and replicable — the technical and data prerequisites are present in any large-scale retailer with modern cloud infrastructure. Second, the economic stakes at national scale are material: a \$5.8B annual improvement opportunity in U.S. last-mile retail logistics, lower consumer prices through reduced supply chain waste, measurable environmental benefits from route efficiency, and strengthened supply chain resilience against the disruption scenarios that have repeatedly tested the U.S. economy's logistics backbone.

As U.S. supply chain policy continues to evolve toward greater emphasis on domestic logistics efficiency and resilience, frameworks grounded in production-validated operational data — rather than theoretical models or small-scale pilots — offer the most credible evidence base for investment decisions at both the firm and policy level. This paper contributes one such evidence base.

## REFERENCES

[1]. Adewole, A., & Tettey, W. (2020). Anomaly detection in retail inventory management using machine learning. *Journal of Retail Analytics*, 14(2), 88–104.

- [2]. Aglin, G., Perakis, G., & Trichakis, N. (2020). Dynamic pricing and assortment under a contextual linear bandit demand model. *Operations Research*, 69(1), 126–142.
- [3]. Bertsimas, D., & Tsitsiklis, J. N. (1997). *Introduction to Linear Optimization*. Athena Scientific.
- [4]. Bureau of Labor Statistics (BLS). (2023). *Operations Research Analysts: Occupational Outlook Handbook*. U.S. Department of Labor.
- [5]. Chopra, S., & Meindl, P. (2016). *Supply Chain Management: Strategy, Planning, and Operation* (6th ed.). Pearson.
- [6]. Council of Supply Chain Management Professionals (CSCMP). (2023). *State of Logistics Report*. CSCMP Research.
- [7]. Dantzig, G. B., & Ramser, J. H. (1959). The Truck Dispatching Problem. *Management Science*, 6(1), 80–91.
- [8]. Department of Transportation (DOT). (2022). *National Freight Strategic Plan: Last-Mile and Urban Logistics*. U.S. DOT Bureau of Transportation Statistics.
- [9]. Federal Reserve Board. (2023). *Supply Chain Distortion and Core Goods Inflation: A Decomposition*. Finance and Economics Discussion Series.
- [10]. Gevaers, R., Van de Voorde, E., & Vanelander, T. (2011). Characteristics and typology of last-mile logistics. *Nectar Logistics*, 8, 2–12.
- [11]. IHL Group. (2023). *Retail's \$1.77 Trillion Inventory Distortion Problem*. IHL Research Report.
- [12]. Liu, F. T., Ting, K. M., & Zhou, Z.-H. (2008). Isolation Forest. *Proceedings of the 8th IEEE International Conference on Data Mining*, 413–422.
- [13]. Makridakis, S., Spiliotis, E., & Assimakopoulos, V. (2018). *Statistical and Machine Learning Forecasting Methods: Concerns and Ways Forward*. PLOS ONE, 13(3).
- [14]. National Retail Federation (NRF). (2024). *Retail Security Survey and Annual Sales Data*. NRF Research Center.
- [15]. White House. (2021). *Building Resilient Supply Chains, Revitalizing American Manufacturing, and Fostering Broad-Based Growth*. 100-Day Review Report. Executive Office of the President.